Maintaining and Troubleshooting Avaya Analytics™ for Avaya Oceana®

Release 4.3.1.1
Issue 1
March 2026

Notices

You may, at your own risk, assemble a MyDocs collection solely for your own internal business purposes, which constitutes a modification to the original published version of the publications. Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of publications. You agree to defend, indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, your modifications, additions or deletions to the publications.

A single topic or a collection of topics may come from multiple Avaya publications. All of the content in your collection is subject to the legal notices and disclaimers in the publications from which you assembled the collection. For information on licenses and license types, trademarks, and regulatory statements, see the original publications from which you copied the topics in your collection.

Except where expressly stated by Avaya otherwise, no use should be made of materials provided by Avaya on this site. All content on this site and the publications provided by Avaya including the selection, arrangement and design of the content is owned by Avaya and/or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. Avaya owns all right, title and interest to any modifications, additions or deletions to the content in the Avaya publications.

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

“Documentation” means information published in varying media which may include product information, subscription or service descriptions, operating instructions and performance specifications that are generally made available to users of products. Documentation does not include marketing materials. Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of Documentation unless such modifications, additions, or deletions were performed by or on the express behalf of Avaya. End user agrees to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End user.

Link disclaimer

Avaya is not responsible for the contents or reliability of any linked websites referenced within this site or Documentation provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Avaya provides a limited warranty on Avaya hardware and software. Please refer to your agreement with Avaya to establish the terms of the limited warranty. In addition, Avaya’s standard warranty language as well as information regarding support for this product while under warranty is available to Avaya customers and other parties through the Avaya Support website: https://support.avaya.com/helpcenter/getGenericDetails?detailId=C20091120112456651010 under the link “Warranty & Product Lifecycle” or such successor site as designated by Avaya. Please note that if the product(s) was purchased from an authorized Avaya channel partner outside of the United States and Canada, the warranty is provided by said Avaya Channel Partner and not by Avaya.

“**Hosted Service**” means an Avaya hosted service subscription that You acquire from either Avaya or an authorized Avaya Channel Partner (as applicable) and which is described further in Hosted SAS or other service description documentation regarding the applicable hosted service. If You purchase a Hosted Service subscription, the foregoing limited warranty may not apply but You may be entitled to support services in connection with the Hosted Service as described further in your service description documents for the applicable Hosted Service. Contact Avaya or Avaya Channel Partner (as applicable) for more information.

Hosted Service

THE FOLLOWING APPLIES ONLY IF YOU PURCHASE AN AVAYA HOSTED SERVICE SUBSCRIPTION FROM AVAYA OR AN AVAYA CHANNEL PARTNER (AS APPLICABLE), THE TERMS OF USE FOR HOSTED SERVICES ARE AVAILABLE ON THE AVAYA WEBSITE, HTTPS://SUPPORT.AVAYA.COM/LICENSEINFO UNDER THE LINK “Avaya Terms of Use for Hosted Services” OR SUCH SUCCESSOR SITE AS DESIGNATED BY AVAYA, AND ARE APPLICABLE TO ANYONE WHO ACCESSES OR USES THE HOSTED SERVICE. BY ACCESSING OR USING THE HOSTED SERVICE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE DOING SO (HEREINAFTER REFERRED TO INTERCHANGEABLY AS “YOU” AND “END USER”), AGREE TO THE TERMS OF USE. IF YOU ARE ACCEPTING THE TERMS OF USE ON BEHALF A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT THAT YOU HAVE THE AUTHORITY TO BIND SUCH ENTITY TO THESE TERMS OF USE. IF YOU DO NOT HAVE SUCH AUTHORITY, OR IF YOU DO NOT WISH TO ACCEPT THESE TERMS OF USE, YOU MUST NOT ACCESS OR USE THE HOSTED SERVICE OR AUTHORIZE ANYONE TO ACCESS OR USE THE HOSTED SERVICE.

Licenses

The Global Software License Terms (“Software License Terms”) are available on the following website https://www.avaya.com/en/legal-license-terms/ or any successor site as designated by Avaya. These Software License Terms are applicable to anyone who installs, downloads, and/or uses Software and/or Documentation. By installing, downloading or using the Software, or authorizing others to do so, the end user agrees that the Software License Terms create a binding contract between them and Avaya. In case the end user is accepting these Software License Terms on behalf of a company or other legal entity, the end user represents that it has the authority to bind such entity to these Software License Terms.

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, Hosted Service, or hardware provided by Avaya. All content on this site, the documentation, Hosted Service, and the product provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software unless expressly authorized by Avaya. Unauthorized reproduction, transmission, dissemination, storage, or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

Virtualization

The following applies if the product is deployed on a virtual machine. Each product has its own ordering code and license types. Unless otherwise stated, each Instance of a product must be separately licensed and ordered. For example, if the end user customer or Avaya Channel Partner would like to install two Instances of the same type of products, then two products of that type must be ordered.

Third Party Components

The following applies only if the H.264 (AVC) codec is distributed with the product. THIS PRODUCT IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO (i) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD (“AVC VIDEO”) AND/OR (ii) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE HTTP://WWW.MPEGLA.COM.

Service Provider

WITH RESPECT TO CODECS, IF THE AVAYA CHANNEL PARTNER IS HOSTING ANY PRODUCTS THAT USE OR EMBED THE H.264 CODEC OR H.265 CODEC, THE AVAYA CHANNEL PARTNER ACKNOWLEDGES AND AGREES THE AVAYA CHANNEL PARTNER IS RESPONSIBLE FOR ANY AND ALL RELATED FEES AND/OR ROYALTIES. THE H.264 (AVC) CODEC IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO: (i) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD (“AVC VIDEO”) AND/OR (ii) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION FOR H.264 (AVC) AND H 3 (HEVC) CODECS MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE HTTP://WWW.MPEGLA.COM.

Compliance with Laws

You acknowledge and agree that it is Your responsibility to comply with any applicable laws and regulations, including, but not limited to laws and regulations related to call recording, data privacy, intellectual property, trade secret, fraud, and music performance rights, in the country or territory where the Avaya product is used.

Preventing Toll Fraud

“Toll Fraud” is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud intervention

If You suspect that You are being victimized by Toll Fraud and You need technical assistance or support, please contact your Avaya Sales Representative.

Security Vulnerabilities

Information about Avaya’s security support policies can be found in the Security Policies and Support section of https://support.avaya.com/security.

Suspected Avaya product security vulnerabilities are handled per the Avaya Product Security Support Flow (https://support.avaya.com/css/P8/documents/100161515).

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support website: https://support.avaya.com, or such successor site as designated by Avaya.

Contact Avaya Support

See the Avaya Support website: https://support.avaya.com for Product or Cloud Service notices and articles, or to report a problem with your Avaya Product or Cloud Service. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: https://support.avaya.com (or such successor site as designated by Avaya), scroll to the bottom of the page, and select Contact Avaya Support.

Trademarks

The trademarks, logos and service marks (“Marks”) displayed in this site, the Documentation, Hosted Service(s), and product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, its licensors, its suppliers, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the Documentation, Hosted Service(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party.

Avaya is a registered trademark of Avaya LLC.

All non-Avaya trademarks are the property of their respective owners.

Contents

Chapter 1: Introduction 10
Purpose 10
New in this release 10
Support for Inactivity Timeout setting 10
Enhancement in Maintaining Avaya Analytics™ on Avaya Common Services chapter 11
Support for new Cluster Control Manager commands 11
Chapter 2: Maintaining Avaya Analytics™ using the Cluster Control Manager console 12
Avaya Analytics™ Maintenance 12
Setting group filters 12
Configuring Dimensions 14
Configuring Call Originator Redaction 15
Generating CSV files for Agent by Account and Routing Service reports 16
Configuring user management 19
Adding time zones 21
Avaya Analytics™ backups 22
Full backup on a remote server 22
Incremental backup 31
Restoring the database from remote backup 36
Restoring the database from incremental backup 37
Dropping and restoring database schemas 39
Viewing the list of data retention limits 39
Setting data retention limits 40
Configuring data roll-up 42
Configuring and deleting SNMP alarm destinations 43
Generating Management Information Base file for SNMP 45
Registering for certification expiration alarms 46
Configuring Geo database replication for alarming 47
Replacing the authorization certificates in the breeze-security secret 48
Manually replacing the authorization certificates in the breeze-security secret 49
Configuring Routing Service groups for Historical Reporting 52
Checking Routing Service Group feature status 52
Enabling or disabling Routing Service groups feature 52
Configuring System Routing Service group name 53
Historical Reporting user settings 54
Resetting Historical Reporting administrator password 55
Configuring DataWarehouse Password for Historical Reporting 56
Creating a local user for Historical Reporting 57
Adding a local user to a group 58

Removing a local user from a group ... 59
Resetting a Historical Reporting local user password ... 60
Deleting a Historical Reporting local user ... 61
Viewing list of Historical Reporting local users ... 62
Configuring Avaya Analytics™ email distribution services ... 62
Viewing Email Distribution Service configuration ... 63
Updating the governing rule for Historical Reporting ... 64
Viewing and configuring user idle timeouts ... 65
Configuring Avaya Analytics™ LDAP authentication ... 66
Configuring certificates for secure LDAP connection ... 68
Viewing LDAP settings ... 69
Configuring EASG availability for Historical Reporting ... 69
Enabling EASG availability for Historical Reporting ... 70
Disabling EASG availability for Historical Reporting ... 70
Checking EASG availability for Historical Reporting ... 71
Monitoring security certificates ... 71
Creating metadata backups ... 72
Viewing and removing metadata backups ... 74
Scheduling a full metadata backup ... 75
Package migration ... 78
Viewing a list of imported migration packages ... 78
Importing migration packages ... 78
Replacing migration package metadata ... 79
Rolling back an imported migration package ... 80
Checking disk space usage ... 81
Copying Historical Reporting logs to Cluster Control Manager ... 82
Deleting logs from Historical Reporting pods ... 83
Checking High Availability status ... 83
Configuring Okta as Identity Provider for Historical Reporting ... 84
Configuring Active Directory Federation Services as Identity Provider for Historical Reporting ... 86
Tracing of Agents for Historical Reporting ... 88
Viewing the verbose level set for Trace Processor ... 88
Configuring the verbose level for Trace Processor ... 89
Start an Agent Trace ... 89
Stop an Agent Trace ... 90
SAML implementation ... 90
Creating a new reporting user and map to SAML user for Historical Reporting ... 91
Mapping new SAML user in Historical Reporting ... 92
Configuring Web Single Sign-on for importing SAML users automatically into Historical Reporting ... 92
Enable debug level logging for troubleshooting SAML ... 94
Opening a node port on the Analytics database ... 95
Grafana reports for monitoring Crunchy database ... 95

Accessing Grafana reporting system..... 96
Enabling or disabling Zero or Empty Row Suppression parameter..... 97
Chapter 3: Maintaining Avaya Analytics™ on Avaya Common Sevices..... 98
Gracefully powering off a cluster..... 98
Stopping all events entering in to Avaya Analytics™ from Avaya Oceana®..... 101
Powering on a cluster..... 101
Monitoring the system status..... 103
Updating the network configuration of Cluster Control Manager and cluster nodes..... 106
Adding CPU and memory to a node..... 108
Reducing CPU and memory in a node..... 109
Increasing SDS disk size for a node..... 109
Restoring or replacing a cluster node..... 111
Adding a node to a cluster..... 112
Managing cluster VPN..... 112
Resize PVC using a spreadsheet..... 113
Resizing PVC using a spreadsheet..... 114
Registering a product with Avaya..... 115
Using a remote desktop session to launch the noVNC service..... 116
Resetting the remote desktop session password..... 117
Ending the remote desktop session..... 117
Accessing help for the remote desktop session..... 118
HTTP(S) outbound proxy configuration..... 118
Configuring proxy settings when deploying Cluster Control Manager..... 118
Using the ccmNetSetup command to configure proxy settings..... 119
Enabling or disabling file integrity validation..... 120
Creating index patterns on Kibana..... 121
Viewing AIDE logs in the Kibana logging interface..... 122
Linux audit rule updates..... 124
Security warning banner configuration..... 124
Backup and restore operations..... 124
Backing up Common Services..... 125
Restoring Avaya Common Services in an online deployment environment..... 127
Restoring Avaya Common Services in an offline air gap environment..... 129
Restoring Cluster Control Manager..... 132
Detaching from the restore SSH session..... 134
Reattaching to the restore SSH session..... 134
Deleting virtual machine remnants after restoring Avaya Common Services..... 135
Restoring or replacing a cluster node..... 135
Chapter 4: Troubleshooting Avaya Analytics™ from the Cluster Control Manager console..... 137
Troubleshooting Avaya Analytics™ from the Cluster Control Manager console..... 137
Capturing logs..... 137
Troubleshooting database issues..... 139

Restarting Historical Reporting 142
Unable to configure Historical Reporting with LDAP 142
The mstr_srv pod fails to start 143
The mstr_srv pod fails to get to a running 1/1 state 144
A Historical Reporting pod fails to start 146
Data does not display in Agent by Routing Service report after a node restarts 147
Multiple pods in Terminating state after node state changes to Not Ready 147
Troubleshooting general issues 148
Changing Programmatic Account Password 157
Restarting Avaya Analytics™ after an Avaya Oceana® restart 160
Backup files gets recreated after deleting from Cluster Control Manager 162
Error message seen in backup logs for Incremental backups 162
Unable to correctly restore backups after changing Postgres database password 163
Avaya Analytics DR Monitoring tool for replication and failover 163
Chapter 5: Troubleshooting third-party certificates issues 165
Specifying the identity certificates used for generating CSRs 165
Specifying trust stores for adding a third-party CA certificate 168
Ingress-gateway is not displayed in the list of identity certificates 171
List of trust stores in Common Service Platform 171
List of identity certificates in Common Service Platform 172
List of secrets 172
View the contents of a secret 172
Decode the content of a secret 172
View the contents of a PEM file 172
View the contents of a CSR file 173
Services fail to consume certificates renewed by Certificate Manager service 173
Chapter 6: Troubleshooting Avaya Analytics™ on Avaya Common Services 175
Cluster is in an unusable state (clients cannot connect to the cluster) 175
Cluster Failed to Install Due to Invalid Cluster Configuration 176
Service failure during an installation or upgrade 176
Upgrade process stalls 176
Terminal Shell Timed Out 177
DRS anti-affinity rule error during cluster installation 177
Creating a cluster node anti-affinity rule 178
Third-party Certificates Are Not Being Used 178
License node error message displayed at login 178
Reported Alarms Not Seen on NMS 180
Recovering a deleted virtual machine 180
Pods and services do not recover after a suspended VM is resumed 181
Pod crash alarms 181
Pod remains in init state 181
Logging in and out of Kibana 182
Error when logging in to Kibana 182

Restarting Kafka pod after multiple K8s node restarts ... 183
Cluster Control Manager commands ... 184
Cluster Control Manager core commands ... 184
ccm archive command options ... 192
ccm backup and ccm backup schedule command options ... 194
ccm restore command options ... 196
ccm report command ... 197
ccm report log file operations ... 202
Certificate Manager commands ... 205
Internal cluster certificate commands ... 207
Common Services commands ... 209
Eventing commands ... 210
Miscellaneous commands ... 210
ccm swhistory command ... 215
pvcCleanup command ... 217
ccm resizePVC command ... 218
clusterVPN commands ... 219
Command help output examples ... 220
Chapter 7: Troubleshooting Avaya Analytics™ web issues ... 228
Unable to access the Avaya Analytics™ webpage ... 228
Unable to log in to the Avaya Analytics™ Historical Reporting ... 228
Recover / Reset default Historical Reporting Administrator account password ... 229
Unable to log in to Avaya Analytics™ using LDAP ... 229
Unable to run an Avaya Analytics™ historical report with selected interval range ... 230
Routing Services no longer appear on the Routing Service Group ... 230
Voice channel activity does not increase on the Routing Service Summary report ... 230
Deleted agent groups data appears in historical reports that supervisors generate ... 231
MSTR-SRV pods stuck in PodInitializing status ... 231
Copying analyticsdb-node secrets into mstr namespace ... 232
Chapter 8: Troubleshooting Avaya Analytics™ migration issues ... 233
Migration from same source to other target error ... 233
Chapter 9: Troubleshooting Messaging ... 234
Verifying the status of the Messaging channel ... 234
Contact not getting created in Avaya Oceana® or messages not flowing into Avaya Oceana® ... 235
Async messages not getting to the customer ... 236
Chapter 10: Troubleshooting Avaya Analytics™ upgrade issues ... 238
Historical and real-time reporting fails after an upgrade ... 238
Upgrade roll back ... 239
Rolling back the Avaya Analytics™ software ... 239
Rolling back the custom reports ... 240
Rolling back to the previous version of the database ... 241
Avaya Common Services upgrade revert ... 242
Reverting a failed upgrade to the previous release ... 242

Unable to view the full list of routing services after pumpup 243
Chapter 11: Troubleshooting Avaya Analytics™ installation and post install script 244
Troubleshooting Avaya Analytics™ installation errors 244
Avaya Workspaces real-time dashboard does not update after installation 245
The ccm release orca analytics command fails 245
Troubleshooting Analytics online and offline deployment issues 246
Use of unknown CA certificate 246
Unauthorized or failed logins 246
Avaya Analytics™ chart and images fail to download 246
Avaya Analytics™ chart and images fail to upload 247
In real-time reporting the Routing Service Group producer is not visible 247
Historical data not available in the database and on Agent Trace reports in MSTR 248
Error mapping user for SAML 248
Avaya Workspaces does not load real-time data 249
Chapter 12: Troubleshooting common issues 250
Macros of deployment spreadsheet are disabled by administrator 250
Issues in localized spreadsheet 250
Error getting initialization data for producer: undefined 251
Restarting Avaya Analytics™ after applying Avaya Oceana® patches 251
Realtime reports and Historical reports are inaccessible after node outage 252
Deleting ORCA product results in PV and data deletion 253
Restrictions to vCenter user account password used for Avaya Analytics™ 253
Chapter 13: Resources 254
Documentation 254
Finding documents on the Avaya Support website 255
Avaya Documentation Center navigation 256
Training 257
Viewing Avaya Mentor videos 260
Support 261
Using the Avaya InSite Knowledge Base 261

Chapter 1: Introduction

Purpose

This document contains procedures for the routine maintenance and troubleshooting of Avaya Analytics™. Routine maintenance practices include scheduling backup and restoration, post-installation configuration, daily monitoring, and verification testing. It also contains regulatory information, safety precautions, configuration, and administration best practices.

The troubleshooting information in this document is not intended to replace or replicate the dynamic information stored in a product or solution knowledge base that the Avaya Support team maintains. Instead, it provides procedures on how to overcome and resolve common issues. It also describes troubleshooting tools and techniques.

This document is intended for people who perform Avaya Analytics™ maintenance and troubleshooting tasks, are familiar with the solution, and who are trained to handle software errors. To handle issues that are not documented in this document, contact Avaya support.

New in this release

Avaya Analytics™ Release 4.3 contains the following features and enhancements:

Related links

[Enhancement in Maintaining Avaya Analytics on Avaya Common Services chapter](#) on page 11
[Support for new Cluster Control Manager commands](#) on page 11

Support for Inactivity Timeout setting

Inactivity Timeout Setting in Avaya Analytics™ offers two timeout settings. The user session Idle time is a configurable period in seconds that users can remain idle. After that time expires, the user is locked out of the session and has to log back into Avaya Analytics™ Historical to continue their session. The user session idle time defaults to 1,800 seconds (30 mins). The web session is idle and expires when the idle time exceeds. The user must log in with their credentials if the web session expires.

Enhancement in Maintaining Avaya Analytics™ on Avaya Common Services chapter

There are enhancements in procedures under chapter *Maintaining Avaya Analytics™ on Avaya Common Services*. Backup and restore Avaya Common Services procedures are also updated.

Related links

New in this release on page 10

Support for new Cluster Control Manager commands

Following new Cluster Control Manager commands are supported in Avaya Analytics™ 4.3:

- clusterVPN commands
- ccm resizePVC command
- Resizing a PVC

Related links

New in this release on page 10

Chapter 2: Maintaining Avaya Analytics™ using the Cluster Control Manager console

Avaya Analytics™ Maintenance

When you install Avaya Analytics™, the installation configures the default settings automatically. To change these default settings, use the procedures in this chapter.

To complete these procedures, you must log in to the Cluster Control Manager (CCM) console as a customer user. The customer username varies for every organization. You must then switch to the root user. Root user access might not be required for some steps as indicated in the respective procedures.

To complete some of these procedures, you require access to:

- Cluster Control Manager IP address
- System Manager IP address
- Authorization Service Avaya Breeze® node SIPs
- Avaya Oceana® Cluster 1 Avaya Breeze® Node SIPs
- Avaya Control Manager

Note:

The post-install script might take some time to display the options after you run the **`ccm release orca analytics`** command or might display the following message: `Failed to get auth token`. Wait for some time or type the command again.

Setting group filters

About this task

You must set group filters for supervisors to monitor their respective assigned agents. By default, the agent group filter is enabled, and the routing service group filter is disabled.

Note:

When you enable or disable a routing service group filter, the supervisor experiences a disconnection in Avaya Workspaces on the real-time reports because the data publisher gets restarted.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Deployment** by pressing the corresponding number.
5. To select the **Realtime Configuration** option, enter the corresponding number.
6. To select the **Realtime Group Filters** option, enter the corresponding number.
7. To view the existing status of the agent group filter, select the **Show current Group Filter** option, enter the corresponding number.

   Depending on whether the group filter is enabled or disabled, the Group Filters screen displays the following:

   - `Group Filter is currently enabled.`
   - `Group Filter is currently disabled.`
8. Return to the previous page by entering `b`.
9. Quit the current page by entering `q`.
10. Return to the main menu by entering `m`.
11. To view the existing routing service group filter, select the **Show current Routing Service Group Filter** option, enter the corresponding number.

    Depending on whether the group filter is enabled or disabled, the Group Filters screen displays the following:

    - `Routing Service Group Filter is currently enabled.`
    - `Routing Service Group Filter is currently disabled.`
12. Return to the previous page by entering `b`.
13. Quit the current page by entering `q`.
14. Return to the main menu by entering `m`.
15. To set an agent group filter, select the **Set Group Filter** option, enter the corresponding number.
16. In the **Proceed to setting group filter value** field, type `y`.

    Entering `n` cancels the operation.

17. In the **Set Group Filter** field, type `enable` or `disable` and press **Enter**.
18. Return to the previous page by entering `b`.
19. Quit the current page by entering `q`.
20. Return to the main menu by entering `m`.
21. To set the routing service group filter, select the **Set Routing Service Group Filter** option, enter the corresponding number.
22. In the **Proceed with setting routing service group filter** field, type `y`.

    Entering `n` cancels the operation.
23. In the **Set Routing Service Group Filter** field, type `enable` or `disable` and press **Enter**.
24. Return to the previous page by entering `b`.
25. Quit the current page by entering `q`.
26. Return to the main menu by entering `m`.

Configuring Dimensions

About this task

To see all configured dimensions in real-time reports regardless of their activity, you can enable the display of all dimensions (including inactive dimensions). By default, the display of inactive dimensions is disabled.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Deployment** by pressing the corresponding number.
5. To select the **Realtime Configuration** option, enter the corresponding number.
6. To select the **Dimension Configuration** option, enter the corresponding number.
7. To view the existing status of **Configured dimensions**, select the **Show current settings** option, enter the corresponding number.
8. To enable the display of all dimensions (including inactive dimensions), enter the corresponding number.

Warning:

It is not recommended to enable the feature at peak times because it triggers the restart of admin-data-service and requests for UCA pump up events.

9. In the **Proceed to enable display of all dimensions (including inactive dimensions)? [y/n]** field, type `y` to confirm.
10. In the **Please create a list of the services you wish to enable display of all dimensions (including inactive dimensions)** option, select the required services and type 99.
11. To disable the display of inactive dimensions, enter the corresponding number.
12. In the **Proceed to disable display of inactive dimensions? [y/n]** field, type `y` to confirm.
13. In the **Please create a list of the services you wish to disable display of inactive dimensions** option, select the required services and type 99.
14. Return to the previous page by entering `b`.
15. Quit the current page by entering `q`.
16. Return to the main menu by entering `m`.

Configuring Call Originator Redaction

About this task

You must configure Call Originator Redaction to protect customer data stored in the database and shown in Realtime reports. When you enable this feature, Avaya Analytics™ redacts the originating caller number in your reports.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`

4. Select **Deployment** by pressing the corresponding number.
5. To select the **GDPR** option, enter the corresponding number.
6. To view the current settings for the Call Originator Redaction feature, select the **Show current GDPR configuration** option, enter the corresponding number.

   Depending on whether the Call Originator Redaction feature is enabled or disabled, the screen displays the following:

   - GDPR is currently enabled.

- GDPR is currently disabled

7. Return to the previous page by entering `b`.
8. Quit the current page by entering `q`.
9. Return to the main menu by entering `m`.
10. To enable the Call Originator Redaction feature, select the **Enable GDPR configuration option**, enter the corresponding number.
11. In the **Proceed with disabling GDPR for Historical and Realtime reporting value** field, type `Y` or `N` to cancel the operation.

    Warning:

    Setting the GDPR value requires CDR pod(s) restart.

12. To disable the Call Originator Redaction feature, select the **Disable GDPR configuration** option, enter the corresponding number.
13. In the **Proceed with disabling GDPR for Historical and Realtime realtime reporting value** field, type `y`.Entering `n` cancels the operation.

    Warning:

    Setting the GDPR value requires CDR pod(s) restart.

14. Return to the previous page by entering `b`.
15. Quit the current page by entering `q`.
16. Return to the main menu by entering `m`.

Generating CSV files for Agent by Account and Routing Service reports

About this task

To share historical reporting data with Avaya Experience Platform™ (Public Cloud Workforce Engagement) and non-Avaya Experience Platform™ (Public Cloud Workforce Engagement) solutions, you can configure Avaya Analytics™ to generate `CSV` files for the Agent By Account and Routing Service historical reports. Perform the following procedure to enable `CSV` file generation.

Warning:

Complete this procedure only if your solution includes Avaya Experience Platform™ (Public Cloud Workforce Engagement). If you enable `CSV` file generation when Avaya Experience Platform™ (Public Cloud Workforce Engagement) is not part of your solution, this can impact Avaya Analytics™ performance.

After you complete this procedure, Avaya Analytics™ generates the `CSV` files at the following locations:

- Routing Service
  - The location you configured on the dedicated external host server
- Agent By Account
  - The location you configured on the dedicated external host server

For Agent by Account, you can enable daily OR interval CSV producer. For the Routing Service processor, you can enable only interval CSV producer. Interval Producer puts the data into the CSV file at an interval of 15 minutes. The daily producer generates a new CSV file with a unique name once daily.

Before you begin

Ensure that the location to store the `CSV` files already exists on the host server and the specified user has write permissions to the directory.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Deployment** by pressing the corresponding number.
5. To select the **CSV Producer** option, enter the corresponding number.
6. To view the current CSV producer configuration, select the **Show CSV Producer settings** option, enter the corresponding number.

   Depending on the CSV producer configuration status, the CCM console displays the following:

   - If both the CSV producers are disabled:

     `CSV producer for Agent by Account processor is currently disabled.`

     `CSV producer for Routing Service processor is currently disabled.`
   - If one or both of the CSV producers are enabled, the CCM console displays settings that were configured for this producer. For example:

     ```
     oceanalytics.csv-producer.external-storage.host=[host]
     oceanalytics.csv-producer.external-storage.port=[port]
     oceanalytics.csv-producer.external-storage.username=[username]
     oceanalytics.csv-producer.external-storage.password=[password]
     oceanalytics.csv-producer.external-storage.path=[path]
     ```
7. Return to the previous page by entering `b`.
8. Quit the current page by entering `q`.

9. Return to the main menu by entering `m`.
10. To configure daily or interval CSV producer for Agent by Account processor, select the **Configure CSV producer for Agent by Account processor** option, enter the corresponding number.
11. Select one of the following:
    - To configure interval CSV producer for Agent by Account processor, select **Configure interval CSV producer for Agent by Account processor** option, enter the corresponding number.
    - To configure daily CSV producer for Agent by Account processor, select **Configure daily CSV producer for Agent by Account processor** option, enter the corresponding number.

    Note:

    You can enable the interval or the daily CSV producer at a time. When you enable one, the other producer gets disabled if it is currently enabled. After you successfully enable the CSV Producer, the configuration maps are updated without restart of the agentbyaccount pods.
12. In the **Enter the host IP address** field, type the IP address of the dedicated host server where the scripts write the `CSV` files and press **Enter**.
13. In the **Enter the port number used for connecting to the host** field, type the port number used for connecting to the host and press **Enter**.
14. In the **Enter the username for connecting to the host** field, type the name of the specified user with the write permissions to the directory and press **Enter**.
15. In the **Enter the users password** field, type the password of the specified user and press **Enter**.
16. In the **Enter the path** field, enter the path to store the `CSV` files.

    Note:

    Ensure that the path you enter to store the files must already exist on the host server, and the specified user must have the write permissions to the directory.

    After you successfully enable the CSV Producer, the configuration maps are updated without restarting the `agentbyaccount` pods.
17. Return to the previous page by entering `b`.
18. Quit the current page by entering `q`.
19. Return to the main menu by entering `m`.
20. To verify the CSV Producer settings status after enabling it, select the **Show CSV Producer settings** option, enter the corresponding number.
21. To disable CSV producer for Agent by Account processor, select the **Disable CSV producer for Agent by Account processor** option, enter the corresponding number after

step 10. The configuration maps are updated without restart of the `agentbyaccount` pods.

22. To configure interval CSV producer for Routing Service processor, select the **Configure CSV producer for Routing Service processor** option, enter the corresponding number.
23. To enable interval CSV producer for Routing Service processor or to change CSV producer settings for Routing Service processor, select the **Configure interval CSV producer for Routing Service processor** option, enter the corresponding number. Then, complete steps 12 through 16.
24. To disable interval CSV producer for Routing Service processor, select the **Disable CSV producer for Routing Service processor** option, enter the corresponding number.

    The configuration maps are updated without restart of the `routingservice` pods.
25. Return to the previous page by entering `b`.
26. Quit the current page by entering `q`.
27. Return to the main menu by entering `m`.

Configuring user management

About this task

You can use the user management option in the Analytics administration script to manage your database users. You can also create and remove a read-only database user using this option.

For example, you can create a read-only database user to connect to the database and view database tables, fields, or values.

Before you begin

Open the database port to connect to the database using the database administration tool.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. Select the **User Management** option, enter the corresponding number.
6. To create a read-only database user, select the **Create read-only database user** option by entering the corresponding number.
7. In the **Proceed to Read-only user creation** option, type `y`.

Entering `n` cancels the operation.

8. In the **Username** field, type the user name for the read-only database user and press **Enter**.
9. In the **Password** field, type the password for the read-only database user and press **Enter**.
10. In the **Re-enter password** field, type the password again for the read-only database user and press **Enter**.

    The CCM console displays the message that the account is created.
11. Return to the previous page by entering `b`.
12. Quit the current page by entering `q`.
13. Return to the main menu by entering `m`.
14. To view a list of the read-only database users, select the **List all read-only database users** option by entering the corresponding number.

    The page displays the list of database users.
15. Return to the previous page by entering `b`.
16. Quit the current page by entering `q`.
17. Return to the main menu by entering `m`.
18. To change the password for a read-only database user, select the **Change the password for a read-only database user** option by entering the corresponding number.
19. In the **Proceed to password change** field, type `y`.

    Typing `n` cancels the operation.
20. At the prompt, enter the username of the user for whom you are changing the password.
21. Enter and confirm the new password for the new read-only database user.

    The CCM console displays the message that the password is updated.
22. Return to the previous page by entering `b`.
23. Quit the current page by entering `q`.
24. Return to the main menu by entering `m`.
25. To remove a read-only database user, select the **Remove a read-only database user** option by entering the corresponding number.
26. In the **Proceed to Read-only user removal** field, type `y`.

    Entering `n` cancels the operation.
27. At the prompt, enter the username for the user you want to delete.

    This script deletes the account of the specified read-only database user.
28. Return to the previous page by entering `b`.

29. Quit the current page by entering `q`.
30. Return to the main menu by entering `m`.

Adding time zones

About this task

View the time zones that the Avaya Analytics™ database supports, and set the time zones required for your Avaya Analytics™ solution.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. To select the **Time Zones** option, enter the corresponding number.
6. In the **Proceed to Timezone search** field, enter `y`.

   Entering `n` cancels the operation.
7. In the **Time Zone** field, enter the required timezone that you entered when adding a time zone.

   The value is case-sensitive.
8. Return to the previous page by entering `b`.
9. Quit the current page by entering `q`.
10. Return to the main menu by entering `m`.
11. To select the **Add a new Time Zone** option, enter the corresponding number.
12. In the **Proceed with adding Timezone** field, type `y`.

    Entering `n` cancels the operation.
13. Enter the time zone that you want to add.

    The time zone must match one of the existing time zones in the list of supported time zones. This value is case-sensitive.
14. Return to the previous page by entering `b`.
15. Quit the current page by entering `q`.
16. Return to the main menu by entering `m`.

17. To select the **List current Time Zones** option, enter the corresponding number.

    The CCM console displays a list of the time zones that are set currently.
18. Return to the previous page by entering `b`.
19. Quit the current page by entering `q`.
20. Return to the main menu by entering `m`.
21. To select the **Remove a Time Zone** option, enter the corresponding number.
22. In the **Proceed to Timezone removal** field, type `y`.

    Entering `n` cancels the operation.
23. In the **Time Zone** field, enter the timezone that you want to remove.

    This value must match the `timezone_name` column from the existing timezones list. This value is case-sensitive.
24. Return to the previous page by entering `b`.
25. Quit the current page by entering `q`.
26. Return to the main menu by entering `m`.

Avaya Analytics™ backups

You must configure daily and weekly backups of Historical databases and ensure that the output files are stored securely.

Note:

Full backup on local server is not supported.

Related links

Restoring the database from remote backup on page 36
Restoring the database from incremental backup on page 37

Full backup on a remote server

Managing remote server connection settings

About this task

You can configure, view, or remove the remote server connection settings used to take a full backup of the Avaya Analytics™ database.

Cluster Control Manager archive destination uses the remote server details for backups. Changing the remote server here overwrites it for Cluster Control Manager archive Destination.

Before you begin

SFTP is used to export backups and requires:

- The external server to be accessible from the cluster
- A user login credentials for the external server
- Path to the location on the external server to store the backups
- Read and write privileges for the user to the path location
- Use of Port 22

Note:

Avaya Analytics™ does not manage the destination server. Therefore, Avaya Analytics™ does not enforce any retention policy for backups on this remote server.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. To select the **Database backup** option, enter the corresponding number.
6. To select the **Remote backups** option, enter the corresponding number.
7. To configure the remote server connection settings, enter the corresponding number to select the **Update remote server connection settings** option.
8. Enter the following settings of the remote server as required:
   a. Enter **Remote Destination IP/FQDN** of the remote server.

      The default port is 22.
   b. Enter **Remote Destination username login** of the remote server. You must have read and write permission on the remote server.
   c. Enter **Remote Destination path**, the directory path on remote server, where the database is exported. For example: `/temp/backups`.

      Note:

      The directory path must already exist on the remote server.
   d. In the **Do you wish to proceed?[y/n]:** field, type `y` to confirm the settings.
   e. Enter **Remote User Password** of the remote server
   f. **Re-enter the Remote User Password** of the remote server.

      The remote server connection settings are updated.

9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.
12. To view the existing remote server settings, enter the corresponding number to select the **Show remote server connect settings** option.

    The remote server settings are displayed.
13. Return to the previous page by entering `b`.
14. Quit the current page by entering `q`.
15. Return to the main menu by entering `m`.
16. To remove the existing remote server settings, enter the corresponding number to select the **Remove remote server connection settings** option.
17. In the **The Remote Destination configuration will be erased. Do you wish to proceed? [y/n]:** field, type `y`.

    Remote server connection settings are removed.
18. Return to the previous page by entering `b`.
19. Quit the current page by entering `q`.
20. Return to the main menu by entering `m`.

Related links

Avaya Analytics backups on page 22

Configuring remote backup

About this task

You can run an immediate full backup of the Avaya Analytics™ crunchy Postgres database to a remote server. The remote server can be Windows or Linux server, and is located outside of the cluster. Use pgdump tool to take the full database backup.

Avaya recommends to use remote backups as these are stored outside the Avaya Analytics™ Cluster to help avoid loosing backups.

For scheduled backups, see the *Managing scheduled backups* section in this document.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. To select the **Database backup** option, enter the corresponding number.

6. To export the full backup to a remote server outside of CCM, do the following to configure the remote server:
   a. To select the **Update Remote Server connection settings** option, enter the corresponding number.
   b. In the **Remote Destination IP/FQDN** field, enter the IP of the remote server.

      For example: **Remote Destination IP/FQDN**: 10.10.10.123
   c. Enter **Remote Destination username login** of the remote server. You must have read and write permission on the remote server.
   d. Enter **Remote Destination path**, the directory path on the remote server, where the database is exported. For example: `/temp/backups`.

      For example, to add Linux server path, enter: `/temp/backups`

      To add Windows server path, enter: `/home/myuser/analytics_backups`
   e. In the **Do you wish to proceed?[y/n]:** field, type `y` to confirm the settings.
   f. Enter **Remote User Password** of the remote server
   g. **Re-enter the Remote User Password** of the remote server.

      The Remote server is configured.
7. Return to the previous page by entering `b`.
8. To take a backup of the database, enter the corresponding number to select the **Remote backup** option.
9. To select the **Run a Remote backup** option, enter the corresponding number.
10. In the **Proceed to running a remote backup** option, type `y`.

    Entering `n` cancels the operation.
11. In the **Confirm server details [y/n]:** field, type `y`.

    The full database backup to the remote server is completed. The backup is first stored on the database PV. It is compressed and then transferred to the remote server in `.gz` file format. After the backup is moved to remote server, the PV is cleaned, and no space is occupied on the cluster.
12. Return to the previous page by entering `b`.
13. Quit the current page by entering `q`.
14. Return to the main menu by entering `m`.

Related links

Avaya Analytics backups on page 22

Scheduling a full database backup on remote server

About this task

You can schedule a full database backup on a remote server based on your choice of time, such as daily, weekly, monthly, and yearly. You can also choose the hour for the backup.

Before you begin

You must configure the remote server. See the *Managing Remote Server connection settings* section in this document.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. To select the **Database backup** option, enter the corresponding number.
6. To run a remote backup, enter the corresponding number to select the **Remote Backups** option.
7. If the remote server is not configured, then enter the corresponding number and configure the remote server to select **Update Remote Server connection settings**.
8. To select the **Update/Create a schedule for Remote Backups** option, enter the corresponding number.
9. In the **Proceed with remote scheduled backup update** field, type `y`.

   Entering `n` cancels the operation.
10. In the **How often should the backup be executed [daily/weekly/monthly/yearly]** field, type the required duration.

    Depending on your selection, follow the next steps:

For daily backups:

11. In the **How often should the backup be executed [daily/weekly/monthly/yearly]** field, type `daily` and press **Enter**.

    The value is case-sensitive.
12. To select the time in hour for the backup, in the **What hour of the day [0–23]** field, type the time and press **Enter**.

    The time selection is in 24–hour format. For example, if you type 2, then the backup happens at 2 am.
13. To select the time in minute for the backup, in the **What minute of the hour [0–59]** field, type the time and press **Enter**.

For example, if you type 2, then the backup happens at 2 minutes past 2 am.

14. Return to the previous page by entering `b`.
15. Quit the current page by entering `q`.
16. Return to the main menu by entering `m`.

For weekly backups:

17. In the **How often should the backup be executed [daily/weekly/monthly/yearly]** field, type `weekly` and press **Enter**.

    The value is case-sensitive.

18. To select the day of the week for the backup, in the **What day of the week (Sunday to Saturday) [0–6]** field, type the number of the day and press **Enter**.

    For example, to select Monday, type `1`.

19. To select the time in hour for the backup, in the **What hour of the day [0–23]** field, type the time and press **Enter**.

    The time selection is in 24-hour format. For example, if you type 2, then the backup happens at 2 am.

20. To select the time in minute for the backup, in the **What minute of the hour** field, type the time and press **Enter**.

    For example, if you type 2, then the backup happens at 2 minutes past 2 am.

21. Return to the previous page by entering `b`.
22. Quit the current page by entering `q`.
23. Return to the main menu by entering `m`.

For monthly backups:

24. In the **How often should the backup be executed [daily/weekly/monthly/yearly]** field, type `monthly` and press **Enter**.

    The value is case-sensitive.

25. To select the day for the backup, in the **What day of the month [1–31]** field, type the number of the day and press **Enter**.

    For example, to select 2nd day of every month, type `2`

26. To select the time in hour for the backup, in the **What hour of the day** field, type the time and press **Enter**.

    The time selection is in 24-hour format. For example, if you type 2, then the backup happens at 2 am.

27. To select the time in minute for the backup, in the **What minute of the hour** field, type the time and press **Enter**.

    For example, if you type 2, then the backup happens at 2 minutes past 2 am.

28. Return to the previous page by entering `b`.
29. Quit the current page by entering `q`.
30. Return to the main menu by entering `m`.

For yearly backups:

31. In the **How often should the backup be executed [daily/weekly/monthly/yearly]** field, type `yearly` and press **Enter**.

    The value is case-sensitive.

32. To select the month for the backup, in the **What month of the year (January to December) [1–12]** field, type the number corresponding to the month and press **Enter**.

    For example, for selecting October, type `10`.

33. To select the day for the backup, in the **What day of the month [1–31]** field, type the number of the day and press **Enter**.

    For example, to select 2nd day of every month, type `2`

34. To select the time in hour for the backup, in the **What hour of the day** field, type the time and press **Enter**.

    The time selection is in 24-hour format. For example, if you type 2, then the backup happens at 2 am.

35. To select the time in minute for the backup, in the **What minute of the hour** field, type the time and press **Enter**.

    For example, if you type 2, then the backup happens at 2 minutes past 2 am.

36. Return to the previous page by entering `b`.
37. Quit the current page by entering `q`.
38. Return to the main menu by entering `m`.

Related links

Avaya Analytics backups on page 22

Managing scheduled remote backups

About this task

Use this procedure to manage your scheduled remote backups, such as deleting and viewing the list of scheduled backups.

To schedule a full backup on the remote server, see the *Scheduling a full database backup on remote server* section in this document.

Before you begin

Stop all events coming to the Avaya Analytics™ database from Avaya Oceana®. For steps on stopping the events, see Stopping all events entering in to Avaya Analytics from Avaya Oceana on page 101.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. To select the **Database backup** option, enter the corresponding number.
6. To select the **Remote backups** option, enter the corresponding number.
7. To view the list of currently scheduled full backups, enter the corresponding number to select the **List schedule for running Remote Backups** option.
8. Return to the previous page by entering `b`.
9. Quit the current page by entering `q`.
10. Return to the main menu by entering `m`.
11. To delete a scheduled backup, enter the corresponding number to select the **Delete scheduled for running remote backups** option.

    The Delete scheduled of the remote backup page displays the option to enter the backup details that you want to remove.
12. In the **Select which schedule you would like to delete** field, enter the number corresponding to the scheduled backup that you want to remove from the list.

    Note:

    When you upgrade Avaya Analytics™ to a new version, all the existing backup schedules are deleted automatically. After the upgrade, you need to create new schedules for database backup.
13. Return to the previous page by entering `b`.
14. Quit the current page by entering `q`.
15. Return to the main menu by entering `m`.

Related links

Avaya Analytics backups on page 22

Restoring the database from remote backup

About this task

You can restore a full backup of the database from the remote server using the Avaya Analytics™ Administration script.

Note:

Restoring backup is only supported when the backup is taken on the same Avaya Analytics™ release as a target for the restore. For example, a backup taken on Avaya Analytics™ 4.1.1.0 can only be restored on the Avaya Analytics™ 4.1.1.0.

Before you begin

Warning:

You must restore backups only during maintenance window periods. To prevent data loss, ensure that no other activities occur in the database during the backup process.

Do not restore a database backup of the previous Avaya Analytics™ release version.

- Stop all events coming to the Avaya Analytics™ database from Avaya Oceana® and stop interval data being written to the Avaya Analytics™ database. Run the following commands:
  - To stop data flow to the cluster, run the following command:

    ```
    kubectl scale --replicas=0 deployment orca-ref-input-adaptor
    ```
  - To stop interval data, run the following command:

    ```
    kubectl scale --replicas=0 deployment orca-interval-controller
    ```
- Check your backup file permissions.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. To select the **Database Restore** option, enter the corresponding number.
6. To select the **Restore from remote backup** option, enter the corresponding number.
7. In the **Proceed to restore a remote backup** field, type `y`.

   Entering `n` cancels the operation.
8. The details of the remote server is displayed. In **Confirm server details [y/n]** field, type `y`.
9. In the **Enter the name of the backup file on the Remote server** field, enter the backup file that you need to restore.

   The file extension should be in `.gz` format, and the file naming convention is `full_backup_yyyy-mm-dd_hh-mm.bkp.gz`.

   The script restores the latest database backup from the remote server to the desired location on CCM.
10. Return to the previous page by entering `b`.

11. Quit the current page by entering `q`.
12. Return to the main menu by entering `m`.

Next steps

After the database backup is restored, enable the events coming to the Avaya Analytics™ database from Avaya Oceana® using the following commands:

- To start data flow to the cluster, run the following command:

  ```
  kubectl scale --replicas=2 deployment orca-ref-input-adaptor
  ```

- To start interval data, run the following command:

  ```
  kubectl scale --replicas=2 deployment orca-interval-controller
  ```

  Note:

  For non-HA systems, the replica count is 1 instead of 2.

Related links

Avaya Analytics backups on page 22

Incremental backup

Incremental backup overview

Using incremental backup, also called PgBackRest backup, you can schedule a daily automatic backup of the Avaya Analytics™ Postgres Cluster from the previous backup taken till the current day. Incremental database backup is compressed and stored on the pre-configured NFS.

Note:

Avaya Analytics™ Postgres Cluster includes both Avaya Analytics™ data and Historical Reporting custom report data.

The default schedule of incremental backup is the following:

- Base backup: It is a full database backup. By default, the backup is scheduled every Sunday at 04:00 a.m.
- Incremental backup: It is a backup taken based on the latest base backup. By default, the backup is scheduled daily at 00:00 a.m.

Important:

Incremental backups are based on the local system timezone.

Note:

Incremental backups (pgBackRest) and Remote backups (pgdump) are not interoperable.

Retention Policy

The retention policy is applicable by default on incremental backups. Retention policy enforces below:

- Up to one full backup can be stored on BackRest PV. When a new full backup is stored, the oldest backup is deleted from the NFS.
- There is no limit for incremental backup storage. Incremental backups cannot be independently deleted because all prior incremental backups (back to the last full backup) are needed to reconstruct the backup. For example, if you have seven incremental backups in a row, you cannot delete the first four.
- When a full backup is deleted, all the incremental backups based on this full backup is automatically deleted.
- There is no retention policy for full remote backups, as they are stored off the cluster.

Note:

There is an exception to the retention policy. When performing a restore, a backup is created at the end of the restore process. This backup is used for restoring the replica database and is a part of the pgBackRest operation. When this backup is created, it does not have any retention setting configured to not delete the oldest backup.

Related links

Avaya Analytics backups on page 22

Configuring incremental backup

About this task

You can run an immediate incremental backup and view a list of incremental backups.

Warning:

You should schedule backups during low traffic times or less busy as they can use resources on the database.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. To select the **Database backup** option, enter the corresponding number.
6. To take an incremental backup of the database, select the **Incremental backups** option by entering the corresponding number.
7. To select the **Reset Baseline for Incremental Backups** option, enter the corresponding number.

When you reset the baseline for Incremental backup, CCM takes a full pgBackRest backup on which the future Incremental backups will be based on.

Note:

You must run **Reset Baseline for Incremental Backups** before running the Incremental backup.

8. To select the **Run an Incremental backups** option, enter the corresponding number.

   Run this option if you want to manually run an ad-hoc Incremental backup.

9. In the **Proceed to run an incremental backup** option, type `y`.

   Entering `n` cancels the operation.

   An incremental backup is completed. The backup file is compressed in `.gz` file format and stored on NFS.

10. Return to the previous page by entering `b`.
11. Quit the current page by entering `q`.
12. Return to the main menu by entering `m`.
13. To view the list of the incremental database backups, select the **List incremental backups** option by entering the corresponding number.

    The page displays the list of incremental backups. The latest backup is listed at the bottom of the list with its timestamp.

14. Return to the previous page by entering `b`.
15. Quit the current page by entering `q`.
16. Return to the main menu by entering `m`.

Related links

Avaya Analytics backups on page 22

Scheduling an incremental database backup

About this task

You can schedule an incremental database backup based on your choice of time. The default schedule of incremental backups is the following:

- Base backup (full backup): Every Sunday at 04:00 am UTC
- Incremental backup: Daily at 00:00 am UTC

Note:

The full backup and incremental backup cannot be scheduled at the same time. There must be adequate time in between them, as both backups cannot run at the same time.

It is recommended not to have more than 7 incremental backups before running a full backup. The recommended schedule is to have 1 full back up per week and 1 incremental backup daily.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. To select the **Database backup** option, enter the corresponding number.
6. To select the **Incremental Backups** option, enter the corresponding number.
7. To select the **Update schedule for Incremental Backups** option, enter the corresponding number.
8. In the **Proceed with incremental scheduled backup update** field, type `y`.

   Entering `n` cancels the operation.

   The default schedules of incremental backup and base backup are listed.
9. In the **Which Backup Schedule do you want to update [full/incremental]:** field, type `incremental` or `full`, whichever backup you want to update.

   Depending on your selection, follow the next steps:

For daily incremental backups, do the following:

10. To select the time in hour for the backup, type the time in the **What hour of the day [0–23]** field and press **Enter**.

    The time selection is in 24–hour format. For example, if you type 2, the backup occurs at 2 a.m.
11. To select the time in minutes for the backup, type the time in the **What minute of the hour [0–59]** field and press **Enter**.

    For example, if you type 2, the backup occurs at 2 minutes past 2 a.m.
12. Return to the previous page by entering `b`.
13. Quit the current page by entering `q`.
14. Return to the main menu by entering `m`.

For weekly full backups, do the following:

15. To select the day of the week for the backup, type the number of the day in the **What day of the week (Sunday to Saturday) [0–6]** field and press **Enter**.

    For example, to select Monday, type `1`.

16. To select the time in hour for the backup, type the time in the **What hour of the day [0–23]** field and press **Enter**.

    The time selection is in 24–hour format. For example, if you type 2, the backup occurs at 2 a.m.

17. To select the time in minutes for the backup, type the time in the **What minute of the hour [0–59]** field and press **Enter**.

    For example, if you type 2, the backup occurs at 2 minutes past 2 a.m.

    Note:

    You can create one weekly full backup schedule and one daily incremental backup schedule. The schedule can be updated but not deleted.

18. Return to the previous page by entering `b`.
19. Quit the current page by entering `q`.
20. Return to the main menu by entering `m`.

Related links

Avaya Analytics backups on page 22

View list of scheduled incremental backups

About this task

Use this procedure to view the list of scheduled incremental backups.

Before you begin

Stop all events coming to the Avaya Analytics™ database from Avaya Oceana®. For steps to stop the events, see the relevant section in this document.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`

4. Select **Database** by pressing the corresponding number.
5. To select the **Database backup** option, enter the corresponding number.
6. To select the **Incremental backups** option, enter the corresponding number.
7. To view the list of currently scheduled full backups, select the **Show schedule for incremental backups** option by entering the corresponding number.

Note:

When you upgrade Avaya Analytics™ to a new version, all the backup schedules are deleted automatically. After the upgrade, you must create new schedules for database backup.

8. Return to the previous page by entering `b`.
9. Quit the current page by entering `q`.
10. Return to the main menu by entering `m`.

Related links

Avaya Analytics backups on page 22

Restoring the database from remote backup

About this task

You can restore a full backup of the database using the database restore option.

Note:

Do not restore a database backup of previous Avaya Analytics™ release version.

Restore needs to be done within a maintenance windows.

Warning:

Restore will delete all current data in the database and restore the contents of the backup.

Before you begin

- Stop all events coming to the Avaya Analytics™ database from Avaya Oceana®. For steps on stopping the events, see Stopping all events entering in to Avaya Analytics from Avaya Oceana on page 101.
- Check your backup file permissions.

Important:

The database must be in a healthy state before performing the restore. For example, no alarms must be triggered and all the pods must be in running state.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. To select the **Database Restore** option, enter the corresponding number.
6. Select the **Restore from a Remote Backup** option by entering the corresponding number.

7. In the **Proceed to backup restore** field, type `y`.

   Entering `n` cancels the operation.

8. In the **Confirm remote server details [y/n]:** field, type `y`.

   Entering `n` cancels the operation.

9. At the prompt, enter the full file name of the backup file located on the remote server that you want to restore.

   For example, `full_backup_2021-02-04_11-16.bkp.gz`.

   The script restores the database backup to the desired location.

10. Return to the previous page by entering `b`.
11. Quit the current page by entering `q`.
12. Return to the main menu by entering `m`.

Related links

Avaya Analytics backups on page 22

Restoring the database from incremental backup

About this task

You can restore an incremental database backup using the database restore option.

Note:

Do not restore a database backup of previous Avaya Analytics™ release version.

Before you begin

- Restore needs to be done within a maintenance window.

  Warning:

  Restore deletes all the current data in the database and restores the contents of the backup.

- Stop all events coming to the Avaya Analytics™ database from Avaya Oceana®. For steps to stop the events, see the relevant section in this document.
- Check your backup file permissions.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`

4. Select **Database** by pressing the corresponding number.

5. To select the **Database Restore** option, enter the corresponding number.
6. To select the **Restore from an Incremental Backup** option, enter the corresponding number.
7. In the **Proceed to backup restore** field, type `y`.

   Entering `n` cancels the operation. On entering `y`, the following warning message is displayed:

   ```
   Restoring Analytics Database from backup. Note that this will wipe
   all data currently on the database and replace with the contents of
   the backup.
   ```
8. In the **Confirm you want to continue** field, type `y`.

   The list of full and incremental backups is displayed. For example,

   Backup # 1 - Full backup at 2022-06-02 11:08:38

   Backup # 2 - Incremental backup at 2022-06-02 14:42:18
9. Enter the backup number that you want to restore in **From the backup you want to restore, please enter the Backup # (e.g. 1):** field.

   The script restores the latest database backup to the desired location.

   Note:

   If you see the following warning message in the logs, you can ignore it:

   ```
   WARNING: archiving write-ahead log file "0000000x.history"
   failed too many times, will try again later
   ```
10. After a successful restore, run a full base backup.

    Note:

    This step is mandatory after a successful restore.

    The first full backup after a restore shows the following error in the logs, you can ignore it:

    ```
    [WARN: a timeline switch has occurred since the
    20220615-124126F_20220616-144614I backup, enabling delta checksum\n
    HINT: this is normal after restoring from backup or promoting a
    standby.
    ```
11. Return to the previous page by entering `b`.
12. Quit the current page by entering `q`.
13. Return to the main menu by entering `m`.

Related links

Avaya Analytics backups on page 22

Dropping and restoring database schemas

About this task

After installing Avaya Analytics™, you can send test traffic to Avaya Oceana® for validating reports. When you are satisfied with the results, you can delete the test data using the steps in this procedure and start working afresh.

⚠ Warning:

These steps delete all data from the Avaya Analytics™ database.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. To select the **Drop and Restore Database Schemas** option, enter the corresponding number.

   The CCM console displays the following warning:

   `Continuing will remove all data from Database.`
6. In the **Confirm [y/n]** field, type `y`.

   Entering `n` cancels the operation.

   The CCM console starts installing the schemas and restarts the Avaya Analytics™ database pod.
7. Return to the previous page by entering `b`.
8. Quit the current page by entering `q`.
9. Return to the main menu by entering `m`.

Viewing the list of data retention limits

About this task

You can view a list of the existing data retention limits and also know the maximum limits of all the items for data retention.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.

2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. To select the **Data retention limits** option, enter the corresponding number.
6. To view a list of the database retention limits, select the **List Data retention limits** option by entering the corresponding number.
7. In the **Proceed to data retention limit list** field, type `y`.

   Entering `n` cancels the operation.

   The CCM console displays the list of the items for which you can view the retention limits. The options are:

   - Call Detail Records
   - Interval Data
   - Daily Data
   - Monthly Data
   - Login/Logout Data
   - Agent Trace Data
   - List all
8. In the **Select from the above which data retention limit you would like to list** field, enter the corresponding number of the item from the list you want to view.
9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.

Setting data retention limits

About this task

You can reset the default data retention limits for call detail records, interval data, daily data, monthly data, and login or logout data.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

```
ccm release orca analytics
```

4. Select **Database** by pressing the corresponding number.
5. To select the **Data retention limits** option, enter the corresponding number.
6. To select the **Set Data Retention limits** option, enter the corresponding number.
7. In the **Proceed to setting data retention limits** option, type `y`.

   Entering `n` cancels the operation.

   The CCM console displays the list of the items for which you want to set the data retention limits. The options are:

   - Call Detail Records
   - Interval Data
   - Daily Data
   - Monthly Data
   - Login/Logout Data
   - Agent Trace Data
   - List all

8. In the **Select from the above which data retention limit you would like to set** field, enter the corresponding number of the item from the list you want to set.

For Call Detail Records

9. In the **Enter the new limit value** field, enter the value in days before which you want the system to remove the CDR data.

   The maximum value is `365` days.

10. Return to the previous page by entering `b`.
11. Quit the current page by entering `q`.
12. Return to the main menu by entering `m`.

For Interval Data

13. In the **Enter the new limit value** field, enter the value in months before which you want the system to drop the partitions for the interval tables.

    The maximum value is `12` months.

14. Return to the previous page by entering `b`.
15. Quit the current page by entering `q`.
16. Return to the main menu by entering `m`.

For Daily Data

17. In the **Enter the new limit value** field, enter the value in months before which you want the system to drop the partitions for the interval tables.

The maximum value is `60` months.

18. Return to the previous page by entering `b`.
19. Quit the current page by entering `q`.
20. Return to the main menu by entering `m`.

For Monthly Data

21. In the **Enter the new limit value** field, enter the value in months before which you want the system to drop the partitions for the interval tables.

    The maximum value is `9999` months.

22. Return to the previous page by entering `b`.
23. Quit the current page by entering `q`.
24. Return to the main menu by entering `m`.

For Agent Trace Data

25. In the **Enter the new limit value** field, enter the value in days before which you want the system to remove the Trace date.

    The maximum value is `365` days.

26. Return to the previous page by entering `b`.
27. Quit the current page by entering `q`.
28. Return to the main menu by entering `m`.

For Login/Logout Data

29. In the **Enter the new limit value** field, enter the value in months before which you want the system to drop the partitions for the interval tables.

    The maximum value is `12` months.

30. Return to the previous page by entering `b`.
31. Quit the current page by entering `q`.
32. Return to the main menu by entering `m`.

Configuring data roll-up

About this task

Use this procedure to configure the daily and monthly roll-up delay duration. The daily roll-up duration is the time to wait after a day is closed before the data roll-up and the monthly roll-up duration is the time to wait after a month is closed before the data roll-up. The duration is calculated in minutes.

Note:

Though the menu displays `Configure Daily Roll-up`, you can also use this procedure to configure the delay for monthly roll up.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. To select the **Configure Daily Roll-Up** option, enter the corresponding number.
6. To view the list of the current roll-up delay, select the **Show the current Roll-Up delay** option by entering the corresponding number.

   The CCM console displays the current value in minutes.
7. Return to the previous page by entering `b`.
8. Quit the current page by entering `q`.
9. Return to the main menu by entering `m`.
10. To select the **Set Daily Roll-up delay** option, enter the corresponding number.
11. In the **Proceed to setting rollup delay** field, type `y`.

    Entering `n` cancels the operation.
12. In the **Enter the new delay value** field, enter the value in minutes.

    The maximum value is 4320 minutes, which is 3 days.
13. Return to the previous page by entering `b`.
14. Quit the current page by entering `q`.
15. Return to the main menu by entering `m`.

Configuring and deleting SNMP alarm destinations

About this task

You must configure alarms with Avaya™ Services as a destination.

To receive notifications on security or fault-related events, configure the details of your Network Management System (NMS), such as the authentication protocol configured for NMS and listening port of the configured NMS.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. To select the **SNMP** option, enter the corresponding number.
5. To select the **Configure a destination for sending SNMP alarms** option, enter the corresponding number.
6. In the **Proceed with SNMP Desctination config** field, type **y**.

   Entering `n` cancels the operation.
7. In the **Hostname/IP of the configured NMS destination** field, type the hostname or IP address of your NMS and press **Enter**.
8. In the **Listening port of the configured NMS** field, enter the port number of your NMS.

   The default option is `162`.
9. In the **Username of the SNMPv3 user configured for the NMS** field, type the username of an SNMP user that can access the NMS and press **Enter**.
10. In the **Password of the SNMPv3 user configured for the NMS** field, type the password of an SNMP user that can access the NMS and press **Enter**.
11. In the **Authentication protocol configured for NMS** field, type the NMS authentication protocol in use and press **Enter**.

    The options are:

    - `SHA`
    - `MD5`
12. In the Privacy protocol configured for NMS, type the NMS privacy protocol in use and press **Enter**.

    The options are:

    - `AES`
    - `DES`
13. In the **Authentication password** field, type the NMS privacy protocol password and press **Enter**.

    Leave this field blank if you do not have a password.
14. In the **Privacy password** field, type the NMS authentication password and press **Enter**.

    Leave this field blank if you do not have a password.

    The CCM console displays the new NMS settings.

15. Review the NMS settings.
16. If you are satisfied with the settings, in the **Please confirm your settings** field, type `y` and press **Enter**.

    This CCM console adds the SNMP destination and displays the alarm details.
17. Return to the previous page by entering `b`.
18. Quit the current page by entering `q`.
19. Return to the main menu by entering `m`.
20. To view the list of the current destinations for sending SNMP alarms, select the **List the current destination for sending SNMP alarms** option by entering the corresponding number.
21. Return to the previous page by entering `b`.
22. Quit the current page by entering `q`.
23. Return to the main menu by entering `m`.
24. To send a test SNMP alarm to a destination, select the **Send a test SNMP alarm to destination** option by entering the corresponding number.

    Avaya Analytics™ sends a test alarm to the configured destination.
25. Return to the previous page by entering `b`.
26. Quit the current page by entering `q`.
27. Return to the main menu by entering `m`.
28. To delete an SMNP configured destination, select the **Delete a configured destination address** option by entering the corresponding number.
29. In the **Proceed to SNMP destination deletion** field, type `y`.
30. In the **Enter the destinations IP address** field, type the IP address of the destination that you want to delete and press **Enter**.
31. Return to the previous page by entering `b`.
32. Quit the current page by entering `q`.
33. Return to the main menu by entering `m`.

Generating Management Information Base file for SNMP

About this task

Management Information Base (MIBs) files describe the Object Identifiers (OIDs) for SNMP alarms. Use the following procedure to generate the MIB file.

Procedure

1. Log in to Cluster Control Manager (CCM) console as the customer user.
2. To switch to the root user, type `su` and press **Enter**.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. To select the **SNMP** option, enter the corresponding number.
5. To select the **Generate MIB** option, enter the corresponding number.

   The system displays the following message:

   `/tmp/Analytics.MIB file is created successfully`

   `/tmp/CommonServices.MIB file is created successfully`

   The MIB file is created in the `/tmp` directory.
6. Return to the previous page by entering `b`.
7. Quit the current page by entering `q`.
8. Return to the main menu by entering `m`.

Registering for certification expiration alarms

About this task

You can register to get notified when a certificate expires. The alarm notifies you 8 weeks before the registered certificate expires, so that you get enough time to renew it.

Before you begin

Configure SNMP destination to receive alarms and to get monitored.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. To select the **SNMP** option, enter the corresponding number.
5. To select the **Certification Expiration Alarming** option, enter the corresponding number.
6. To register the certificates for SNMP alarms, type the number corresponding to the required certificate.

The options are:

- Oceana REF Certificate
- Oceana Authentication Certificates
- Breeze Certificate

7. Depending on your selection, in the **Proceed with <certificate> registration** field, type `y` or `n`, where certificate is the type of the *certificate* you select in the previous step.
8. If you type `y`, in the **Please enter the name of the keystore including the full path to the <certificate> file** field, type the location of the file.

   For example, `/home/cust/sslRef/ref-input-adaptor.jks`

   Here, *certificate* is the type of the certificate you selected earlier.
9. If you type `n`, CCM console takes you back to the registration page.
10. In the **Please enter the password to this keystore** field, type the password.

    The certified is now registered for expiration notification alarms.

Configuring Geo database replication for alarming

About this task

The SNMP alarm notifies you of a Geo streaming failure on the crunchy database. An SQL query that returns the Geo status triggers the alarm every set period. Enable the alarm only for Geo sites. You need to enable the alarm on the Primary and DR Geo sites. Both sites can trigger the alarm independently. You can receive the alarm twice at the SNMP destination. The alarm log is located at `/var/log/avaya/ccm/analytics_geo_db.log`.

You can enable or disable the alarm using the `Analytics Administration` script.

Note:

Disable the alarm before the known network outages or scheduled maintenance windows to the Avaya Analytics™ database and enable the alarm after that.

Before you begin

Configure the SNMP destination to receive alarms and send a test alarm to confirm it is working. For steps on configuring the SNMP alarm destination, see Configuring and deleting SNMP alarm destinations on page 43.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`

4. To select the **SNMP** option, enter the corresponding number.
5. To select the **Geo Database Replication for Alarming** option, enter the corresponding number.
6. Do one of the following:
   a. To enable the **Geo Database Replication for Alarming** option, enter the corresponding number and, at the prompt, set the frequency for the alarm.

      The alarm frequency must be greater than zero. The default value is `10` minutes. For example, if you want to trigger the alarm every hour, set `60` minutes.

   b. To disable the **Geo Database Replication for Alarming** option, enter the corresponding number.
7. Return to the previous page by entering `b`.
8. Quit the current page by entering `q`.
9. Return to the main menu by entering `m`.

Replacing the authorization certificates in the breeze-security secret

If you have replaced the authorization certificates on Avaya Oceana®, you can manually replace the certificates in the breeze-security secret. You can perform this procedure instead of repeating the Avaya Breeze® certificate creation process. For more information about the manual replacement see, Manually replacing the authorization certificates in the breeze-security secret on page 49.

Alternatively, you can use the `Analytics Administration` script and perform the following tasks to recreate the Avaya Analytics™ certificates:

1. Create the certificate signing request for Avaya Oceana® authentication.
2. Get identity token.
3. Import signed certificate for Avaya Oceana® authentication.
4. Retrieve the identity certificates.
5. Create Avaya certificates keystore.
6. Restart Avaya Authentication.
7. Restart Streams REST.
8. Restart Data Publisher.
9. Restart the Open Kafka interface.

For more information about the procedures mentioned above, see *Deploying Avaya Analytics™ for Avaya Oceana®*.

Manually replacing the authorization certificates in the breeze-security secret

About this task

You need to add the authorization certificates to the new directory and create a new secret using the content of the old secret.

Recreating the secret requires the secret content input in cleartext. If the corresponding information is not available in the documentation, you need to decode all token values of the existing secret from base64 to cleartext. Note the output to prepare for the next steps.

Note:

This is a manual procedure that involves kubectl commands and requires familiarity with the Avaya Analytics™ environment. Alternatively you can perform the procedures from *Deploying Avaya Analytics™ for Avaya Oceana®*. For more information, see Replacing the authorization certificates in the breeze-security secret on page 48.

Before you begin

Backup your existing certificates directory. For example, **`cp -r /home/cust/ssl /home/cust/ssl_backup`**.

Procedure

1. To create a new directory and copy the content from the existing directory, run the following commands:

   ```
   mkdir /home/cust/ssl_new
   cp /home/cust/ssl/clienttruststore.jks /home/cust/ssl_new
   cp /home/cust/ssl/clientkeystore.jks /home/cust/ssl_new
   ```

   You need to create a new directory to store certificate files and copy the content from the existing certificate directory.

2. To change to the new directory and create a backup of the existing secret, run the following commands:

   ```
   cd /home/cust/ssl_new
   kubectl get secret breeze-security -o yaml > breeze-security.backup
   ```

3. To pull the new certificates from SMRG and Avaya Breeze® Avaya Oceana® Cluster 2 nodes, run the following commands:

   ```
   echo -n | openssl s_client -connect <SMGR_IP_ADDRESS>:443 | sed -ne '/-
   BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > <SMGR_IP_ADDRESS>.crt
   echo -n | openssl s_client -connect
   <BREEZE_1_SIP_ENTITY_IP_ADDRESS>:443 | sed -ne '/-BEGIN CERTIFICATE-
   /,/-END CERTIFICATE-/p' > <BREEZE_1_SIP_ENTITY_IP_ADDRESS>.crt
   echo -n | openssl s_client -connect
   <BREEZE_2_SIP_ENTITY_IP_ADDRESS>:443 | sed -ne '/-BEGIN CERTIFICATE-
   /,/-END CERTIFICATE-/p' > <BREEZE_2_SIP_ENTITY_IP_ADDRESS>.crt
   ```

   If there is no Avaya Oceana® Cluster 2, run the commands for three Avaya Oceana® Cluster 1 nodes.

4. In the existing certificate directory, compare the content of the CRT files with the corresponding files.
5. To delete the changed certificates from `clienttruststore.jks`, run the following commands:

```
keytool -delete -alias SMGR_CRT -keystore clienttruststore.jks
keytool -delete -alias BREEZE_1 -keystore clienttruststore.jks
keytool -delete -alias BREEZE_2 -keystore clienttruststore.jks
```

6. To add new certificates to `clienttruststore.jks`, run the following commands:

```
keytool -import -noprompt -alias SMGR_CRT -keystore
clienttruststore.jks -file <SMGR_IP_ADDRESS>.crt
keytool -import -noprompt -alias BREEZE_1 -keystore
clienttruststore.jks -file <BREEZE_1_SIP_ENTITY_IP_ADDRESS>.crt
keytool -import -noprompt -alias BREEZE_2 -keystore
clienttruststore.jks -file <BREEZE_2_SIP_ENTITY_IP_ADDRESS>.crt
```

7. For each node, export authentication services identity certificates from SMGR:
   a. Go to **Services** > **Inventory** > **Manage Elements**.
   b. Select the Avaya Breeze® node with installed authentication services.
   c. Click **More Actions** > **Manage Identity Certificates**.
   d. Click **Authorization** > **Export**.
   e. Rename the exported file to `AvayaBreezeNode1.pem`.

      Name files according to the node number. For example, `AvayaBreezeNode2.pem`, `AvayaBreezeNode3.pem`.
8. Copy the files to the certificate folder on Cluster Control Manager.
9. To convert the certificate files into binary format, run the following commands:

```
openssl x509 -outform der -in AvayaBreezeNode1.pem -out
AvayaBreezeNode1.der
openssl x509 -outform der -in AvayaBreezeNode2.pem -out
AvayaBreezeNode2.der
```

10. To create a new Avaya Breeze® certificate store, run the following commands:

```
keytool -import -alias AvayaBreezeNode1 -keystore BreezeCerts -file
AvayaBreezeNode1.der
keytool -import -alias AvayaBreezeNode2 -keystore BreezeCerts -file
AvayaBreezeNode2.der
```

11. To list the content of the existing secret, run the following command:

    `kubectl get secret breeze-security -o yaml`

    The following is a sample output of the command:

```
server1breeze1TokenEndpoint:
aHR0cHM6Ly8xMDAuOTYuOTIuMTEzOjk0NDMvc2VydmljZXMvQXV0aG9yaXphdGlvblNlcnZpY2UvdG9rZW
4=
server1breeze2TokenEndpoint:
aHR0cHM6Ly8xMDAuOTYuOTIuMTE1Ojk0NDMvc2VydmljZXMvQXV0aG9yaXphdGlvblNlcnZpY2UvdG9rZW
4=
server1clientKey: UXRuZTdfM1dUc0txUnZOdHVlSDFtdw==
server2breeze1TokenEndpoint: ""
server2breeze2TokenEndpoint: ""
```

```
server2clientKey: ""
smgrPemAlias: YXZheWFicmVlemVub2RlMSxhdmF5YWJyZWV6ZW5vZGUy
storePwd: bXlwYXNzd29yZA==
```

12. To decode the token values from step 11 from base64 to clear text, run the following command **`echo -n 'token values' | base64 -d`**

    For example, **`echo -n 'bXlwYXNzd29yZA==' | base64 -d`**, where `bXlwYXNzd29yZA` is the store password.

    You can also use the corresponding information from the documentation if it is available.

13. To delete the existing secret, run the following command:

    `kubectl delete secret breeze-security`

    Ensure you have made a backup of the existing secret at the beginning of this procedure.

14. Create a new secret using the parameters of the old secret.

    The following is a sample command:

```
kubectl create secret generic breeze-security --from-
literal=storePwd=<passwd> --from-file=/home/cust/ssl_new/clientkeystore.jks -
--from-file=/home/cust/ssl_new/clienttruststore.jks --from-
literal=server1clientKey=<server1clientkey> --from-
literal=server1breeze1TokenEndpoint=<server1_breeze1_token_endpoint> --from-
literal=server1breeze2TokenEndpoint=<server1_breeze2_token_endpoint> --from-
file=/home/cust/ssl_new/BreezeCerts --from-
literal=smgrPemAlias=<smgrPemAlias>
```

15. Restart the following pods:

    - `breeze-auth`
    - `streams-rest`
    - `data-publisher`
    - `open-interface`

    You can restart the pods manually or using the `Analytics Administration` script. For more information, see Restarting a specific pod on page 148.

16. **(Optional)** Clean up the installation directories.

    You can copy the new keystore files to the original certificate directory and remove the new directory.

17. **(Optional)** To restore the original secret from the backup, run the following command:

    `kubectl apply -f breeze-security.backup`

Configuring Routing Service groups for Historical Reporting

You can incorporate Routing Service groups into Avaya Analytics™ Historical Reporting using the `Historical Routing Service Groups Configuration` option. Configuring Routing Services groups in Historical Reporting allows supervisors to view historical information about the routing services in the Routing Service groups to which they are assigned.

You can assign supervisors to a group called System Routing Service group, so that supervisors can bypass individual Routing Service group filters and access all routing services in all groups. Such supervisors will have unrestricted access to all routing service information across routing service groups.

Checking Routing Service Group feature status

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Deployment** by pressing the corresponding number.
5. Select the **Historical Routing Service Groups Configuration** option by entering the corresponding number.
6. Select the **Check feature status** option by entering the corresponding number.

   The CCM console displays the current Routing Service group name and feature status as `true` or `false`, indicating whether the feature is currently enabled or disabled.
7. Return to the previous page by entering `b`.
8. Quit the current page by entering `q`.
9. Return to the main menu by entering `m`.

Enabling or disabling Routing Service groups feature

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Deployment** by pressing the corresponding number.

5. Select the **Historical Routing Service Groups Configuration** option by entering the corresponding number.
6. Select the **Enable/Disable Historical Routing Service groups feature** option.
7. Select one of the following options:
   a. Select **Enable Historical Routing Service groups feature** to enable the feature.

      Note:

      If you enable this feature, the `SPECIFIED_RESOURCE` data will not display in the report.

   b. Select **Disable Historical Routing Service groups feature** to disable the feature.
8. Type `y` to proceed with enabling or disabling the feature.

   Depending on whether you selected `a` or `b` in step 6, the CCM console displays a confirmation that the feature was successfully enabled or disabled.
9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.

Configuring System Routing Service group name

About this task

Use this procedure to configure System Routing Service group name.

Note:

Avaya recommends that you first create the Routing Service group in ACM to comply with the ACM group naming policy.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Deployment** by pressing the corresponding number.
5. Select the **Historical Routing Service Groups Configuration** option by entering the corresponding number.
6. Select the **Configure system group name** option by entering the corresponding number and typing `y`.
7. Enter a name in the **Please input the system name for Routing Service group** field.

Note:

You cannot proceed by leaving this field blank.

If you enter a valid name for the group, the CCM console displays that the System Routing Service name has been successfully saved. The console also displays that the System Routing Service group name is configured only if the Routing Service group feature is enabled.

8. Return to the previous page by entering `b`.
9. Quit the current page by entering `q`.
10. Return to the main menu by entering `m`.

Historical Reporting user settings

The following table outlines the Avaya Analytics™ defined user privileges.

Persona	Consumer	Basic users	Advanced users
You can	• Run dossiers, reports, and documents. • Perform basic operations, such as sort, pivot, drill, and export.	• All the capabilities provided to the consumer. • Create new dossiers, reports, and documents. • Modify or customize Avaya Analytics™ canned reports. • Access and import data from local machine for file formats, such as excel file, CSV files, and JSON files.	• All the capabilities provided to the basic users. • Interact with rest APIs to push external data, extract data subsets, and create datasets. • Connect and access data from a range of third-party databases.
You cannot	• Save or design new dossiers, reports or documents. • Create objects like prompts, filters, and metrics.	• Create objects like prompts and filters. • Create schema objects such as attributes and facts. • Leverage rest APIs. For example, access data pushed through the API.	• Create schema objects such as attributes, facts, and hierarchies.

Resetting Historical Reporting administrator password

About this task

You can change a range of default historical reporting settings, such as managing historical reporting users after resetting your administrator password. The default administrator password is listed in the deployment spread sheet.

Important:

Do not change the Historical Reporting administrator password using the Historical Analytics GUI. If you do so, the underlying passwords could get out of sync.

If the password policy is enabled, it is applied on administrator password, then you must to reset the password as per your choice.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. Select **Historical Reporting Configuration** by pressing the corresponding number.
6. To change the default administrator password for historical reporting, select the **Administrator Password reset** option by entering the corresponding number.
7. In the **Proceed with admin password reset?** field, type `y` or `n`.

   Entering `n` cancels the operation.

   When you type `y`, CCM displays a prompt to enter the current administrator password.
8. At the prompt, enter the current administrator password.

   The default administrator password is listed in the deployment spread sheet.
9. Enter and confirm the new password.

   Adminstrator password is configured. The mstr-srv and mstr-web pods are restarted. Historical reporting is restarted.
10. Return to the previous page by entering `b`.
11. Quit the current page by entering `q`.
12. Return to the main menu by entering `m`.

Configuring DataWarehouse Password for Historical Reporting

About this task

You can change the default DataWarehouse password (whuserservice) for the Postgres historical reporting database.

Note:

If the password policy is enabled, it is applied on DataWarehouse password, then you must to reset the password as per your choice.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. Select **Historical Reporting Configuration** by pressing the corresponding number.
6. Select the **Configure DataWarehouse Password for Historical Reporting** option by entering the corresponding number.
7. In the **Proceed with DataWarehouse password config?** option, type `y` or `n`.

   Typing `n` cancels the operation.

   When you type `y`, CCM displays a prompt to enter the credentials to the DataWarehouse database.
8. Enter the DataWarehouse database password.
9. Enter and confirm the new password.

   Note:

   When the DataWarehouse password is changed, it invalidates all the backups taken with the old password. Therefore, ensure that you take a full backup of the Postgres historical reporting database immediately after changing the password.

   The DataWarehouse (whuserservice) password is changed. The mstr-srv pod is restarted. Historical reporting is restarted.
10. Return to the previous page by entering `b`.
11. Quit the current page by entering `q`.
12. Return to the main menu by entering `m`.

Creating a local user for Historical Reporting

About this task

You can create a local user and add the user to a group for the user to avail the respective Avaya Analytics™ defined user privileges. However, only a user configured with supervisor role in Avaya Control Manager can get access to the historical reports.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. Select **Historical Reporting Configuration** by pressing the corresponding number.
6. To create a local user for Historical Reporting, select the **Create a Historical Reporting local user (valid ACCCM Reporting Supervisor)** option by entering the corresponding number.
7. In the **Proceed with user creation** option, type `y`.

   Entering `n` cancels the operation.
8. At the prompt, enter the administrator password.
9. In the **Username** field, type the username of the supervisor who is a valid supervisor in Avaya Control Manager.
10. Enter the new user name and password.

    Note:

    When you assign a name to this user, ensure that a user of the same name does not already exist on the LDAP server. If a local user in Historical Reporting has the same user name as an existing LDAP user, then, after this LDAP user logs in to Avaya Analytics™, the LDAP user profile gets linked to the local user profile of the same name in Historical Reporting. As a result, the privileges of the existing LDAP user are enhanced to the same level as the privileges of this newly created local user in Historical Reporting.

    After linking the LDAP user to the local user with the same name in Historical Reporting, the local user profile no longer appears in the list of Historical Reporting users, because the list cannot include users with LDAP accounts.
11. Confirm the password.
12. In the **New Users Fullname** field, type the new user's full name.

    The CCM console displays that message that the user is created.

13. Return to the previous page by entering `b`.
14. Quit the current page by entering `q`.
15. Return to the main menu by entering `m`.

Next steps

Add user to an Avaya Analytics™ group.

Adding a local user to a group

About this task

After you create a historical reporting user, you must add them to one of the following groups: Consumer, Basic, or Advanced.

Before you begin

Create a local user. See Creating a local user for Historical Reporting on page 57.

Note:

When you assign a name to this user, ensure that a user of the same name does not already exist on the LDAP server. If a local user in Historical Reporting has the same user name as an existing LDAP user, then, after this LDAP user logs in to Avaya Analytics™, the LDAP user profile gets linked to the local user profile of the same name in Historical Reporting. As a result, the privileges of the existing LDAP user are enhanced to the same level as the privileges of this newly created local user in Historical Reporting.

After linking the LDAP user to the local user with the same name in Historical Reporting, the local user profile no longer appears in the list of Historical Reporting users, because the list cannot include users with LDAP accounts.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`

4. Select **Historical Reporting** by pressing the corresponding number.
5. Select **Historical Reporting Configuration** by pressing the corresponding number.
6. To add a local user to a group, select the **Add a local user (valid ACCCM Reporting Supervisor) to a group** option by entering the corresponding number.
7. In the **Proceed with adding user to group?** option, type `y`.

   Typing `n` cancels the operation.

8. In the **Please enter the current Administrator password for Historical Reporting Administrator Password** option, enter the password.
9. In the **Username** option, enter the user name that you want to add to a group.
10. Enter the name of the group you want to add the user to.

    You can enter `Consumer`, `Basic`, or `Advanced`.

    The user can now access Avaya Analytics™ Web and run historical reports.
11. Return to the previous page by entering `b`.
12. Quit the current page by entering `q`.
13. Return to the main menu by entering `m`.

Removing a local user from a group

About this task

After you create a historical reporting user and add them to one of the following groups: Consumer, Basic, or Advanced, you can also remove the user from a group.

Before you begin

The local user must exists in one of the following three groups:

- Consumer
- Basic
- Advanced

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. Select **Historical Reporting Configuration** by pressing the corresponding number.
6. To remove a user from a group, select the **Remove a local user (valid ACCCM Reporting Supervisor) from a Group** option by entering the corresponding number.
7. In the **Proceed with removing user to group?** option, type `y`.

   Typing `n` cancels the operation.
8. At the prompt, enter the current administrator password.

9. In the **Username** option, enter the user name that you want to remove from the group.
10. Enter the name of the group you want to remove the user from.

    You can enter `Consumer`, `Basic`, or `Advanced`.

    The script removes the user from the Avaya Analytics™ user group.
11. Return to the previous page by entering `b`.
12. Quit the current page by entering `q`.
13. Return to the main menu by entering `m`.

Resetting a Historical Reporting local user password

About this task

After you create a historical reporting user and add them to one of the following groups: Consumer, Basic, or Advanced, you can also reset the user password from a group.

Before you begin

The local user must exist in one of the following three groups:

- Consumer
- Basic
- Advanced

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. Select **Historical Reporting Configuration** by pressing the corresponding number.
6. To reset a historical reporting local user password, select the **Resetting a Historical Reporting local user password** option by entering the corresponding number.
7. In the **Proceed with Resetting a Historical Reporting local user password?** option, type `y`.

   Typing `n` cancels the operation.
8. At the prompt, enter the current administrator password.
9. At the prompt, enter the user name of the historical reporting local user that you want to reset.

10. At the prompt, enter a new password for the historical reporting local user.
11. At the prompt, confirm the new password for the historical reporting local user.
12. Return to the previous page by entering `b`.
13. Quit the current page by entering `q`.
14. Return to the main menu by entering `m`.

Deleting a Historical Reporting local user

About this task

`You can delete a Historical Reporting local user.`

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. Select **Historical Reporting Configuration** by pressing the corresponding number.
6. To delete a historical reporting local user, select the **Delete a Historical Reporting local user (valid ACCCM Reporting Supervisor)** option by entering the corresponding number.
7. In the **Proceed with Delete local User** option, type `y`.

   Typing `n` cancels the operation.
8. At the prompt, enter the username of the historical reporting local user that you want to delete.

   The local user is deleted.
9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.

Viewing list of Historical Reporting local users

About this task

You can view the list of all the historical reporting users.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. Select **Historical Reporting Configuration** by pressing the corresponding number.
6. To view list of historical reporting local users, select the **List all Historical Reporting local users** option by entering the corresponding number.
7. In the **Proceed with List all Historical Reporting local users?** option, type `y`.

   Typing `n` cancels the operation.

   A list of the historical reporting local users is displayed.
8. Return to the previous page by entering `b`.
9. Quit the current page by entering `q`.
10. Return to the main menu by entering `m`.

Configuring Avaya Analytics™ email distribution services

About this task

You can configure the Avaya Analytics™ email distribution services to send and receive emails of Avaya Analytics™ reports in excel and pdf formats.

Note:

HTML data format is currently not supported.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`

4. Select **Historical Reporting** by pressing the corresponding number.
5. Select **Historical Reporting Configuration** by pressing the corresponding number.
6. Select **Email Distribution Service Configuration** option by entering the corresponding number.
7. Select **Set up an Email Distribution Service** option by entering the corresponding number.
8. In the **Proceed to distribution services config?** option, type `y`.

   Typing `n` cancels the operation.
9. In the **SMTP Server** field, type the FQDN of the SMTP server and press **Enter**.
10. In the **Port** field, type the port for the email server and press **Enter**.
11. In the **E-mail** field, type the source email address and press **Enter**.
12. In the **E-mail display name** field, type the display name for the source email address.

    This email address is the one that you want Avaya Analytics™ to display when the email recipient replies to any mails. For example, `Analytics Services`.

    Note:

    Once you press **Enter**, the updates are saved and the mstr-srv pod automatically restarts. You cannot review the changes after this step. If you are sure, press **Enter**.

    After you enter the details, press **Enter**.

    The CCM console displays the message that updating the distribution settings is complete.
13. Return to the previous page by entering `b`.
14. Quit the current page by entering `q`.
15. Return to the main menu by entering `m`.

Viewing Email Distribution Service configuration

About this task

You can view the current Email Distribution Service configuration using Cluster Control Manager (CCM).

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`

4. To select the **Historical Reporting** option, enter the corresponding number.
5. To select the **Historical Reporting Configuration** option, enter the corresponding number.
6. To select the **Email Distribution Service Configuration** option, enter the corresponding number.
7. To select the **Show email Distribution Service Configuration** option, enter the corresponding number.

   The CCM console displays current Email Distribution Service settings.
8. Return to the previous page by entering `b`.
9. Quit the current page by entering `q`.
10. Return to the main menu by entering `m`.

Updating the governing rule for Historical Reporting

About this task

You can update the default governing rule to change the view of a historical report.

Note:

Ensure that you apply the appropriate filters at the input prompts so that you get optimum results.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. Select **Historical Reporting Configuration** by pressing the corresponding number.
6. To update the default governing rule for returning the number of records for a report, select the **Governing Rule** option.

   At the prompt to enter the Administrator password, enter your administrator password.
7. In the **Proceed with setting governing rules?** option, type `y`.

   Typing `n` cancels the operation.
8. In the **Maximum number of rows** field, enter the maximum number of rows that you can view in a historical report.

The default value is 32,000 rows.

9. In the **Maximum number of elements** field, enter the maximum number of elements returned in prompt values.
10. In the **Maximum number of intermediate rows** field, enter the maximum number of intermediate result rows to view in a report.

    The CCM console displays that the updates are complete.
11. Return to the main menu by entering `m`.
12. Quit the current page by entering `q`.

Viewing and configuring user idle timeouts

About this task

You can view and configure the Web Uses Session Idle time and the web server session timeout using this procedure.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. Select **Historical Reporting Configuration** by pressing the corresponding number.
6. To set the users session idle time timeout, select the **Set User Session Idle Timeout** option.
7. In the **Proceed with Set User Session Idle Timeout?** option, type `y`.

   Typing `n` cancels the operation.

   At the prompt, in enter the Administrator password, enter your administrator password.

   The current timeout settings are returned to the console.
8. In the **Continue to set user session idle time [y/n]** option, type `y`.

   Typing `n` cancels the operation.
9. In the **Please enter new web user session idle time in seconds** field, enter the new web user idle time timeout value in seconds. For example, 600 seconds equals 10 minutes.

10. In the **Please enter new web session timeout in minutes** field, enter the new web session timeout value in minutes.

    Note:

    A restart of the mstr-srv and mstr-web pods are needed to update the values.

11. To return to the previous page, type `b` and press **Enter**.
12. To quit from the current page, type `q` and press **Enter**.
13. To return to the main menu, type `m` and press **Enter**.

Configuring Avaya Analytics™ LDAP authentication

About this task

You can configure Lightweight Directory Access Protocol (LDAP) authentication by using the post-install scripts. The connection to the LDAP server must be secure.

When you configure a user for LDAP authentication in Avaya Analytics™, you must add them to one of the following groups: *Consumer*, *Basic*, or *Advanced*.

Note:

- When you create, and assign a name to a local user in Historical Reporting, ensure that a user of the same name does not already exist on the LDAP server.
  - If a local user in Historical Reporting has the same user name as an existing LDAP user, after the LDAP user logs in to Avaya Analytics™, the LDAP user's profile gets linked to the local user profile of the same name in Historical Reporting. As a result, the privileges of the LDAP user are enhanced to the same level as the privileges of the newly created local user in Historical Reporting.
  - After linking the LDAP user to the local user with the same name in Historical Reporting, the local user profile no longer appears in the list of Historical Reporting users, because the list cannot include users with LDAP accounts.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   ```
   ccm release orca analytics
   ```

4. Select **Historical Reporting** by pressing the corresponding number.
5. Select **LDAP configuration** by entering the corresponding number.
6. In the **Proceed to LDAP over SSL config?** option, enter `y`. Entering `n` cancels LDAP configuration.

7. In the **LDAP host FQDN address** field, enter the host name of the LDAP server.
8. In the **LDAP port number** field, enter the port number of the LDAP.
9. In the **LDAP user DN** field, enter the distinguished name of the trusted LDAP authentication user for the LDAP repository searches.
   - For example, `CN=test,CN=Users,DC=test,DC=TEST,DC=COM`
10. In the **LDAP password** field, enter the password for the trusted LDAP authentication user.
    - For example, `LDAP_PWD`
11. In the **LDAP server vendor name** field, enter `ADS`.
12. In the **LDAP authentication method** field, enter `BINDING`.
13. In the **LDAP search root DN** field, enter the root distinguished name on the LDAP server.
    - For example, `CN=Users,DC=test,DC=TEST,DC=COM`
14. In the **LDAP search filter user** field, enter the LDAP search filter for importing users in a batch.
    - For example, `(&(objectclass=person) (sAMAccountName=#LDAP_LOGIN#))`
15. In the **LDAP search filter group** field, enter the LDAP search filter for importing groups in a batch.
    - For example, `(&(objectclass=group) (member=#LDAP_DN))`
16. In the **basicLDAPLINK DN** field, enter the input parameter for linking the Basic group to the LDAP group.
    - For example, `CN=Basic,CN=Users,DC=test,DC=TEST,DC=COM`
17. In the **advancedLDAPLINK DN** field, enter the input parameter for linking the Advanced group to the LDAP group.
    - For example, `CN=Advanced,CN=Users,DC=test,DC=TEST,DC=COM`
18. In the **consumerLDAPLINK DN** field, enter the input parameter for linking the Consumer group to the LDAP group.
    - For example, `CN=Consumer,CN=Users,DC=test,DC=TEST,DC=COM`
19. To confirm your settings, enter `y` or `n`.
    - If you enter `y`, the CCM console displays the message: `LDAP configuration added`.
      - By default, running this script also links LDAP to groups.
      - The MicroStrategy server pod restarts in 10 minutes.
    - If you enter `n`, you are returned to the beginning of the LDAP configuration prompts, allowing you to update your LDAP configuration settings.
20. Return to the previous page by entering `b`.

21. Quit the current page by entering `q`.
22. Return to the main menu by entering `m`.

Configuring certificates for secure LDAP connection

About this task

If you used the post install script to configure Avaya Analytics™ Lightweight Directory Access Protocol (LDAP) authentication, you must also add and configure certificates for the Historical Reporting secure LDAP connection. You can do that using the following process.

Before you begin

1. Configure Avaya Analytics™ LDAP authentication.
2. Get the `.pem` file from your LDAP server. You can get the `.pem` from your IT personnel or system administrator.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. Copy the `.pem` file into a directory. For our examples we assume that the tmp directory.
4. Rename the certificate as `LDAP_CACERT` using the following command: **`mv <filename>.pem LDAP_CACERT`** where `<filename>` is the current name of the file.
   - For example: `mv cacert.pem LDAP_CACERT`
5. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
6. Select **Historical Reporting** by pressing the corresponding number.
7. Select the **Configure Certificate for LDAP using SSL** option by entering the corresponding number.
8. In the **Proceed LDAP over SSL cert config?** option, type `y`. Typing `n` cancels the operation.
9. In the **Please confirm that you renamed the LDAP CA certificate to LDAP_CACERT and saved to a location in CCM** field, type **y** and press **Enter**.
10. In the **Please enter the LDAP_CACERT file name including the path** field, enter the location and filename of the renamed LDAP certificate file.
    - For example, `tmp/LDAP_CACERT`
    - The MicroStrategy server pod restarts in 10 minutes.

11. **(Optional)** To cancel the renaming of the LDAP CA certificate, in the **Please confirm that you renamed the LDAP CA certificate and saved to the Historical Reporting Certificate directory** field, type **y** and press **Enter**.
    - Canceling allows renaming of the LDAP CA certificate to `LDAP_CACERT` and saving it to a location on CCM.
12. Return to the previous page by entering `b`.
13. Quit the current page by entering `q`.
14. Return to the main menu by entering `m`.

Viewing LDAP settings

About this task

You can view current LDAP settings using Cluster Control Manager (CCM).

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. To select the **Historical Reporting** option, enter the corresponding number.
5. To select the **LDAP Configuration** option, enter the corresponding number.
6. To select the **Show LDAP settings** option, enter the corresponding number.

   The CCM console displays current LDAP settings.
7. Return to the previous page by entering `b`.
8. Quit the current page by entering `q`.
9. Return to the main menu by entering `m`.

Configuring EASG availability for Historical Reporting

You can enable or disable Enhanced Access Security Gateway (EASG) availability for Historical Reporting.

Enabling EASG availability for Historical Reporting

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. To switch to the root user, type `su` and press **Enter**.
3. Use the following command to run the Avaya Analytics™ Administration script:

   ```
   ccm release orca analytics
   ```

4. Enter the corresponding number to select the **Historical Reporting** option.
5. Select the **Historical reporting EASG availability** option by entering the corresponding number.
6. Select the **Enable EASG** option by entering the corresponding number.
7. In the **Historical Reporting enable easg availability?** field, type `y` or `n`.

   Typing `n` cancels the operation.
8. At the prompt, **Do you want to enable EASG [y/n]** type `y` or `n`.

   Typing `n` cancels the operation.
9. To go back to the previous page, type `b` and press **Enter**.
10. To leave the current page, type `q` and press **Enter**.
11. To return to the main menu, type `m` and press **Enter**.

Disabling EASG availability for Historical Reporting

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. To switch to the root user, type `su` and press **Enter**.
3. Use the following command to run the Avaya Analytics™ Administration script:

   ```
   ccm release orca analytics
   ```

4. Enter the corresponding number to select the **Historical Reporting** option.
5. Select the **Historical reporting EASG availability** option by entering the corresponding number.
6. Select the **Disable EASG** option by entering the corresponding number.
7. In the **Historical Reporting disable easg availability?** field, type `y` or `n`.

   Typing `n` cancels the operation.

   Typing `y` displays a warning on the console.
8. At the prompt, **Do you want to disable EASG [y/n]** type `y` or `n`.

   Typing `n` cancels the operation.

9. To go back to the previous page, type `b` and press **Enter**.
10. To leave the current page, type `q` and press **Enter**.
11. To return to the main menu, type `m` and press **Enter**.

Checking EASG availability for Historical Reporting

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. To switch to the root user, type `su` and press **Enter**.
3. Use the following command to run the Avaya Analytics™ Administration script:

   ```
   ccm release orca analytics
   ```

4. Enter the corresponding number to select the **Historical Reporting** option.
5. Select the **Historical reporting EASG availability** option by entering the corresponding number.
6. Select the **Check EASG availability** option by entering the corresponding number.
7. In the **Historical Reporting check easg availability?** field, type `y` or `n`.

   Typing `n` cancels the operation.

   The current EASG settings display on the console.
8. To go back to the previous page, type `b` and press **Enter**.
9. To leave the current page, type `q` and press **Enter**.
10. To return to the main menu, type `m` and press **Enter**.

Monitoring security certificates

About this task

Service engineers can use the `Analytics Administration` script in Cluster Control Manager (CCM) to monitor whether:

- Security certificates are configured correctly.
- The Avaya Analytics™ internal certificates (For example, the internal certificates between Input Adapter, Kafka, and Measure Processors) are due to expire.

You can also configure Avaya Analytics™ to raise an alarm at least 8 weeks before the expiry date of these internal certificates.

Complete the following procedure to monitor the Analytics certificate expiry date:

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. To select the **General** option, enter the corresponding number.
6. To select the **Analytics Certificate expiry date** option, enter the corresponding number.
7. In the **Proceed with Analytics Certificate expiry date** field, type `y`.

   Typing `n` cancels the operation.

   Avaya Analytics™ confirms the date on which the certificates are due to expire.
8. Return to the previous page by entering `b`.
9. Quit the current page by entering `q`.
10. Return to the main menu by entering `m`.

Creating metadata backups

About this task

You can create a backup of a historical reporting database for storing your reports. This backup is required to restore the data later during troubleshooting or upgrades.

Note:

You must restart the mstr-srv pod after a successful restore of the historical metadata.

Historical reporting metadata backups are saved to the analyticsdb-node pods.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. To select the **Backup Metadata** option, enter the corresponding number.
6. To create a backup in the metadata node, select the **Backup Metadata** option by entering the corresponding number.

7. In the **Proceed with historical reporting metadata backup** option, enter `y`.

   The CCM displays the message that the backup is successful with the details of the location of the backup folder.
8. Restart the mstr-srv pod after a successful restore of the historical metadata.
9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.
12. To export the backups to CCM, select the **Export Backups** option by entering the corresponding number.
13. In the **Proceed with backup export** option, enter `y`.

    Entering `n` cancels the operation.

    The CCM console displays the path to the location where the backup is exported.
14. Return to the previous page by entering `b`.
15. Quit the current page by entering `q`.
16. Return to the main menu by entering `m`.
17. To view the list of Historical Reporting metadata backups, select the **List Backups** option by entering the corresponding number.
18. Return to the previous page by entering `b`.
19. Quit the current page by entering `q`.
20. Return to the main menu by entering `m`.
21. To restore the historical metadata, select the **Restore Metadata** option by entering the corresponding number.
22. In the **Proceed with Hstorical Reporting Metadata restore** field, enter `y`.

    Warning:

    Restoring the data from the backup wipes out all the current data and replaces the content from the backup.
23. In the **Specify the location of the backup file you want to restore** field, type one of the following and press **Enter**.

    Note:

    The values are case-sensitive.

    - `CCM`: If the file is located in the CCM machine.
    - `POD`: If the file is located on the analyticsdb-node pods.

24. If you enter CCM, in the **Enter the full path and file name of the backup file located on CCM** field, type the required details and press **Enter**.

    For example:`/home/cust/historical_md_backups/full_backup_2020_01_31_15_38_01.bkp`

25. If you enter POD, in the **Enter the name of the backup file you wish to restore from** field, type the name of the file and press **Enter**.

    The back up is restored to historical reporting node.

26. Return to the previous page by entering `b`.
27. Quit the current page by entering `q`.
28. Return to the main menu by entering `m`.

Viewing and removing metadata backups

About this task

Use this procedure to view a list of existing scheduled metadata backups and to remove any of these backups from the list.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`

4. Select **Historical Reporting** by pressing the corresponding number.
5. To select the **Backup Metadata** option, enter the corresponding number.
6. To select **Schedule Metadata Backup** option, enter the corresponding number.
7. To view the list of currently scheduled metadata backups, select the **List currently scheduled Metadata Backups** option, enter the corresponding number.
8. Return to the previous page by entering `b`.
9. Quit the current page by entering `q`.
10. Return to the main menu by entering `m`.
11. To remove a scheduled backup, select the **Remove a Scheduled Backup** option by entering the corresponding number.
12. In the **Proceed with backup schedule removal** field, enter `y`.

    The CCM console displays the list on incremental backups.

13. In the **Select which schedule you would like to remove** field, enter the number corresponding to the incremental backup you want to remove.
14. Return to the previous page by entering `b`.
15. Quit the current page by entering `q`.
16. Return to the main menu by entering `m`.

Scheduling a full metadata backup

About this task

You can schedule a full metadata backup based on your choice of time period, such as daily, weekly, and monthly. You can also choose the hour for the backup.

Before you begin

Stop all events coming to the Avaya Analytics™ database from Avaya Oceana®. For steps on stopping the events, see Stopping all events entering in to Avaya Analytics from Avaya Oceana on page 101.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. To select the **Backup Metadata** option, enter the corresponding number.
6. To create a backup in the Metadata node, select the **Schedule a Metadata Backup** option by entering the corresponding number.
7. In the **Proceed to backup schedule creation** field, enter `y`.

   Entering `n` cancels the operation.

   The Schedule a Full database backup page displays the following message:

   ```
   How often should the backup be executed [daily/weekly/monthly/
   yearly]
   ```

   Depending on your selection, follow the next steps:

For daily backups

8. In the **How often should the backup be executed [daily/weekly/monthly/yearly]** field, type `daily` and press **Enter**.

   The value is case sensitive.

9. To select the time in hour for the backup, in the **What hour of the day [0–23]** field, type the time and press **Enter**.

   The time selection is in 24–hour format. For example, if you type 2, then the backup happens at 2am.

10. To select the time in minute for the backup, in the **What minute of the hour [0–59]** field, type the time and press **Enter**.

    For example, if you type 2, then the backup happens at 2 minutes past 2 am.

11. Return to the previous page by entering `b`.
12. Quit the current page by entering `q`.
13. Return to the main menu by entering `m`.

For weekly backups

14. In the **How often should the backup be executed [daily/weekly/monthly/yearly]** field, type `weekly` and press **Enter**.

    The value is case sensitive.

15. To select the day of the week for the backup, in the **What day of the week (Sunday to Saturday) [0–6]** field, type the number of the day and press **Enter**.

    For example, to select Monday, type `1`.

16. To select the time in hour for the backup, in the **What hour of the day [0–23]** field, type the time and press **Enter**.

    The time selection is in 24-hour format. For example, if you type 2, then the backup happens at 2am.

17. To select the time in minute for the backup, in the **What minute of the hour** field, type the time and press **Enter**.

    For example, if you type 2, then the backup happens at 2 minutes past 2 am.

18. Return to the previous page by entering `b`.
19. Quit the current page by entering `q`.
20. Return to the main menu by entering `m`.

For monthly backups

21. In the **How often should the backup be executed [daily/weekly/monthly/yearly]** field, type `monthly` and press **Enter**.

    The value is case sensitive.

22. To select the day for the backup, in the **What day of the month [1–31]** field, type the number of the day and press **Enter**.

    For example, to select 2nd day of every month, type `2`

23. To select the time in hour for the backup, in the **What hour of the day** field, type the time and press **Enter**.

    The time selection is in 24-hour format. For example, if you type 2, then the backup happens at 2 am.

24. To select the time in minute for the backup, in the **What minute of the hour** field, type the time and press **Enter**.

    For example, if you type 2, then the backup happens at 2 minutes past 2 am.

25. Return to the previous page by entering `b`.
26. Quit the current page by entering `q`.
27. Return to the main menu by entering `m`.

For yearly backups

28. In the **How often should the backup be executed [daily/weekly/monthly/yearly]** field, type `yearly` and press **Enter**.

    The value is case sensitive.

29. To select the month for the backup, in the **What month of the year (January to December) [1–12]** field, type the number corresponding to the month and press **Enter**.

    For example, for selecting October, type `10`.

30. To select the day for the backup, in the **What day of the month [1–31]** field, type the number of the day and press **Enter**.

    For example, to select 2nd day of every month, type `2`

31. To select the time in hour for the backup, in the **What hour of the day** field, type the time and press **Enter**.

    The time selection is in 24-hour format. For example, if you type 2, then the backup happens at 2am.

32. To select the time in minute for the backup, in the **What minute of the hour** field, type the time and press **Enter**.

    For example, if you type 2, then the backup happens at 2 minutes past 2 am.

33. Return to the previous page by entering `b`.
34. Quit the current page by entering `q`.
35. Return to the main menu by entering `m`.

Package migration

Viewing a list of imported migration packages

About this task

Using the package migration option in the Cluster Control Manager (CCM) console, you can migrate historical reporting projects from one environment to another. You can also use this option to troubleshoot issues, such as changing metrics, attributes or reports on an installed environment.

Before you begin

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. To select the **Package Migration** option, enter the corresponding number.
6. To view the list of imported migration packages, select the **List Imported Migration Packages** option by entering the coresponding number.

   The CCM console displays the list of imported migration packages.
7. Return to the previous page by entering `b`.
8. Quit the current page by entering `q`.
9. Return to the main menu by entering `m`.

Importing migration packages

About this task

Using the package migration option in the Cluster Control Manager (CCM) console, you can migrate historical reporting projects from one environment to another. You can also use this option to troubleshoot issues, such as changing metrics, attributes or reports on an installed environment.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.

5. To select the **Package Migration** option, enter the corresponding number.
6. In the **Import Migration Package** option, enter the corresponding number.
7. To proceed with the package migration, in the **Proceed package migration config** option, enter `y`.

   Entering `n` cancels the process.
8. In the **Please confirm the package migration file (.mmp) is saved to the default location**, type `y`
9. In the **Please enter the migration package file (.mmp) name including the path** field, type the package file name and the path of the file.

   For example, `/home/cust/historicalPackageMigration/avaya_analytics_u_1_4.0.1_21052020.mmp`
10. In the **Please enter the package migration type** field, type one of the following:

    - `P`

      The Release package contains the complete Historical Reporting project.
    - `U`

      The Update package contains the updated historical reports.
    - `C`

      The Custom package contains the custom historical reports.
11. Return to the previous page by entering `b`.
12. Quit the current page by entering `q`.
13. Return to the main menu by entering `m`.

Replacing migration package metadata

About this task

Using the package migration option in the Cluster Control Manager (CCM) console, you can migrate historical reporting projects from one environment to another, such as from development to production. You can also use this option to troubleshoot issues, such as changing metrics, attributes or reports on an installed environment.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.

5. To select the **Package Migration** option, enter the corresponding number.
6. Select the **Migration Package — Replace Metadata** by entering the corresponding number.
7. In the **Proceed metadata replacement config** field, enter `y`.

   Entering `n` cancels the process.
8. In the **Please confirm you saved the new Metadata.sql file to a location on CCM** field, enter `y`.

   Entering `n` cancels the process.
9. In the **Please enter the Metadata.sql name including the path** field, type the required details

   For example, `/home/cust/historicalPackageMigration/avaya_analytics_u_1_4.0.1_21052020.sql`
10. In the **Please enter the Historical Reporting Administrator User password for the new Metadata.sql** field, type the required password.
11. At the prompt, type the current Historical Reporting administrator password for your system.
12. In the **Please enter the Historical Reporting Database Metadata User password**, type the Historical Reporting database metadata password for your system
13. Return to the previous page by entering `b`.
14. Quit the current page by entering `q`.
15. Return to the main menu by entering `m`.

Rolling back an imported migration package

About this task

Using the package migration option in the Cluster Control Manager (CCM) console, you can migrate historical reporting projects from one environment to another. You can also use this option to troubleshoot issues, such as changing metrics, attributes or reports on an installed environment.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. To select the **Package Migration** option, enter the corresponding number.

6. Select the **Migration Package - Rollback an imported migration package** option by entering the corresponding number.
7. To proceed with the package migration, in the **Proceed package migration config** field, enter `y`.

   Entering `n` cancels the process.
8. In the **Please enter the imported migration package name including the path you would like to roll back** field, type the package file name and the path of the file.

   For example, `/home/cust/historicalPackageMigration/avaya_analytics_u_1_4.0.1_21052020.mmp`

   The imported file gets deleted from the backup.
9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.

Checking disk space usage

About this task

Use this procedure to view the disk space sage on the Historical Reporting pods.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. To select the **Logging** option, enter the corresponding number.
6. To select the **Logging Volumes Available** option, enter the corresponding number.
7. In the **Proceed to disk space usage reporting** field, enter `y`.

   Entering `n` cancels the operation.

   The CCM console displays the list of the Historical Reporting pods.
8. In the **Select which pod to check logging disk usage** option, type the name of the pod and press **Enter**.

   For example, `mstr-srv-75fc6954bc-948sn`.

   The CCM console displays the used and available volume for the pod.

9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.

Copying Historical Reporting logs to Cluster Control Manager

About this task

Use this procedure to copy logs from Historical Reporting pods to Cluster Control Manager.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. To select the **Logging** option, enter the corresponding number.
6. To copy the logs, select the **Copy Historical Reporting Logs to ccm** option by entering the corresponding number.
7. In the **Proceed with copying logs to CCM** field, type `y`.

   Entering `n` cancels the operation.

   The CCM console displays the list of available pods with their log files.
8. In the **Select which pod to copy logs to CCM from** field, enter the name of the pod.

   For example, `mstr-srv-75fc6954bc-948sn`.
9. In the **Select which log file to copy to CCM**, enter the log file name.

   The CCM console displays the message that the selected log file is copied to the `historicalReportingLogs` directory in CCM.
10. Return to the previous page by entering `b`.
11. Quit the current page by entering `q`.
12. Return to the main menu by entering `m`.

Deleting logs from Historical Reporting pods

About this task

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Historical Reporting** by pressing the corresponding number.
5. To select the **Logging** option, enter the corresponding number.
6. In the **Proceed to pod logs deletion** option, type `y` and press **Enter**.

   Entering `n` cancels the operation.

   The CCM console displays a list of the available Historical Reporting pods.
7. In the **Select which pod to delete logs from** field, enter the name of the pod.

   For example, `mstr-srv-75fc6954bc-948sn`.

   The CCM console displays a list of the log files and their respective volume details.
8. In the **Select which log file to delete** field, enter the name of the log file.

   For example:`/mnt/log/mstr/scripts.log/mnt/log/mstr/DSSErrors.log`
9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.

Checking High Availability status

About this task

You can use the post-install script to check the High Availability (HA) status of your deployment. The status check runs at a regular interval and detects issues in a primary-primary or standby-standby scenario and automatically fixes issues, if any.

Before you begin

You must have a deployment scenario with HA.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.

2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. To select the **Geo/High Availability** option, enter the corresponding number.
5. To select the **High Availability Options** option, enter the corresponding number.
6. To run a HA status check, select the **Configure scheduled Pod HA status checks** option by entering the corresponding number.
7. In the **Proceed to scheduled POD HA status checks config** field, type `y`.

   Entering `n` cancels the operation.
8. To set the interval at which you want the HA check, in the **Please enter the interval** field, type the value in minutes and press **Enter**.

   You can enter any value between 15 to 60 minutes.
9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.
12. To remove the HA status check, select the **Remove a scheduled Pod HA status check** option by entering the corresponding number.
13. In the **Proceed with scheduled Pod HA status check removal** field, type `y`.

    Entering `n` cancels the operation.
14. Return to the previous page by entering `b`.
15. Quit the current page by entering `q`.
16. Return to the main menu by entering `m`.

Configuring Okta as Identity Provider for Historical Reporting

About this task

Use this procedure to create a SAML integration application with Okta for SAML authentication with Historical Reporting. The terms used here are specific to Okta as an Identity Provider (IDP).

Before you begin

You need a unique identifier of the web application to be recognized by the IDP. For example, `AnalyticsWebHome`.

You need the URL that IDP sends and receives SAML requests and responses. For example, `https://cluster.fqdn/AvayaAnalytics`.

Procedure

1. Log in to the **Okta** application.
2. Select **Create New App** to create a new SAML application.
3. In **Create a New Application Integration**, select **Platform** > **Web**.
   a. Select **Sign on method**.
   b. Click **Create**.
4. In **General Settings**, add an **App name for the SAML integration**.
5. In **SAML Settings** page, select the following:
   a. Select the **Single sign on URL property**, add the URL from which an IDP sends and receives SAML requests and responses.

      For example, `https://cluster.fqdn/AvayaAnalytics` and append on `/saml/SSO`
   b. In the **Audience URL** (SP Entity ID) property, add the unique identifier of the web application to be recognized by the IDP.

      For example, `AnalyticsWebHome`.
   c. Add **Additional attribute statements** as needed. For example,

      `Name Name format Value`

      `DisplayName unspecified user.displayName`

      `Email unspecified user.email`
   d. Add **Group attribute statements**. For example,

      `Name Name format Value`

      `Groups unspecified Matches regex .*` (matches all groups)
6. Preview the SAML Assertion and click **Next**.
7. On Okta feedback, click **Finish** to complete the SAML application.
8. Go to **Assignments** > **Groups** > **Assign**.

   If the web admin page is configured to be protected by SAML, groups must be assigned to this role on Okta.
   a. On the Assign SAML application to Groups, select groups and click **Assign**.
   b. Click **Done**.
9. To download the IDPMetadata for your SAML application, on the **Sign On** tab, click the **Identity Provider metadata** link. This generates the metadata.
10. Save the generated metadata as `IDPMetadata.xml`.

11. Copy the `IDPMetadata.xml` to CCM using WinSCP.

    This file is used for SAML configuration on Historical Reporting.

Configuring Active Directory Federation Services as Identity Provider for Historical Reporting

About this task

You can configure a relying party trust for SAML authentication with Historical Reporting. The terms used here are specific to Active Directory Federation Services (ADFS) on Windows Server 2016 as an Identity Provider (IDP).

Before you begin

You must have the unique identifier of the web application to be recognized by the IDP. For example, `AnalyticsWebHome`.

URL from which IDP sends and receives SAML requests and responses. For example, `https://cluster.fqdn/AvayaAnalytics`.

Procedure

1. Log in to **Active Directory Federation Services (ADFS)** server and open the **ADFS Management UI**.
2. Click **Add Relying Party Trust** to add a new relying party trust.
3. On the Welcome page, select the **Claims aware** option and click **Start**.
4. In the Select Data Source window, select **Enter data about the relying party manually** option and click **Next**.
5. In the Specify Display Name window, add a name for this relying party trust and click **Next**.
6. In the Configure Certificate window, leave blank if not applicable and click **Next**.
7. In the Configure URL window, click **Enable support for the SAML 2.0 WedSSO protocol** option, and then add the URL from which IDP sends and receives SAML requests and responses. Click **Next**.
8. In the Configure URL window, add the URL from which IDP sends and receives SAML requests and responses. Click **Next**.

   For example, `https://cluster.fqdn/AvayaAnalytics` and append on `/saml/SSO`
9. In the Configure Identifiers window, add the unique identifier of the web application to be recognized by the IDP and click **Next**.

   For example, `AnalyticsWebHome`.
10. In the Choose Access Control Policy, select the required access control policy and click **Next**.

11. In the Ready to Add Trust window, confirm your settings and click **Next**.
12. In the Finish window, check the configure claims issuance policy for this application and click **Close**.
13. In **Edit Claim Issuance Policy** for your relying party trust, click **Add Rule**.
14. In the Select Rule Template window, select **Send LDAP Attributes as Claims** from the dropdown and click **Next**.
15. In the Configure Rule window, add the rule name. For example, Name ID.

    Identify the Attribute store as Active Directory from the drop down menu.
16. To map the SAM-Account-Name to Name ID of the user from the LDAP Attribute, select or type **SAM-Account-Name**, and for the Outgoing Claim type, select **Name ID**.
17. To map the Distinguished name of the user from the LDAP Attribute, select or type **distinguishedname**, and for the Outgoing Claim type, select or type **DistinguishedName**.
18. To map the E-Mail Address of the user from the LDAP Attribute, select or type **E-Mail-Addresses**, and for the Outgoing Claim type, select or type **EMail**.
19. To map the DisplayName of the user from the LDAP Attribute, select or type **Display-Name**, and for the Outgoing Claim type, select or type **DisplayName**.
20. To map the Groups name of the user from the LDAP Attribute, select or type **Token-Groups – Unqualified Names**, and for the Outgoing Claim type, select or type **Groups**.

    Note:

    The following strings must match with attributes as provided when configuring SAML using post-install steps:

    - DistinguishedName must match the User display name attribute
    - EMail must match the User email address attribute
    - DisplayName must match the User distinguished name attribute
    - Groups must match the User group attribute
21. To add a SAML logout endpoint, do the following:
    a. Select **Properties** to edit the relying party trust.
    b. Go to the **Endpoints** tab.
    c. Select **Add SAML**,
    d. On the Add Endpoint window, select **SAML logout** from dropdown list.
    e. For Binding select **POST** from dropdown list.
22. For the trusted URL, add the URL from which IDP sends and receives SAML requests and responses. Click **Ok**.

    `https://cluster.fqdn/AvayaAnalytics` and append on `/saml/SingleLogout`
23. To export the IDP `metadata.xml`, enter the URL of the ADFS server.

`https://adfsServer.domain.com/FederationMetadata/2007-06/FederationMetadata.xml`

The Federation Metadata.xml file downloads to your default downloads location.

24. Rename the downloaded IDP metadata file as `IDPMetadata.xml` and copy it using WinSCP file to the CCM.

Tracing of Agents for Historical Reporting

Tracing an agent records the activities of an agent, state changes, the time when these events occurred and the duration. Agent Trace reporting can help supervisors evaluate how agents are spending their time.

Viewing the verbose level set for Trace Processor

About this task

Use this procedure to view the current verbose level set for orca-trace-measure-processor. The verbose levels are:

- 1 (NONE) – no events are processed for trace enabled agents.
- 2 (USER) – only user events are processed for trace enabled agents.
- 3 (INTERACTION) - only call related events are processed for trace enabled agents.
- 4 (ALL)(Default) – both user and call related events are processed for trace enabled agents.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. To select the **Trace Report Configuration** option, enter the corresponding number.
6. To select the **Show current Verbose Level** option, enter the corresponding number.

   The current verbose level for the trace processor is displayed on the screen.
7. Return to the previous page by entering `b`.
8. Quit the current page by entering `q`.
9. Return to the main menu by entering `m`.

Configuring the verbose level for Trace Processor

About this task

Use this procedure to set the verbose level for the orca-trace-measure-processor. The verbose levels are:

- 1 (NONE) – no events are processed for trace enabled agents.
- 2 (USER) – only user events are processed for trace enabled agents.
- 3 (INTERACTION) - only call related events are processed for trace enabled agents.
- 4 (ALL)(Default) – both user and call related events are processed for trace enabled agents.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. To select the **Trace Report Configuration** option, enter the corresponding number.
6. To select the **Configure Verbose Level** option, enter the corresponding number.

   The current verbose level for the trace processor is displayed on the screen.
7. In the **Proceed to configure verbose level** field, type `y`.

   Entering `n` cancels the operation.
8. In the **Please enter the value of the verbose level that you wish to use (from 1 to 4)**, enter a new value for the verbose level.
9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.

Start an Agent Trace

About this task

You must create a user group in ACM named AgentTraceGroup and assign the agents who are required to be traced to this group. The **`orca-trace-measure-proc`** service traces all agents assigned to this group, who are logged-in to Avaya Workspaces.

Procedure

1. Log in to the Avaya Control Manager console as an administrator.
2. On the Avaya Control Manager webpage, click **Groups**.

3. In the navigation pane, click **Location** > **Parent group**.
4. Click **Add ACM Group**.
5. In the **ACM Group Details** section, do the following:
   a. In the **Entity Type** field, select `User`.
   b. In the **Name** field, enter the name `AgentTraceGroup`.
   c. In the **Description** field, enter a description of the entity type.
   d. In the **Systems** field, select `Oceana - Analytics`.
6. Click the **Entity Assignment** tab.
7. In the **Available** section, select the agents to be traced and move to **Selected** section by clicking on arrow icon.
8. Click **Save**.

Stop an Agent Trace

About this task

To stop tracing an agent, you must remove the agent from the AgentTraceGroup user group in ACM. When an agent is unassigned from AgentTraceGroup, tracing stops for that agent immediately, even when the agent is still logged onto Avaya Workspaces and handling calls. Stopping agent trace does not delete the trace records for that agent.

Procedure

1. Log in to the Avaya Control Manager console as an administrator.
2. On the Avaya Control Manager webpage, click **Groups**.
3. In the navigation pane, go to **AgentTraceGroup** user group.
4. Click the **Entity Assignment** tab.
5. In the **Selected** section, select the agents to stop tracing and move to the **Available** section by clicking on arrow icon.
6. Click **Save**.

SAML implementation

To integrate SAML for existing local/LDAP reporting users, follow the procedure *Mapping SAML users to Historical Reporting local users or LDAP users* in the *Deploying Avaya Analytics™ guide* .

To integrate SAML for new reporting users, follow the procedure *Creating a new reporting user and map to SAML user for Historical Reporting* in the *Maintaining and Troubleshooting Avaya Analytics™ guide*.

Note:

When adding a new local reporting user, you do not need to assign a role (Advanced, Basic, or Consumer) to these users. The role is inherited for the Active directory or the Identity Provider when SAML is configured.

If the Web Single Sign-on for importing SAML users is not enabled, you must manually create new local users in Historical Reporting. This is not enabled by default. To configure automatic import/creation of SAML users, see the *Configure Web Single Sign-on for importing SAML users automatically* section in the *Maintaining and Troubleshooting Avaya Analytics™ guide*.

Creating a new reporting user and map to SAML user for Historical Reporting

About this task

This task outlines how to create a new reporting user to map to SAML user for Avaya Analytics™ reporting.

Note:

Only a user configured with a supervisor role in Avaya Control Manager can get access to the historical reports.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. To select the **Historical Reporting** option, enter the corresponding number.
5. To select the **SAML** option, enter the corresponding number.
6. To select **Map SAML users to Historical reporting**, enter the corresponding number.
7. To create a local user for Historical reporting, select the **Create a Historical Reporting local user (valid ACCCM Reporting Supervisor)** option by entering the corresponding number.
8. In the **Proceed with user creation** option, type `y`. Entering `n` cancels the operation.
9. At the prompt, enter **Administrator password**.
10. Enter **Username** of a valid supervisor in Avaya Control Manager.
11. Enter **New username** and **Password**.
12. Confirm the password.
13. In the **New Users Full Name** field, type the full name of the new user.

    The new user is created.

14. Return to the previous page by entering `b`.
15. Return to the main menu by entering `m`.
16. Quit the current page by entering `q`.

Mapping new SAML user in Historical Reporting

About this task

This procedure outlines how to map a new reporting user with a SAML user in Avaya Analytics™ reporting.

Note:

Only a user configured with a supervisor role in Avaya Control Manager can get access to the historical reports.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. To select the **Historical Reporting** option, enter the corresponding number.
5. To select the **SAML** option, enter the corresponding number.
6. To select the **Map SAML users to Historical reporting**, enter the corresponding number.
7. To select the **Map a single SAML users to Historical reporting** option, enter the corresponding number.
8. Enter **Username** of the local user.
9. Enter **Name ID** of the SAML user from the SAML assertion.

   The entered SAML user is mapped with Historical reporting local user. The CCM console displays a message for the SAML user confirming user login and trusted login.
10. Return to the previous page by entering `b`.
11. Return to the main menu by entering `m`.
12. Quit the current page by entering `q`.

Configuring Web Single Sign-on for importing SAML users automatically into Historical Reporting

About this task

This procedure outlines how to configure Web Single Sign-on configuration to enable automatically import of SAML users for Historical Reporting.

Before you begin

- Permission is required from the customer to connect to their CCM using the developer tool.
- A node port needs to be opened on the customer's CCM.
- The Historical Reporting Administrator credentials are available.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. Open a nodeport to the mstr-srv service using the following command:

   ```
   sudo kubectl expose service -n mstr mstr-srv --type=NodePort --name=mstr-nodeport
   ```

4. Find the exposed service with following command:

   ```
   sudo kubectl get services -n mstr | grep mstr
   mstr-nodeport …. <none>        12345:31127/TCP, …..
   ```

5. Open developer tool. On the **Folder List** menu right-click and select the **New Project Source**.
6. For project source enter a **String Name**.
7. For Server name, enter the **FQDN / IP** of the cluster.
8. Update **Port number** for the exposed port number created on CCM.
9. Click **Ok**
10. In **Folder List** menu, right-click on the **New project source** and select **Connect to Project Source**.
11. Enter **Administrator credentials** and click **Ok**.
12. In the **Folder List** menu, right-click on the **New project source** and select **Configure Intelligence Server**.
13. In a new window in the **Categories** menu, open **Web Single Sign-on category**.
14. Select the check boxes to **Allow user to log on**, **Import user at logon**, and **Synch user at logon**.
15. Click **Ok**.
16. In the **Folder List** menu, right-click on the **New project source** and select **Disconnect from project source**.
17. Log in to the Cluster Control Manager (CCM) console as the customer user.
18. Switch to being the root user by entering the command `su`.
19. Close the node port on the cluster using the following command:

    ```
    sudo kubectl delete service -n mstr mstr-nodeport
    ```

Enable debug level logging for troubleshooting SAML

About this task

Use this procedure to enable debug level logging to troubleshoot SAML.

Before you begin

SAML must be enabled.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To list the Historical reporting pods, use the following command:

   ```
   sudo kubectl get pods --namespace mstr
   ```

4. To enter into the mstr-web pod, use the following command:

   ```
   sudo kubectl exec -it --namespace mstr mstr-web-podNumber -- /bin/
   bash
   ```

5. The SAML logs are in the `/opt/tomcat/logs/SAML/ directory`.
6. Edit the SAML logging properties file on the following path:

   ```
   vi /opt/tomcat/webapps/AvayaAnalytics/WEB-INF/classes/
   log4j2.properties
   ```

7. Change SAML loggers to debug level. For example, auth.saml and opensaml.

   ```
   logger.b.name = .auth.saml
   logger.b.level = debug
   logger.d.name = org.opensaml
   logger.d.level = debug
   ```

8. Click **Save** and close.
9. To restart the web server process, use the following command:

   ```
   /opt/tomcat/bin/shutdown.sh
   ```

10. To confirm the process is stopped and restarted, use the following command:

    ```
    ps -ef | grep tomcat
    ```

11. Repeat steps 6 through 10 to change the log level to error.

    ```
    logger.d.name = org.opensaml
    logger.d.level = error
    ```

Opening a node port on the Analytics database

About this task

Use this procedure to open a node port on the Avaya Analytics™ database. Once the node port is opened, you can choose to patch that port to a specific port in the range 30000-32767.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. Select the **Open a port on the database** option by entering the corresponding number. A port is opened on the database.
6. In the **Confirm that you want to patch the database port xx to specific port in range 30000-32767** field, type `y`.

   Entering `n` cancels the operation.
7. In the **Please enter port in range 30000-32767** field, enter a new value for the port.
8. Return to the previous page by entering `b`.
9. Quit the current page by entering `q`.
10. Return to the main menu by entering `m`.

Grafana reports for monitoring Crunchy database

Avaya Analytics™ 4.3 release supports monitoring the Crunchy database using the Grafana reporting system.

The following reports are supported:

- CrunchyData Clusters Backup Details: This report looks into the pgBackrest repository and provides information on full and incremental backups, including size and time taken to create. It also includes information on WAL file archiving.
- CrunchyData CRUD Operations Details: This report tracks database usage from the Create, Read, Update, and Delete statements. It helps to identify operations that are costly on the system.
- CrunchyData Database Activity Details: This report details many items, including active connections, idle time, DB cache usage, locks, commits, and replication lag.

- CrunchyData Database Query Statistics: This report monitors queries on the system and helps to isolate particular queries that take time or consume a lot of system memory.
- CrunchyData Pods Resource Usage: This report provides kubernetes-level information narrowed down for Crunchy-specific pods. It tracks memory, CPU, and disk usage for Crunchy pods.
- CrunchyData Postgresql Service Health: This report shows an overview of the system covering resource usage and system errors.

You can also view the PostgreSQL and Crunchy Container Dashboard in the Grafana reporting system.

Note:

You can use these reports to monitor all the Crunchy instances within the kubernetes cluster. By default, there is analyticsdb and authdbservice for reporting. You must select which instance you are interested in from the drop-down menu at the top of each report. You can also select the individual pod or database from a drop-down menu as per your requirement.

Related links

Accessing Grafana reporting system on page 96

Accessing Grafana reporting system

About this task

Through this procedure, you can use the Grafana reporting system to access the reports and monitor the crunchy database. You need to query the CCM to find the IP address, username, and password of the Grafana reporting system and then use these credentials to log in to the Grafana reporting system.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To find the keepalived-vip **IP address** for Grafana access, run the following command:

   ```
   kubectl get cm common-services-kube-keepalived-vip -o yaml | grep gateway
   ```

   An IP Address is displayed, such as:

   ```
   10.134.139.133: istio-system/istio-ingressgateway
   ```

4. To find **Username**, run the following command:

   ```
   kubectl get secret --namespace default common-services-grafana -o
   jsonpath="{.data.admin-user}" | base64 --decode; echo
   ```

   The default username is admin.

5. To find **Password**, run the following command:

   ```
   kubectl get secret --namespace default common-services-grafana -o
   jsonpath="{.data.admin-password}" | base64 --decode; echo
   ```

6. Open a webbrowser and enter URL **https://<vip>/grafana**. Substitute the above-mentioned IP address in place of <vip> in the URL and click **Enter**.
7. Enter the above-mentioned Username and Password in the Grafana reporting system webpage.

   You are logged in to the Grafana reporting system.
8. Click **Dashboards** > **Manage** to access crunchy database reports.

Related links

Grafana reports for monitoring Crunchy database on page 95

Enabling or disabling Zero or Empty Row Suppression parameter

About this task

You can enable or disable the Zero or Empty Row Suppression parameter from the analytics script during installation or upgrade.

Procedure

1. Log in to Cluster Control Manager as a user.
2. To switch to the root user, type `su` and press `Enter`.
3. To run the Avaya Analytics™ Administration script, use the following command:

   ```
   ccm release orca analytics
   ```
4. Enter a number corresponding to the **Deployment** option.
5. Enter a number corresponding to the **Prevent writing empty intervals to the database** option configuration.
6. Select the **Configure the prevent empty intervals writing to the database** parameter to configure this parameter for some of the interval processors separately.
7. Select the **Configure the prevent empty intervals writing to the database** parameter for all interval processors option to configure this parameter for all of the interval processors.

Chapter 3: Maintaining Avaya Analytics™ on Avaya Common Sevices

Gracefully powering off a cluster

About this task

Use this procedure to gracefully shut down your solution cluster.

Warning:

- If you shut down your cluster using a different method than what is described in this procedure, file corruption can occur.

Before you begin

- You must plan a maintenance window to perform this task.
- Stop call events across the solution.
  - To stop all Avaya Oceana® traffic for an Avaya Analytics™ cluster, run the `kubectl scale deployment orca-ref-input-adaptor --replicas=0` command using an account with root privileges.
- Back up Common Services using the `ccm backup` command. If you do not perform a backup, you risk a full reinstallation of the solution.
- Back up product application data as described in your solution documentation.

Procedure

1. Log into Cluster Control Manager.
2. Check the health of the pods:
   a. Enter `k get pods -A`
   b. If any pod is listed as other than *Running* or *Completed*, do not proceed any further. Contact Avaya Support.
   c. Enter `ccm status --health`
   d. Check that all products are listed as *Healthy*.
   e. Check that historical reporting is working.
   f. Check that the supervisor desktop shows data.

3. Create a metadata backup:
    a. Enter the following commands:

```
ccm release orca analytics
    Historical Reporting
        Backup Metadata
            Backup Metadata
            Export Metadata
```

    b. Wait for the command to complete before continuing.
    c. Note the name and location of the backup.
    d. Use winscp or a similar tool to copy that backup off ccm to a safe location.
4. Create a database backup:
    a. Enter the following commands:

```
ccm release orca analytics
    Database
        Database Backup
            Remote Backups
```

    b. Wait for the command to complete before continuing.
    c. Verify that the new backup file exists on the external backup server.
5. Identify which node runs the registry-pod:
    a. Enter `kubectl get pods -n image-registry -o wide`
    b. Note which master node has the registry-pod.
6. Stop the data flow to the cluster:
    a. Enter `kubectlscale --replicas=0 deployment orca-ref-input-adaptor`
    b. Verify that the *input ref-adaptor* pod has stopped by entering `kubectl get pods | grep orca-ref-input-adaptor` until no *orca-ref-input-adaptor* pods are shown.
7. Create a CCM backup:
    a. Enter `ccm backup`
    b. Wait for the command to complete before continuing.
    c. Verify that the new backup file exists on the external backup server.
8. Stop the authorization database:
    a. Gracefully shutdown the Common Services database by entering `pre-infra-upgrade`
    b. Wait for the command to complete before continuing.
9. Determine the node roles:
    a. Enter `ccm version -k`
    b. Determine role of each node (`worker` or `controller-worker`).

c. Make a note of the node roles as you need them during later procedures.

10. Determine which nodes contain a second disk and which are diskless.

    a. Run the `checkInfra -Sd | grep LVM_THIN` command.

```
[cust@flex190-129 ~]$ checkInfra -Sd | grep LVM_THIN
| pool_sds              | flex190-132.dr.example.com | LVM_THIN
| vg_sds/sds_thinpool |   341.00 GiB |    464.76 GiB |
True          | Ok    |            |
| pool_sds              | flex190-133.dr.example.com | LVM_THIN
| vg_sds/sds_thinpool |   341.00 GiB |    464.76 GiB |
True          | Ok    |            |
[cust@flex190-129 ~]$
```

    b. Note the disk status of each node. Determine which node is not listed. That node is the diskless node.

11. As root user, enter `kubectl get pods -n image-registry -o wide`
12. Make a note of the cluster node hosting the image registry.
13. Log in to vCenter as an administrator or with the account used to deploy the cluster.
14. Click on the **VMs and Templates** tab.
15. Power off the `worker` nodes:

    a. Right-click on the first `worker` node and click **Power** > **Shut Down Guest OS**.

    b. Check the VM status in vCenter. Wait until the node has powered down before continuing to the next node.

    c. Repeat these steps for the next `worker` node until all `worker` nodes are powered off.

16. Power off the `controller-worker` node without a second disk:

    a. Right-click on the `controller-worker` node without a second disk and click **Power** > **Shut Down Guest OS**.

    b. Check the VM status in vCenter. Wait until the node has powered down.

17. Power off the `controller-worker` nodes with second disks:

    Right-click on each controller-worker node containing a second disk and click **Power** > **Shut Down Guest OS**. You can power these off together, no wait time is needed between nodes.

18. Power off Cluster Control Manager:

    a. Click on the **VMs and Templates** tab.

    b. Locate and click on **Cluster Control Manager** in the folder you designated for the cluster.

    c. Right-click **Cluster Control Manager** and then click **Power** > **Shut Down Guest OS**.

Stopping all events entering in to Avaya Analytics™ from Avaya Oceana®

About this task

Use this procedure to turn off the input adapter to stop all traffic from Avaya Oceana® to Avaya Analytics™.

Before you begin

Complete this procedure during a maintenance window.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To stop data flow to the cluster, run the following command:

   ```
   kubectl scale --replicas=0 deployment orca-ref-input-adaptor
   ```
4. To verify that the input ref-input-adaptor pod is terminated, run the following command:

   ```
   kubectl get pods | grep orca-ref-input-adaptor
   ```

   Ensure that the command displays nothing.

Powering on a cluster

About this task

Use this procedure to gracefully power on your solution cluster. You can also use this procedure to power on after an unexpected power down. Complete this procedure during a maintenance window.

Important:

Do not manually delete virtual machine cluster nodes from vCenter.

Procedure

1. Log in to vCenter as an administrator or with the account used to deploy the cluster.
2. Click the **VMs and Templates** tab.
3. Power on Cluster Control Manager.
   a. Locate and click Cluster Control Manager in the folder you designated for the cluster.
   b. Right-click Cluster Control Manager and then click **Power** > **Power On**.
   c. Wait until the VM has initialized and you can log in to Cluster Control Manager before proceeding to the next step.

4. Enter `swversion`.
5. Locate and click on each node virtual machine in the folder you designated during your cluster deployment.
6. Power on the nodes. in the following order:
    a. Right-click on the master node that runs the registry-pod. Click **Power** > **Power On** and waitn for the node to allow login.
    b. Right-click on next `controller-worker` node containing a second disk and click **Power** > **Power On**.
    c. Wait 7 minutes.
    d. Right-click on the `controller-worker` node without a second disk and click **Power** > **Power On**.
    e. Right-click each `worker` node and then click **Power** > **Power On**.
7. Wait approximately 10 minutes for all nodes to power on before continuing.
8. If the cluster was shutdown using a gracefully power off:
    a. Enter **`post-infra-upgrade`** to start the authorization service database.
    b. Wait approximately 20 minutes for pods to come up and for storage to sync.
    c. During this time, monitor the pods for *0/1* status and only proceed for the next step when the status of the pods has not changed for about 10 minutes. Remaining pods show *0/1* or *INIT* as they are waiting on the *orca-ref- input-adaptor* pods to come up in next step.
9. For Avaya Analytics™, if you stopped all Avaya Oceana® traffic, restart data flow to the cluster.

    This step requires an account with root privileges.

    - For non-HA Avaya Analytics™, enter:

      ```
      kubectl scale --replicas=1 deployment orca-ref-input-adaptor
      ```

    - For HA Avaya Analytics™, enter:

      ```
      kubectl scale --replicas=2 deployment orca-ref-input-adaptor
      ```

    a. Wait for the command to complete before continuing. This can take 90 minutes.
    b. Monitor the status of the pods by entering the following command to list those pods that are not fully running as yet:

      ```
      k get pods -A | egrep -v "1/1|2/2|3/3|4/4|5/5|6/6|Completed"
      ```

    c. Wait for all pods to show *Running* or *Completed*. Check using `k get pods -A`
10. Verify the operational status of the cluster
    a. Run `ccm smoke-test`
    b. If `ccm smoke-test` fails, run it again in 10 minutes. If the test continuously fails for over an hour, contact your technical support representative.

11. Check the health of the pods:
    a. Enter `k get pods -A`
    b. If any pod is listed as other than *Running* or *Completed*, do not proceed any further. Contact Avaya Support.
    c. Enter `ccm status --health`
    d. Check that all products are listed as *Healthy*.
    e. Check that historical reporting is working.
    f. Check that the supervisor desktop shows data.

Monitoring the system status

About this task

When you are experiencing problems or just want to verify that the system is in a healthy state, you can check:

- The state of Cluster Control Manager
- Cluster and solution chart status
- Solution health
- Status of Kubernetes processes on each cluster node
- Individual pod status
- Alarm service status
- Avaya services registration of the system

While using this procedure to monitor your system status, if you find that your system is in a bad or unhealthy state, contact Avaya support personnel for assistance.

Procedure

1. Check the health of Cluster Control Manager by running the `ccmHealthCheck` command.

   Command output details:

   - Cluster Control Manager FQDN
   - Cluster Control Manager IP address
   - Cluster storage overall status
   - Number of alarms and types of alarms
   - Operations currently running
   - NTP status: Active, Disabled, or Not reachable.

     To fix this issue, check the Resolution output section. If the NTP server is unreachable, you will be directed to run the `ccmNetSetup` command.

- Node license state: Normal, Grace period, Restricted, or Unknown.

  If the license state is Grace period, Restricted, or Unknown, read the description section carefully for additional information.

  Important:

  After the 30-day licensing grace period elapses, the Common Services cluster is uninstalled. Product data is not preserved.

2. Check the status of the cluster and product by running the `ccm status` command.

   Command output details:

   - Cluster status: Not deployed, Deploying, Deployed, Upgrading, Backup/Restore Status
   - List of installed products: Chart, Platform, Release, Revision, Updated on Date, Status, Namespace
   - Environment
   - Data encryption at rest policy
   - Staged products and unstaged products
   - Scheduled backup configuration
   - Archive configuration
   - Cluster storage overall status
   - Number of alarms and types of alarms
   - CSP node license state
   - Operations currently running
3. Check the health status of the products by running the `ccm status --health` command.

   Command output details:

   - List of installed products: Chart, Platform, Release, Revision, Updated on Date, Status, Namespace, Health (Healthy or Unhealthy)
4. Check the status of processes on each cluster node by running the `ccm status --ps` command.

   Command output details includes the status of the following Kubernetes processes on each cluster node:

   - `sdsetcd.service`
   - `etcd.service`
   - `kube-apiserver.service`
   - `kube-controller-manager.service`
   - `kube-scheduler.service`

- keepalived.service
- flanneld.service
- kubelet.service
- kube-proxy.service
- containerd.service

5. Check the status of the pod by running the `ccm status --pod-details` command.

   Command output details include a list of staged products and the product pod status. The following example image shows the output of this command:

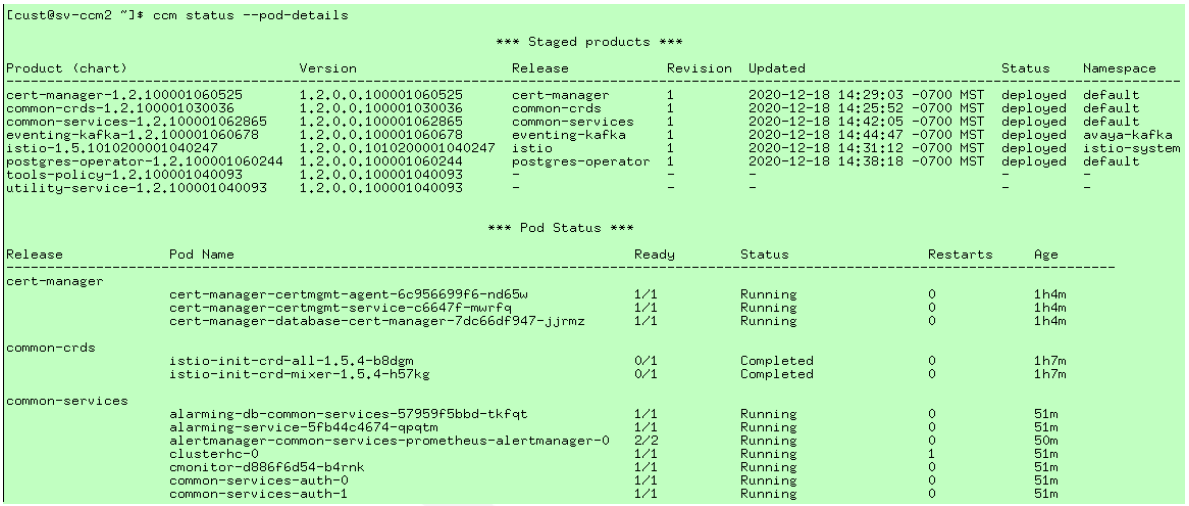

```
[cust@sv-ccm2 ~]$ ccm status --pod-details

                                                *** Staged products ***
Product (chart)                      Version                  Release            Revision  Updated                          Status    Namespace
-------------------------------------------------------------------------------------------------------------------------------------------------
cert-manager-1.2.100001060525        1.2.0.0.100001060525     cert-manager       1         2020-12-18 14:29:03 -0700 MST  deployed  default
common-crds-1.2.100001030036         1.2.0.0.100001030036     common-crds        1         2020-12-18 14:25:52 -0700 MST  deployed  default
common-services-1.2.100001062865     1.2.0.0.100001062865     common-services    1         2020-12-18 14:42:05 -0700 MST  deployed  default
eventing-kafka-1.2.100001060678      1.2.0.0.100001060678     eventing-kafka     1         2020-12-18 14:44:47 -0700 MST  deployed  avaya-kafka
istio-1.5.1010200001040247           1.2.0.0.1010200001040247 istio              1         2020-12-18 14:31:12 -0700 MST  deployed  istio-system
postgres-operator-1.2.100001060244   1.2.0.0.100001060244     postgres-operator  1         2020-12-18 14:38:18 -0700 MST  deployed  default
tools-policy-1.2.100001040093        1.2.0.0.100001040093     -                  -         -                                -         -
utility-service-1.2.100001040093     1.2.0.0.100001040093     -                  -         -                                -         -


                                                *** Pod Status ***
Release             Pod Name                                                  Ready   Status              Restarts   Age
-------------------------------------------------------------------------------------------------------------------------
cert-manager
                    cert-manager-certmgmt-agent-6c956699f6-nd65w              1/1     Running             0          1h4m
                    cert-manager-certmgmt-service-c6647f-mwrfq                1/1     Running             0          1h4m
                    cert-manager-database-cert-manager-7dc66df947-jjrmz       1/1     Running             0          1h4m

common-crds
                    istio-init-crd-all-1.5.4-b8dgm                            0/1     Completed           0          1h7m
                    istio-init-crd-mixer-1.5.4-h57kg                          0/1     Completed           0          1h7m

common-services
                    alarming-db-common-services-57959f5bbd-tkfqt              1/1     Running             0          51m
                    alarming-service-5fb44c4674-qpqtm                         1/1     Running             0          51m
                    alertmanager-common-services-prometheus-alertmanager-0    2/2     Running             0          50m
                    clusterhc-0                                               1/1     Running             1          51m
                    cmonitor-d886f6d54-b4rnk                                  1/1     Running             0          51m
                    common-services-auth-0                                    1/1     Running             0          51m
                    common-services-auth-1                                    1/1     Running             0          51m
```

   Expectations are as follows:

   - Status is 100% of expected containers running.
   - Restarts in the single digits.
   - Age indicates how long the pod has been running. If you see a recent restart and your cluster has been up for more than 24 hours, in the next step, verify that there are no active alarms for that service or pod.

6. Run the `ccm smoke-test` command.

   The following example shows the output of a successful smoke test:

```
$ ccm smoke-test
Executing Smoke Tests:

This may take a few minutes...
SDS Check:
    SDS PASSED
192.0.2.171
Test Results:

Cluster Check Test
PASS Smoke Test Pod Count Pass. 75/75 Pods Ready
```

```
Ping Test:
PASSED    kube_keepalived_vip UP
PASSED    keepalived_vip UP

Finished executing smoke tests!
```

7. Check the Common Services alarm by running the `ccm release common-services alarmctl -l alarmEvents` command.

   Note:

   Before running this command, make sure the common-service and alarming service are running.

   Command output shows details about alarms.

8. When the cluster is deployed, verify that the system is monitored by Avaya Services by running the `displaySEID` command.

   Command output details include the cluster ID, product name, version, instance, and SEID.

   If the system is not monitored or the cluster is not deployed, details are not provided.

Updating the network configuration of Cluster Control Manager and cluster nodes

About this task

Use this procedure to update the network-related settings of Cluster Control Manager and the cluster nodes:

- CCM network settings:
  - Host Fully Qualified Domain Name (FQDN)
  - Network domain search list
  - Public interface
    - IP address
    - Netmask
    - Gateway IP address
  - Primary Domain Name Server (DNS) server IP address
  - Secondary Domain Name Server (DNS) server IP address
  - Time zone and date format
  - Turn on or off the Network Time Protocol (NTP) server

  - IP/FQDN of primary NTP Server
  - IP/FQDN of secondary NTP Server
- Cluster node settings
  - Cluster node primary DNS server IP address
  - Cluster node secondary DNS server IP address
  - Network domain search list
  - IP/FQDN of cluster node primary NTP Server
  - IP/FQDN of cluster node secondary NTP Server
  - Time zone

Before you begin

- Perform a full backup of Common Services and application data.
- If you are changing the hostname or FQDN:
  - If Cluster Control Manager currently has a third-party signed identity certificate, you must have an updated identity certificate for Cluster Control Manager with the new hostname or FQDN.
  - Make sure the new hostname or FQDN is updated on your DNS with the correct Cluster Control Manager IP address.
- If you are changing the IP address:
  - Make sure the new IP address is updated on your DNS with the Cluster Control Manager FQDN.
  - Make sure the new IP address is compatible with the current/new gateway and netmask.

Procedure

1. Log in to Cluster Control Manager.
2. Run the `ccmNetSetup` command.

   At each step you can choose to change a value or keep the current value.

   Note:

   When changing the Cluster Control Manager IP address, netmask, and gateway IP address, make sure the new values work on the network where Cluster Control Manager is currently connected. If any of the values are not properly configured, you might lose connectivity to the Cluster Control Manager. If you do, you will need to log into the Cluster Control Manager console via vCenter and run this command again with the correct values.

If you changed the Cluster Control Manager hostname or FQDN:

3. If you are using a third-party signed certificate for Cluster Control Manager, upload the new identity certificate with the updated FQDN to the Certificate Manager running in the cluster.

   Upload this certificate only after you run the `ccmNetSetup` command.

4. Complete the Product Initial Registration process again.

Adding CPU and memory to a node

About this task

Use this procedure to increase node CPU and memory resources when not also increasing SDS disk size or a upgrading software. For example, when adding a product.

Before you begin

- You must plan a maintenance window to perform this task.

Procedure

1. Edit the solution configuration spreadsheet that you previously downloaded from the support site. Enter values to increase the CPU and memory resources.

   On the `cluster_config` tab, modify the `node-properties-cpu` and `node-properties-ram` values as required.

2. Power off the cluster nodes.
3. In vCenter, on each cluster node VM, edit the settings to increase the CPU and memory resources as required.
4. Power on the cluster nodes.
5. Log in to Cluster Control Manager using your customer account.
6. Wait approximately 30 minutes, then run `ccm smoke-test` on Cluster Control Manager and confirm that all tests pass. If `ccm smoke-test` fails, run it again in 10 minutes. If the test continuously fails for over an hour, contact your technical support representative.
7. **(Optional)** If you are adding a product as part of this resource change, continue with the product upgrade.

Result

After this procedure, the cluster nodes are running with the new CPU and memory resource allocations.

Related links

Powering on a cluster on page 101
Gracefully powering off a cluster on page 98

Reducing CPU and memory in a node

About this task

Use this procedure to reduce the node CPU and memory resources when not also increasing SDS disk size or a upgrading software. For example, when removing a product.

Before you begin

- You must plan a maintenance window to perform this task.

Procedure

1. Edit the solution configuration spreadsheet that you previously downloaded from the support site. Enter values to reduce the CPU and memory resources.

   On the `cluster_config` tab, modify the `node-properties-cpu` and `node-properties-ram` values as required.

2. Power off the cluster nodes.
3. In vCenter, on each cluster node VM, edit the settings to reduce the CPU and memory resources as required.
4. Power on the cluster nodes.
5. Log in to Cluster Control Manager using your customer account.
6. Wait approximately 30 minutes, then run `ccm smoke-test` on Cluster Control Manager and confirm that all tests pass. If `ccm smoke-test` fails, run it again in 10 minutes. If the test continuously fails for over an hour, contact your technical support representative.

Result

After this procedure, the cluster nodes are running with the new CPU and memory resource allocations.

Related links

Powering on a cluster on page 101
Gracefully powering off a cluster on page 98

Increasing SDS disk size for a node

About this task

This procedure increases the Software-Defined Storage (SDS) disk size for all disk nodes in the cluster. You cannot decrease SDS disk size.

Note:

vCenter does not enable you to resize SDS disks on virtual machines that have snapshots. To resize an SDS disk, remove any virtual machine snapshots before proceeding with this procedure.

Procedure

1. Resize the SDS disk in vCenter.

   This step can be done without powering off the cluster nodes.

2. Log in to Cluster Control Manager with an account that has root privileges.

3. On Cluster Control Manager, run the following command:

   ```
   kubectl exec --namespace=piraeus deployment/piraeus-op-piraeus-operator-cs-controller -linstor sp l
   ```

4. Note the nodes that have a storage pool of `pool_sds`.

5. Open the VMware Infrastructure client and log in to vCenter or the ESX host machine.

6. Right-click the cluster node VM from the output above.

7. Click **Edit settings**.

8. Select **Virtual Disk**.

   This is Hard Disk 2, with a size similar to the SDS Disk Size in current solution spreadsheet.

9. Enter the new size for the virtual hard disk based on the footprint summary table in the updated deployment spreadsheet.

10. Run the `ccm upgrade spec <solution spreadsheet name>.xlsx --infra` command.

11. Wait approximately 30 minutes, then run `ccm smoke-test` on Cluster Control Manager and confirm that all tests pass. If `ccm smoke-test` fails, run it again in 10 minutes. If the test continuously fails for over an hour, contact your technical support representative.

12. To verify the new disk size, run the `checkInfra -Sd` command and review the SDS capacity.

```
. . .
Storage pools:
+-------------------------------------------------------------------------------------------------------------------------------------------------------
--------------+
| StoragePool          | Node                   | Driver   | PoolName           | FreeCapacity | TotalCapacity | CanSnapshots |
State | SharedName |
|
=======================================================================================================================================================
==============|
| DfltDisklessStorPool | flex190-78.example.com | DISKLESS |                    |              |               | False        |
Ok    |            |
| DfltDisklessStorPool | flex190-79.example.com | DISKLESS |                    |              |               | False        |
Ok    |            |
| DfltDisklessStorPool | flex190-80.example.com | DISKLESS |                    |              |               | False        |
Ok    |            |
| DfltDisklessStorPool | flex190-94.example.com | DISKLESS |                    |              |               | False        |
Ok    |            |
| pool_sds             | flex190-78.example.com | LVM_THIN | vg_sds/sds_thinpool |   521.02 GiB |    599.70 GiB | True         |
Ok    |            |
| pool_sds             | flex190-79.example.com | LVM_THIN | vg_sds/sds_thinpool |   526.18 GiB |    599.70 GiB | True         |
Ok    |            |
+-------------------------------------------------------------------------------------------------------------------------------------------------------
--------------+
. . .
```

Result

After the upgrade completes, the cluster nodes have access to the increased SDS disk capacity.

Restoring or replacing a cluster node

About this task

Use this procedure to replace and join a missing or deleted cluster node into an existing cluster.

Before you begin

- Obtain the cluster node OVA. The version of the OVA must be the same version as the existing cluster nodes.
- Obtain the network settings (FQDN, IP address, gateway, and netmask) used by the original cluster node that is being recovered.
- Obtain the CPU and memory resource values used by the original cluster node that is being recovered.
- Obtain the SDS disk size (Disk 2) used by the original cluster node that is being recovered.
- Obtain the enrollment password configured on Cluster Control Manager.
- If DRS is configured within vCenter to support node anti-affinity, obtain the VM name to be recovered.

Procedure

1. Deploy the cluster node OVA using the same cluster node version, VM name, network settings, enrollment password, and VM resources (CPU, memory and SDS disk) as the original cluster node.

   Before the VM is powered on, use the **Edit Settings** option for the VM to modify its CPU and memory values.
2. If you enabled the HA audit during initial installation, on your vCenter Cluster object, edit the VM/Host Anti-Affinity rule created earlier and add the replacement VM name back into the rule.
3. Power on the cluster node VM.
4. Log in to Cluster Control Manager with your customer account.
5. On Cluster Control Manager, run `ccm restore node` to recover and join the missing cluster node into the cluster.
6. Wait approximately 30 minutes, then run `ccm smoke-test` on Cluster Control Manager and confirm that all tests pass. If `ccm smoke-test` fails, run it again in 10 minutes. If the test continuously fails for over an hour, contact your technical support representative.
7. If the old cluster node VM still exists (powered down), delete the VM using **Delete from Disk**.

Adding a node to a cluster

About this task

Use this procedure to add a worker node to a cluster.

Procedure

1. Edit the solution spreadsheet.

   Under the `cluster_config` tab modify the number of worker nodes (`worker_count`).

   Only worker nodes can be added to the cluster, unless you are migrating from a non-HA cluster to an HA cluster. The `master_count` must be 1 for a non-HA cluster and 3 for an HA cluster.

   Update the deployment spreadsheet to add one or more additional cluster node FQDNs.

2. Deploy the additional cluster nodes in vCenter.

   For worker nodes, do not associate a disk (SDS=0).

   Use the same cluster node version for the new nodes, unless you are adding the nodes as part of a software upgrade.

3. Power on the additional cluster nodes.

   Wait approximately 6 minutes for the `first-boot` script to execute on the new nodes.

4. Log in to Cluster Control Manager using your customer account.

5. Create a directory called artifacts using the following command:

   ```
   mkdir -p ~/artifacts
   ```

6. Use a utility, such as WinSCP, to transfer the solution configuration spreadsheet to the `~/artifacts` directory.

7. Run the `ccm upgrade spec <solution spreadsheet name>.xlsx --infra --force` command.

8. Wait approximately 30 minutes, then run `ccm smoke-test` on Cluster Control Manager and confirm that all tests pass. If `ccm smoke-test` fails, run it again in 10 minutes. If the test continuously fails for over an hour, contact your technical support representative.

Result

The additional worker nodes are part of the cluster.

Managing cluster VPN

About this task

This procedure describes the actions you can take using clusterVPN options.

Procedure

Select one of the following options applicable to your situation:

- To set up cluster VPN on an installed cluster, run the **clusterVPN setup** command.

  The **clusterVPN setup** command enables IPSec and firewall settings on the cluster nodes. This command configures the settings for tunnel access between cluster nodes and sets configurations on all the cluster nodes.

- To start the cluster VPN on an installed cluster, run the **clusterVPN start** command.

  The **clusterVPN start** command starts IPSec and validates that IPSec is running correctly.

  Note:

  The cluster VPN must be set up before starting the service.

- To restart the cluster VPN on an installed cluster, run the **clusterVPN restart** command. The **clusterVPN restart** command restarts the cluster VPN on all cluster nodes.
- To stop the cluster VPN on an installed cluster, run the **clusterVPN stop** command.

  The **clusterVPN stop** command stops the cluster VPN on all cluster nodes.

- To disable the cluster VPN on an installed cluster, run the **clusterVPN disable** command.

  The **clusterVPN disable** command stops the running of cluster VPN on all cluster nodes and deletes the configuration.

- To find the status of the cluster VPN on an installed cluster, run the **clusterVPN status** command.

  The **clusterVPN status** command displays the status of cluster VPN on all the cluster nodes.

  If the **clusterVPN status** command fails for a cluster that has IPSec VPN with one node offline, run the **clusterVPN setup** command and then run the **clusterVPN start** command. If the issue persists, contact your technical support representative.

Related links

clusterVPN commands on page 219

Resize PVC using a spreadsheet

As an administrator, you can use this feature to change the Kubernetes Persistent Volume Claim (PVC) disk size as part of solution upgrade or PVC resize-only upgrade using a spreadsheet.

When resizing, a new PVC is created that is the size specified in the spreadsheet and the existing application data is then copied into the new PVC before removing the old PVC.

Note:

Services using the specified PVCs are taken offline during resize operations so a maintenance window is recommended.

This process is particularly beneficial if you need to make a lot of changes to PVC disk sizes concurrently.

Disk space considerations

- The cluster storage provider must have adequate free space for each resized PVC.

  For example, when resizing 3 PVCs and the existing size of each PVC is 25GB and the desired destination size of each PVC is 40GB, the storage provider requires 120GB of free space. After the command completes, an additional 45GB is consumed.

- If multiple PVCs are specified for resize in a single command, they are executed concurrently. If insufficient storage is available to execute all commands concurrently, resize each PVC individually.

  In the example, 120GB of free space is required at the time of the command. If PVCs are resized individually, the required free space is 70GB.

For details about resizing a PVC disk using the **`ccm resizePVC`** command, see ccm resizePVC command on page 218.

Related links

Resizing PVC using a spreadsheet on page 114

Resizing PVC using a spreadsheet

About this task

You can use the following procedure to resize one or more PVC disks using a spreadsheet and an upgrade operation with the `--products` option.

Procedure

1. Update the spreadsheet attribute size values of the PVC disks.

   Note:

   Be sure to include the correct units if required.

2. Update the version of the service using the PVC disk if required by the solution offer.
3. Run the `ccm upgrade spec <spreadsheet> --products` command.

   This command reads the new PVC size values in the spreadsheet to resize the PVC disk space allocated to a service.

   This command does not have to be performed during the initial installation. The PVC sizes can be modified using the spreadsheet at initial installation, as part of the product's software upgrade, or at any time using the product upgrade option.

4. Wait approximately 30 minutes, then run `ccm smoke-test` on Cluster Control Manager and confirm that all tests pass. If `ccm smoke-test` fails, run it again in 10 minutes. If the test continuously fails for over an hour, contact your technical support representative.

Related links

Resize PVC using a spreadsheet on page 113

Registering a product with Avaya

About this task

When you upgrade a cluster, deploy a new product, or after initial installation, you must update Avaya with the new software currently running on the cluster. This allows Avaya to promptly respond to issues with the software running in the cluster. This process is called Product Initial Registration.

Before you begin

To complete this task, you require the following information:

- IP Address of the Secure Access Link Gateway (SALGW) server. If this Avaya server is not currently installed at your location, contact Avaya.
- Customer Avaya credentials (user name and password). This information is needed by Avaya to ensure you have a valid account with Avaya.
- The Installation Location ID or Location Number. This number is provided to you when you purchase the product.
- Make sure Cluster Control Manager hostname has a full FQDN. Run the command **`hostname -f`** on CCM to find its hostname. If the reported hostname is not a full FQDN, update the hostname.

Procedure

1. Log in to Cluster Control Manager.
2. As root or sroot (via EASG) run the command **`/opt/avaya/ccm/salgw/bin/getSALGWCA -r SALGW-IPAdrress/FQDN`**.

   The command opens the firewall and downloads the certificate from port 7443.
3. Run the command **`/opt/avaya/ccm/salgw/bin/configureSALGW -li <locationNumber> -au <customer ID> -gw <SALGW-IPAddress/FQDN>`**.

   You may have to repeat this step multiple times until Cluster Control Manager has SEIDs for all the installed products.
4. When requested, enter your password.

   The command makes registration requests for all products currently installed on the cluster.

The command waits for the response. For each product that is deployed in the cluster, Avaya provides a SEID upon successful registration. The immediate response may not have the SEID for all the products.

5. Run the command **`/opt/avaya/ccm/salgw/bin/displaySEID`** to ensure Cluster Control Manager received the SEID for all products currently installed on the cluster.

Using a remote desktop session to launch the noVNC service

About this task

Use this procedure on Cluster Control Manager to start a noVNC session, which provides an interface for launching the Log Viewer, Manage Certificates, and Monitoring pages. If the cluster is not deployed, the links to these pages will be disabled.

You can only have one remote desktop session running at a time. The session expires after one hour of inactivity.

Before you begin

Deploy Cluster Control Manager.

Procedure

1. Log in to Cluster Control Manager.
2. Run the `remoteDesktopSession -geometry <width>x<height>` command.

   `-geometry` is optional. By default, the size is 1920x1080.
3. Navigate to the URI `https://xxxxxxxx:6901/vnc.html/vnc_auto.html?password=$PASSWORD` using the Firefox, Internet Explorer, or Chrome browser.
4. Click **Activities** > **Firefox**.

   The Cluster Control Manager landing page opens.
5. **(Optional)** To view logs as a non-EASG user (for example, `cust`), click **Log Viewer - Kibana** on the left of the Cluster Control Manager page.

   The Log Viewer (Kibana) page opens in a new tab.
6. **(Optional)** To view logs as an EASG (`craft`) user:
   a. Click the **M** icon to the right of the browser address bar.
   b. Click **Stop** if it is available next to Configure.
   c. Click **Configure** to modify the headers.
   d. Copy the Attributes information to the corresponding fields on the Sample Modify Headers page.

e. Click **Start** to apply the modification.

f. Return to the Cluster Control Manager page.

g. Click **Log Viewer - Kibana**.

h. Click **Click Here to Launch Log Viewer Kibana page** on the right of the Cluster Control Manager page.

The Log Viewer (Kibana) page opens in a new tab.

7. **(Optional)** To manage certificates as a non-EASG user (for example, `cust`), click **Manage Certificates** on the left of the Cluster Control Manager page.

   The public EJBCA page opens in a new tab.

8. **(Optional)** To manage certificates as an EASG (`craft`) user, click **Manage Certificates** on the left of the Cluster Control Manager page.

   A Notification page lists the steps to launch the **Manage Certificates** page.

   a. Follow the steps that are listed.

   b. Access the Administration page after it opens in a new tab.

9. To monitor, click Monitoring on the left of the Cluster Control Manager page.

   The Grafana login page opens in a new tab.

Resetting the remote desktop session password

Procedure

1. To reset the password, run the `remoteDesktopSession -resetpassword <remoteDesktopSession number>` command.
2. **(Optional)** To get the session number, run the `remoteDesktopSession -list` command.

Ending the remote desktop session

About this task

Use this procedure to end the remote desktop session.

Procedure

1. **(Optional)** To get the session number, run the `remoteDesktopSession -list` command.
2. To end the session, run one of the following commands:
   - Run `remoteDesktopSession -kill` on the session you started.
   - Run `remoteDesktopSession -kill <remoteDesktopSession number>` to end a specific session.

Accessing help for the remote desktop session

About this task

Use this procedure to access help or additional information for your remote desktop session.

Procedure

To access help information, run the `remoteDesktopSession -help` command.

HTTP(S) outbound proxy configuration

You can configure an outbound HTTP or HTTPS proxy if your organization mandates the use of a proxy. Clients are configured to send requests to a proxy. The information in this section does not apply to an environment which either has no proxy or which uses an implicit proxy.

Cluster Control Manager configuration

You can configure the outbound proxy in one of the following ways:

- When deploying the Cluster Control Manager OVF by populating the HTTP(S) Proxy Settings section.
- Using the **`ccmNetSetup`** command on Cluster Control Manager.

The proxy exclusion list contains addresses with which Cluster Control Manager does not communicate through the proxy. If an HTTP or HTTPS proxy is configured, the proxy exclusion list includes `localhost,127.0.0.0/8` by default. As part of the cluster installation, the exclusion list is automatically populated with the Cluster Control Manager, vCenter, cluster FQDN, keepalived virtual IP address, and Kubernetes master host virtual IP address.

The format of the destination address can include an IP address prefix (1.2.3.4), a domain name, or a special DNS label (*). Note the following:

- A domain name can match other names and subdomains.

  For example, with the domains `foo.example.com` and `example.com`, `example.com` matches both the `example.com` and `foo.example.com` domains. However, `.example.com` only matches `foo.example.com`.
- A single asterisk (*) indicates that no proxying is needed.
- You can include a port number with IP address prefixes and domain names.

  An example of an IP address prefix with a port number is `1.2.3.4:80`, and an example of a domain name with a port number is `foo.example.com:80`.

Configuring proxy settings when deploying Cluster Control Manager

Procedure

Populate the following HTTP(S) proxy settings while deploying the Cluster Control Manager OVF.

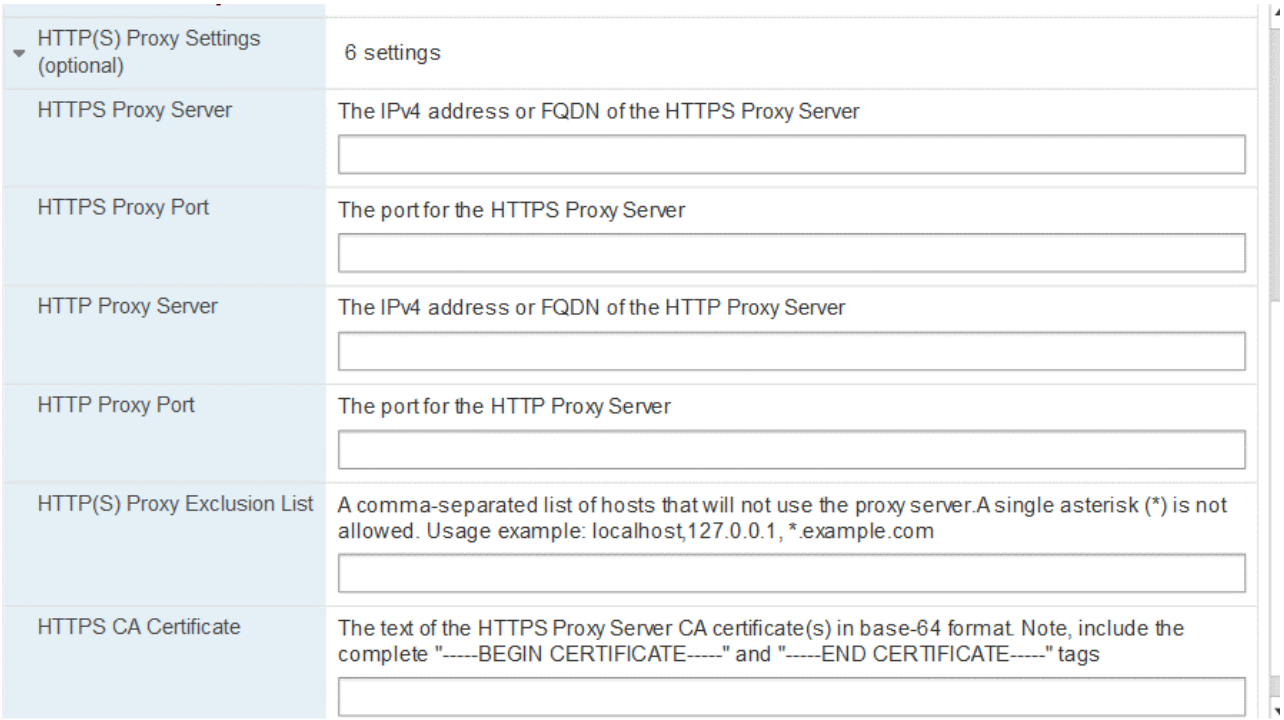

Using the ccmNetSetup command to configure proxy settings

About this task

You can configure proxy settings using the `ccmNetSetup` command. You can use this command anytime to update Cluster Control Manager network attributes. For example, if the FQDN or IP address of the explicit proxy changes, you can run the `ccmNetSetup` command to reflect the new destination.

Before you begin

- If you are configuring an HTTPS proxy, ensure that you have the HTTPS proxy server CA certificate.
- Perform a full backup of Common Services and application data.

Procedure

1. Log in to Cluster Control Manager with your customer account.
2. For an HTTPS proxy, copy the HTTPS proxy server CA certificate to a directory on Cluster Control Manager, such as `/home/cust/`.
3. Run `ccmNetSetup` or `ccmNetSetup --collect-proxy-only`.
   - Running `ccmNetSetup` without any options prompts you to populate all Cluster Control Manager network attributes.
   - Running `ccmNetSetup --collect-proxy-only` prompts you to update proxy-related network attributes only.

4. When you see the prompt `Would you like to configure an HTTP proxy?`, enter `y`.
5. When prompted, provide the following information:
   - HTTPS proxy IP/FQDN
   - HTTPS proxy port
   - HTTP proxy IP/FQDN
   - HTTP proxy port
6. When prompted to enter the absolute file path of the HTTPS proxy server CA certificate, enter the path to the certificate file.

   For example, `/home/cust/<file name>`.
7. When prompted to confirm that the above information is correct, enter `y` if all the information you provided is correct.
8. To complete the proxy configuration, close the current command session and open a new one.

Enabling or disabling file integrity validation

About this task

You can enable file integrity validation on Avaya Common Services servers if you require Advanced Intrusion Detection Environment (AIDE) logs. By default, file integrity validation is disabled.

When file integrity validation is enabled, the server contains additional log files that consume up to 400 MB of additional disk space. This feature also consumes additional CPU while initializing the integrity database and executing the validation. By default, the file integrity validation software runs at 3:00 a.m. every day, but this can be delayed depending on when the validation is enabled.

- If the validation is enabled before 3:00 p.m., the first report is ready by 3:00 a.m. the following day.
- If the validation is enabled after 3:00 p.m., the first report might be ready by 3:00 a.m. the following day or it might be delayed until 3:00 a.m. on the day after.

AIDE logs are only generated if the system detects a problem.

Before you begin

Perform a full backup of Common Services and application data.

Procedure

1. Log in to Cluster Control Manager.
2. To see whether file integrity validation is enabled or disabled, run the `clusterFileIntegrity` command.

3. To enable file integrity validation, run the `clusterFileIntegrity enable` command.

   When file integrity validation is enabled, the system generates AIDE logs when problems are detected.

4. **(Optional)** To disable file integrity validation, run the `clusterFileIntegrity disable` command.

Creating index patterns on Kibana

About this task

An index pattern identifies one or more OpenSearch indexes to explore with the Kibana logging service. Kibana (OpenSearch dashboard) performs a search for the index names that match the specified pattern. An asterisk (*) in the pattern matches zero or more characters. Index patterns enable you to interactively explore and visualize data on Kibana.

Before you begin

- Perform a full backup of Common Services and application data.
- Ensure sufficient disk space is available to store the indexes.
- Ensure the status of all the indexes displays in green.
- Access the Kibana URL to ensure the web page is functional.
- Create the Kibana user and credentials for logging in to the dashboard.

```
ccm release common-services createKibanaUser -u <username> -r READ-WRITE
```

Procedure

1. Log in to Kibana using the URL `https://<cluster-FQDN>/logging`.
2. From the menu, click **Stack Management**.
3. Click **Index Patterns**.

   The index patterns are not created by default. You must create the index patterns.

4. Click **Create Index Pattern**.
5. Specify the index pattern in the following format:

   `fluent-<yyyy.mm>.*`, where *yyyy* is the year and *mm* is the month of the index pattern creation.

   For example: `fluent-2020.06.*`

   Important:

   Do not create generic index patterns like `fluent-*` and `trace-*`.

Tip:

To filter the logs for k8saudit, trace, and ausec, you can create index patterns in the following formats:

`fluent-k8saudit-<yyyy.mm>.*`

`trace-<yyyy.mm>.*`

`ausec-<yyyy.mm>.*`

For example:

`fluent-k8saudit-2020.06.*`

`trace-2020.06.*`

`ausec-2020.06.*`

6. Click **Next step**.
7. From the **Time Filter field name** drop-down, select **@timestamp**.
8. Click **Create Index Pattern**.

   Kibana creates an index pattern.
9. From the menu, click **Discover**.

   Kibana displays all logs for the selected index pattern.

Viewing AIDE logs in the Kibana logging interface

About this task

You can use the Kibana logging interface to view AIDE logs from the Elasticsearch server. Currently, the Kibana interface displays AIDE logs for cluster nodes. You cannot use this procedure to view AIDE logs from Cluster Control Manager.

Before you begin

- Ensure that file integrity validation is enabled. By default, the file integrity validation software runs every day at 3:00 a.m.
- Ensure that you know how to create index patterns in Kibana.

Procedure

1. Log in to Kibana using the URL `https://<cluster-FQDN>/logging`.
2. Create an index pattern called `fluent-k8saudit*` if one does not already exist.

   From the **Time Filter field name** drop-down menu, ensure that you select **@timestamp** and then click **Create index pattern** to complete the process.

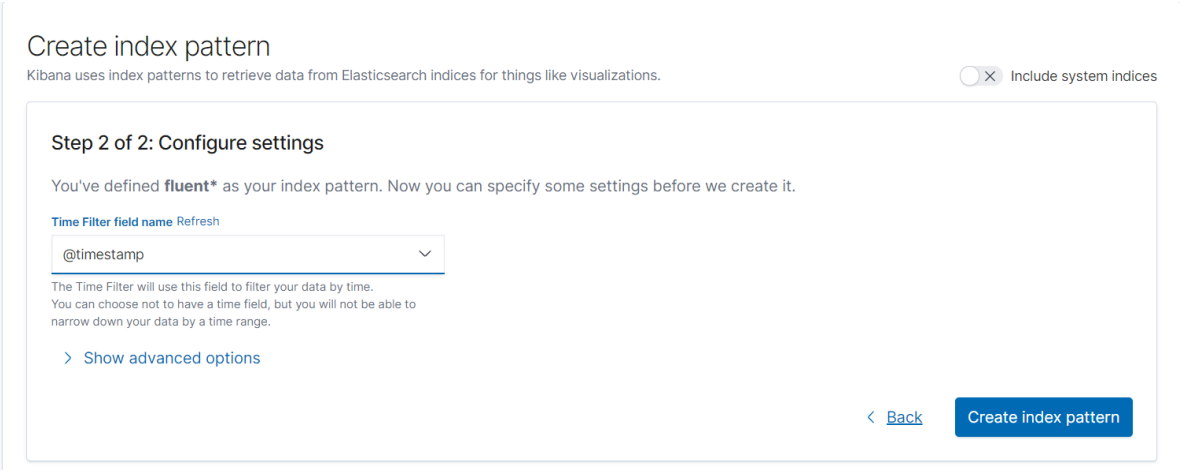

3. Open the Discover view by clicking on the left side of the screen.
4. To view AIDE logs from all servers, do the following:
   a. Type `tag_all:(k8saudit.instanceaide)` in the **Search** field.
   b. Select a time range.
5. To view AIDE logs for specific servers, do the following:

   Run the `k get nodes -o wide` command and find your host IP addresses in the `INTERNAL-IP` column.

```
[root@flex-140 ~]# k get nodes -o wide
NAME                      STATUS   ROLES
AGE     VERSION   INTERNAL-IP      EXTERNAL-
IP   OS-IMAGE                                KERNEL-
VERSION                 CONTAINER-RUNTIME
flex-143.dr.example.com   Ready    controller-worker   5d14h
v1.22.2   192.0.2.143   <none>        Red Hat Enterprise Linux
8.5 (Ootpa)   4.18.0-348.20.1.el8_5.x86_64   containerd://1.5.7
flex-144.dr.example.com   Ready    controller-worker   5d14h
v1.22.2   192.0.2.144   <none>        Red Hat Enterprise Linux
8.5 (Ootpa)   4.18.0-348.20.1.el8_5.x86_64   containerd://1.5.7
flex-145.dr.example.com   Ready    controller-worker   5d14h
v1.22.2   192.0.2.145   <none>        Red Hat Enterprise Linux
8.5 (Ootpa)   4.18.0-348.20.1.el8_5.x86_64   containerd://1.5.7
```

6. **(Optional)** To filter logs by their messages, select **message** from the list of available fields and then click **add**.

Linux audit rule updates

If your enterprise requires additional security, you can work with Avaya support personnel to update the default audit rules for Linux. The option to update audit rules is particularly useful for government solutions using JITC.

Root privileges are required to update audit rules. Contact Avaya support personnel for assistance, but before doing this, take a full backup of Avaya Common Services and application data.

For more information about audit rule options, see https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/security_guide/sec-defining_audit_rules_and_controls. The audit package also includes various preconfigured files, as described in the “Preconfigured Rules Files” section.

Important:

When a node goes down, your audit rule updates are lost. Work with Avaya support personnel to re-apply your audit rule updates after an infrastructure upgrade, a restore, or a node resurrection.

Security warning banner configuration

If your organization requires a customized security warning banner, contact Avaya support personnel for assistance. The security banner is displayed before or directly on the system login screen.

Backup and restore operations

The following sections provide manual backup and restore procedures. You can manually back up Common Services, restore Common Services, or restore Cluster Control Manager.

Important:

You cannot back up or restore cluster nodes through VMware snapshots. Instead, use the procedures in this section. Ensure that you have a recent Cluster Control Manager backup and a solution backup.

Archive destination

You can set the archive destination for backup files to:

- Local: To store backup files in the local directory at `/var/avaya/artifactCache/ccmClusterBackup`.
- Remote: To specify a remote directory for storing backup files.

You can configure the archive destination as follows:

- During the deployment or upgrade process
- While performing the backup procedure in this chapter
- With the `ccm archive config [local|remote]` command

When setting the archive destination to Remote, you must also specify your remote server details.

Scheduled backup

You can also configure scheduled backups. You can configure scheduled backup settings during the deployment or upgrade process, or using the `ccm backup schedule` command with the options you want to configure.

Log files

You can find log files at `/var/log/avaya/ccm/`. The following backup and restore log files are available:

- `backuprestore.log`
- `scheduledbackup.log`

Related links

[ccm archive command options](#) on page 192
[ccm backup and ccm backup schedule command options](#) on page 194
[ccm restore command options](#) on page 196
[ccm report command](#) on page 197

Backing up Common Services

About this task

Use this procedure to generate a backup file for Common Services. This backup does not contain product application data.

The following content is included in this Common Services backup:

- Local users for the system
- Product configurations
- Infrastructure configurations
- System certificates

You can run `ccm backup` on Cluster Control Manager to manually create a backup. You are automatically prompted to perform a backup during a major infrastructure or service change using the `ccm upgrade` or `ccm install` commands.

Avaya recommends taking a backup after any configuration changes to the system or deployed products.

Procedure

1. Log in to Cluster Control Manager.

2. To manually start a backup, run one of the following commands:
   - `ccm backup`: Stores the backup file in either the local or remote directory, depending on the archive destination you configured.

     The local directory location is `/var/avaya/artifactCache/ccmClusterBackup`. The remote location is `<base directory path>/ccmClusterBackup`.
   - `ccm backup --local`: Stores the backup file in the local directory at `/var/avaya/artifactCache/ccmClusterBackup`. This command overrides the configured archive destination.
   - `ccm backup --remote-server "<FQDN/IP> [-p <port>] -u <username> -d <directory path>"`: Stores the backup file in the remote directory you specify. This command overrides the configured archive destination.

     Include the double quotes as shown above for the `--remote-server` option. If you omit these quotes, the command will fail.

     You need to include `-p <port>` if you are using a port other than the default port 22.
3. When you are prompted to proceed with the backup, enter `y`.

   For local backups, confirm that the oldest backup file may be deleted.

   ```
   There are two backup archives located /var/avaya/artifactCache/
   ccmClusterBackup. The oldest one will be deleted
   Do you wish to proceed?
   Response [y, n] => y
   ```
4. If you used the `--remote-server` option in step 2, enter the remote server password when prompted.
5. If you did not configure a backup password, enter one when prompted.

   This password is used to encrypt the generated backup file. You will need this password to perform restore operations.

   The password must be at least 8 characters long and contain at least the following:
   - One upper case character
   - One lower case character
   - One number
   - One special character
6. Wait for the backup to complete.

   You will be notified of the location of the `.tgz` backup file.
7. If the backup file is in the local `/var/avaya/artifactCache/ccmClusterBackup` directory, copy it to a secure location and transfer it off of Cluster Control Manager.

   You are prompted when more than two backups are in the local directory.

For scheduled backups, only the two most recent backup files are kept in the local directory.

8. **(Optional)** To free up space for future use, remove the backup file.

Restoring Avaya Common Services in an online deployment environment

About this task

Use this procedure to recover from a full cluster outage when one or more of the virtual machines in the solution are lost. This procedure restores Cluster Control Manager, infrastructure, and services. It does not restore service application data.

If the CCM and cluster nodes are intact and healthy, you can use your existing Cluster Control Manager and nodes, but if it not intact and healthy you can delete the Cluster Control Manager and nodes and create new.

Note:

This restore process is supported if the cluster or service version in the backup file is Common Services 1.3.0.x or later.

Before you begin

- Perform a backup and locate your backup file.
- Ensure that you have the password for the backup file.
- Assess the node cluster state to determine which option to use for the restoration. For more information about the options, see Assessing cluster state on page 128.

Procedure

1. Log in to Cluster Control Manager.
2. On Cluster Control Manager, run the `screen` command to run the restore process in the background.

   When running the restore process in the background, if you need to detach from the SSH session, see Detaching from the restore SSH session on page 134.
3. Run one of the following commands to start the restore process:

   Note:

   If your backup file is not accessible on a remote server, copy the backup file to Cluster Control Manager using a file transfer utility such as WinSCP.

   You can use the `/tmp` or `/var/avaya/artifactCache/ccmClusterBackup` directory.

   - `ccm restore all --remote-server "<FQDN/IP> [-p <port>] -u <username>" <path to backup file>`: Restores from the remote directory path you specify. This command overrides the configured archive destination.

Include the double quotes as shown above for the `--remote-server` option. If you omit these quotes, the command will fail.

You need to include `-p <port>` if you are using a port other than the default port 22.

- `ccm restore all --local <path to backup file>`: Restores from the local `<path to backup file>` directory. This command overrides the configured archive destination.
- `ccm restore all <path to backup file>`

In these command options, replace `<path to backup file>` with the full path to the backup file. For example, `/var/avaya/artifactCache/ccmClusterBackup/<backup-file>.tgz` or `home/<customer>/<backup-file>.tgz`.

Use the same archive destination (local or remote) that you used for the backup.

4. If you used the `--remote-server` option, enter the remote server password when prompted.
5. If you did not configure the password for the backup file, enter it when prompted.
6. To confirm the restore, type `y`.
7. When prompted, enter your Avaya SSO credentials.
8. Wait for the cluster to fully initialize.
9. Wait approximately 30 minutes, then run `ccm smoke-test` on Cluster Control Manager and confirm that all tests pass. If `ccm smoke-test` fails, run it again in 10 minutes. If the test continuously fails for over an hour, contact your technical support representative.
10. On Cluster Control Manager, run `ccm status --pod-details` and ensure that the status of all pods is Running or Completed.

    When a deployment is complete, the status of all services, except for tools-policy, nfs-provisioner, and utility-service, is Deployed.

Next steps

- The restore process disables file integrity validation. If you enabled file integrity validation previously, you can re-enable it by running the `clusterFileIntegrity enable` command.
- If you imported third-party certificates, the system will not recognize them after the restore. Therefore, you must repeat the import process for the restored cluster to use new third-party certificates.

Assessing cluster state

About this task

Use this procedure to determine whether to deploy the same version of Cluster Control Manager used when performing the backup or delete any surviving virtual machine remnants before restoring Avaya Common Services.

Before you begin

Procedure

1. Determine which course of action is applicable to your situation.
2. Use one of the following options, depending on whether or not your CCM and nodes are intact and healthy.
   - Option 1 - CCM and cluster nodes are intact and healthy:

     Deploy the same version of Cluster Control Manager you used when performing the backup. Ensure that you also use the same configuration, including the same IP address, number of nodes, Node IP address, Node version, FQDN, enrollment password, disk size, CPU, and memory.

     Determine if your backup archive is the same version as your CCM. You can find the Cluster Control Manager version by looking at the backup file name. The format of the backup file name is `<ccm hostname or IP>-<common services product version>-<ccm version>-<timestamp>.tgz`.
   - Option 2 - CCM and cluster nodes are not intact and healthy:

     Delete any surviving virtual machine remnants, such as virtual machine cluster nodes, from vCenter manually. You can power off and delete these virtual machine remnants in any order.

     For information about cleaning up virtual machine remnants from the datastore, see Deleting virtual machine remnants after restoring Avaya Common Services on page 135.

Restoring Avaya Common Services in an offline air gap environment

About this task

In an air gap network environment, use this procedure to recover from a full cluster outage when one or more of the virtual machines in the solution are lost. This procedure restores Cluster Control Manager, infrastructure, and services. It does not restore service application data.

Note:

This restore process is supported if the cluster or service version in the backup file is Common Services 1.3.0.x or later.

Before you begin

- Perform a backup and locate your backup file.
- Ensure that you have the password for the backup file.
- Assess the node cluster state to determine which option to use for the restoration. For more information about the options, see Assessing cluster state on page 128.

Procedure

1. Log in to Cluster Control Manager.

2. Copy the backup file to Cluster Control Manager using a file transfer utility, such as WinSCP.

   You can use the `/tmp` or `/var/avaya/artifactCache/ccmClusterBackup` directory.

3. Change the directory to the location where you transferred the backup onto Cluster Control Manager.
4. Create a temporary directory by running the `mkdir ./temp` command.
5. Run `tar -C ./temp/ -zxf <backup file name> ccm/config/*.gpg` to extract the Cluster Control Manager encrypted backup file.
6. Run `gpg --yes --batch ./temp/ccm/config/*.gpg` to extract the Cluster Control Manager backup `.tar` file from the encrypted backup file.
7. When prompted, enter the password that you used to encrypt the backup file.

   The following is displayed after you enter the password:

   ```
   gpg: CAST5 encrypted data
   gpg: encrypted with 1 passphrase
   gpg: WARNING: message was not integrity protected
   ```

8. Take note of the name of the `.tgz` file in the `./temp/ccm/config/` directory for the next step.

   You can run `ls ./temp/ccm/config/` to find the file.

9. Extract the cluster-config and solution services file, which contains the software artifact information that was deployed in the cluster.

   Run the following commands:

   - `tar -C ./temp/ -zxf ./temp/ccm/config/ <replace_with_the_tgz_file_from_the_previous_step> backup/ clusters/*/releases/initial-service.yaml`
   - `tar -C ./temp/ -zxf ./temp/ccm/config/ <replace_with_the_tgz_file_from_the_previous_step> backup/ clusters/*/meta/cluster-config.yaml`

   For example, where `<replace_with_the_tgz_file_from_the_previous_step>` is `10.10.10.140-1.1.100011050115-202011172342-ccmbackup.tgz`, run the following:

   - `tar -C ./temp/ -zxf ./temp/ccm/config/ 10.10.10.140-1.1.100011050115-202011172342-ccmbackup.tgz backup/ clusters/*/releases/initial-service.yaml`
   - `tar -C ./temp/ -zxf ./temp/ccm/config/ 10.10.10.140-1.1.100011050115-202011172342-ccmbackup.tgz backup/ clusters/*/meta/cluster-config.yaml`

10. Run the following commands to recreate the solution deployment file (`solution-backup.yaml`):
    - `cat ./temp/backup/clusters/*/meta/cluster-config.yaml > ./solution-backup.yaml`
    - `cat ./temp/backup/clusters/*/releases/initial-service.yaml >> ./solution-backup.yaml`
    - `chmod 644 $PWD/solution-backup.yaml`
11. Run `ls $PWD/solution-backup.yaml` to obtain the absolute path to the solution deployment file.
12. Based on the fully qualified path displayed, transfer the `solution-backup.yaml` file to the client computer where the ccm-ctl-agn container is running.
13. Run `rm -rf ./temp` to delete the temporary directory that contained the extracted backup content on Cluster Control Manager.
14. On the client computer where the ccm-ctl-agn container is running, use the `solution-backup.yaml` file to download the Avaya solution software artifacts and then upload them to Cluster Control Manager.
    a. Run `agn download solution-backup.yaml` to download the solution software artifacts.
    b. Configure and start the local chart museum and Docker registry on Cluster Control Manager.
    c. Copy the software artifacts onto Cluster Control Manager.
    d. Upload the software artifacts to the Cluster Control Manager local chart museum and Docker registry.

       You can proceed with the remaining steps in this procedure after all solution software artifacts are uploaded to the Cluster Control Manager chart museum and Docker registry.
15. On Cluster Control Manager, run the `screen` command to run the restore process in the background.

    When running the restore process in the background, if you need to detach from the SSH session, see Detaching from the restore SSH session on page 134.
16. Run one of the following commands to start the restore process:

    Note:

    If your backup file is not accessible on a remote server, copy the backup file to Cluster Control Manager using a file transfer utility such as WinSCP.

You can use the `/tmp` or `/var/avaya/artifactCache/ccmClusterBackup` directory.

- `ccm restore all --remote-server "<FQDN/IP> [-p <port>] -u <username>" <path to backup file>`: Restores from the remote directory path you specify. This command overrides the configured archive destination.

  Include the double quotes as shown above for the `--remote-server` option. If you omit these quotes, the command will fail.

  You need to include `-p <port>` if you are using a port other than the default port 22.

- `ccm restore all --local <path to backup file>`: Restores from the local `<path to backup file>` directory. This command overrides the configured archive destination.
- `ccm restore all <path to backup file>`

In these command options, replace `<path to backup file>` with the full path to the backup file. For example, `/var/avaya/artifactCache/ccmClusterBackup/<backup-file>.tgz` or `home/<customer>/<backup-file>.tgz`.

Use the same archive destination (local or remote) that you used for the backup.

17. If you did not configure the password for the backup file, enter it when prompted.
18. To confirm the restore, type `y`.
19. When prompted for your Avaya SSO and vCenter credentials, enter the credentials for the Cluster Control Manager local Docker registry. Note that this prompt may differ, depending on whether you are restoring on existing virtual machines or restoring on new virtual machines.

    If you receive an `Error logging into docker registry` message while trying to log in, contact Avaya support personnel for assistance.
20. Wait for the cluster to fully initialize.

Next steps

- The restore process disables file integrity validation. If you enabled file integrity validation previously, you can re-enable it by running the `clusterFileIntegrity enable` command.
- If you imported third-party certificates, the system will not recognize them after the restore. Therefore, you must repeat the import process for the restored cluster to use new third-party certificates.

Restoring Cluster Control Manager

About this task

Use this procedure to recover Cluster Control Manager from an outage when the Cluster Control Manager virtual machine is lost. This procedure only restores Cluster Control Manager. It does not restore cluster nodes or application data.

Note:

This restore process is supported if the cluster or service version in the backup file is Common Services 1.3.0.x or later.

Before you begin

- Locate the backup file.

  Use the most recent valid backup file when restoring Cluster Control Manager.
- Ensure that you have the password for the backup file.
- Assess the node cluster state to determine which option to use for the restoration. For more information about the options, see Assessing cluster state on page 128.

Procedure

1. Log in to Cluster Control Manager.
2. Ensure that the IP address and version of Cluster Control Manager matches the IP address in the backup file name.
3. On Cluster Control Manager, run the `screen` command to run the restore process in the background.

   When running the restore process in the background, if you need to detach from the SSH session, see Detaching from the restore SSH session on page 134.
4. Run one of the following commands to restore the Cluster Control Manager:

   Note:

   If your backup file is not accessible on a remote server, copy the backup file to Cluster Control Manager using a file transfer utility such as WinSCP.

   You can use the `/tmp` or `/var/avaya/artifactCache` directory.

   - `ccm restore --remote-server "<FQDN/IP> [-p <port>] -u <username>" <path to backup file> all`: Restores from the remote directory path you specify. This command overrides the configured archive destination.

     Include the double quotes as shown above for the `--remote-server` option. If you omit these quotes, the command will fail.

     You need to include `-p <port>` if you are using a port other than the default port 22.
   - `ccm restore --local all <artifact full path>`: Restores from the local `<path to backup file>` directory. This command overrides the configured archive destination.

   In all of these command options, replace `<path to backup file>` with the full path to the backup file. For example, `/var/avaya/artifactCache/ccmClusterBackup/<backup-file>.tgz` or `home/<customer>/<backup-file>.tgz`.

   Use the same archive destination (local or remote) that you used for the backup.
5. If you used the `--remote-server` option, enter the remote server password when prompted.

6. If you did not configure the password for the backup file, enter it when prompted.
7. To confirm the restore, type `y`.
8. Wait for the restore process to finish.
9. To verify the restore, run the `ccm status` and `swversion` commands.

Next steps

Power off and delete the non-functional Cluster Control Manager, which was lost during the outage, from vCenter manually.

The restore process disables file integrity validation. If you enabled file integrity validation previously, you can re-enable it by running the `clusterFileIntegrity enable` command.

Detaching from the restore SSH session

About this task

When you are running the restore process in the background, you can use this procedure to detach from the SSH session and close it. You can reattach to the session later to reopen it.

Procedure

1. Start a new SSH session to Cluster Control Manager and log in with your customer account.
2. Type `screen -ls` to retrieve the screen ID of the restore session.
3. Type `screen -d <screen id>` to detach the session from the SSH session.
4. Close the SSH session.

Reattaching to the restore SSH session

About this task

After detaching from the restore SSH session, use this procedure if you want to reopen the SSH session.

Procedure

1. Log in to Cluster Control Manager using your customer account.
2. At the CCM prompt, type `screen -ls` to retrieve the screen ID of the session.

   If no screen IDs are listed, the restore is complete.
3. Type `screen -r <screen id>` to reattach to the upgrade session.

Deleting virtual machine remnants after restoring Avaya Common Services

About this task

Use this procedure to delete any remnants in the datastore. The remnants can be deleted in any order.

Procedure

1. Log in to vCenter using the credentials you used to deploy the cluster.

   Alternatively, you can log in using an account that has administrator privileges on the datastore.
2. Using the Storage view, click the datastore that was previously used to deploy the cluster.
3. Click the **Files** tab.
4. Select the virtual machine that you want to delete.
5. Select either **Shut Down Guest OS** or **Power Off**.
6. Click **Yes** on the confirmation dialog.

   The virtual machine is powered off.
7. Select **Delete from Disk** (or similar option) to confirm that all virtual machine remnants are deleted from the virtual client.
8. Click **Yes** on the confirm deletion dialog.
9. Watch the Recents Tasks pane to view the delete progress.

Restoring or replacing a cluster node

About this task

Use this procedure to replace and join a missing or deleted cluster node into an existing cluster.

Before you begin

- Obtain the cluster node OVA. The version of the OVA must be the same version as the existing cluster nodes.
- Obtain the network settings (FQDN, IP address, gateway, and netmask) used by the original cluster node that is being recovered.
- Obtain the CPU and memory resource values used by the original cluster node that is being recovered.
- Obtain the SDS disk size (Disk 2) used by the original cluster node that is being recovered.
- Obtain the enrollment password configured on Cluster Control Manager.
- If DRS is configured within vCenter to support node anti-affinity, obtain the VM name to be recovered.

Procedure

1. Deploy the cluster node OVA using the same cluster node version, VM name, network settings, enrollment password, and VM resources (CPU, memory and SDS disk) as the original cluster node.

   Before the VM is powered on, use the **Edit Settings** option for the VM to modify its CPU and memory values.
2. If you enabled the HA audit during initial installation, on your vCenter Cluster object, edit the VM/Host Anti-Affinity rule created earlier and add the replacement VM name back into the rule.
3. Power on the cluster node VM.
4. Log in to Cluster Control Manager with your customer account.
5. On Cluster Control Manager, run `ccm restore node` to recover and join the missing cluster node into the cluster.
6. Wait approximately 30 minutes, then run `ccm smoke-test` on Cluster Control Manager and confirm that all tests pass. If `ccm smoke-test` fails, run it again in 10 minutes. If the test continuously fails for over an hour, contact your technical support representative.
7. If the old cluster node VM still exists (powered down), delete the VM using **Delete from Disk**.

Chapter 4: Troubleshooting Avaya Analytics™ from the Cluster Control Manager console

Troubleshooting Avaya Analytics™ from the Cluster Control Manager console

You can troubleshoot common issues in Avaya Analytics™ by running the post-install scripts on the Cluster Control Manager (CCM) console. To complete the troubleshooting procedures in this chapter, you must log in to the CCM console as a customer user, where the customer username varies for every organization, and then switch to the root user. Root user access might not be required for some steps as indicated in the respective procedures.

To complete some of these procedures you require access to:

- Cluster Control Manager IP address
- System Manager IP address
- Authorization Service Avaya Breeze® node SIPs
- Avaya Oceana® Cluster 1 Avaya Breeze® Node SIPs
- Avaya Control Manager

Note:

The post-install script might take some time to display the options after you run the **`ccm release orca analytics`** command or might display the following message: `Failed to get auth token`. Wait for some time or type the command again.

Warning:

Do not delete nodes or VMs from your VMware.

Capturing logs

About this task

You can capture the different Avaya Analytics™ services logs to analyze these logs for troubleshooting. You can capture these logs based on a specific time and format.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. To select the **Logs** option, enter the corresponding number.
6. To proceed with capturing the logs, in the **Proceed to Log Capture option?** field, type `y`.

   Typing `n` cancels the operation.

   CCM console displays the following Log Capture Options:

   - All Services (default)
   - REF Input Adaptor
   - Specific Services
7. To capture the logs from all the Avaya Analytics™ services, enter the number corresponding to the **All Services** option in the **Please enter the value of the logs you want to capture** field.
8. To capture the logs from the REF Input Adapter, enter the number corresponding to the **REF Input Adaptor** option in the **Please enter the value of the logs you want to capture** field.

   Note:

   Though the All Services option displays all the logs including the REF Input Adapter logs, this option specifically displays only the REF Input Adaptor logs.
9. To capture the logs for a specific service, enter the number corresponding to the **Specific Services** option in the **Please enter the value of the logs you want to capture** field.
10. In the **Please select a service to add to your list** option, select the required services and type `99`.

    In this step, the user creates a list of the required services by entering the corresponding number from the available list.
11. In the **Please enter the value of the time frame you wish to use** field, enter the number corresponding to the required time.

    The options are:

    - Between 2 specified dates: This option captures all logs between the specified dates.
    - Previous X hours: This option captures the logs for the last specified hours. For example, if you enter 24, then you can capture the logs of the last 24 hours.
12. In the **Please input the date and time which you want the logs to start from** field, enter the required details in iso-8601 format timestamp (local time).

For example, 2019–12–24T20:40:03

13. In the **Please input the date and time which you want the logs to end in** field, enter the required details in iso-8601 format timestamp (local time).

    For example, 2020–01–24T20:40:00

14. In the **Please enter the value of the log format you wish to use** field enter the number corresponding to the required format.

    The options are:

    - Expanded JSON (default)
    - Flattened JSON (easier to read)

15. In the **Are you sure you wish to proceed to capture logs using the above selection?** field type `y` or `n`.

    - If you type `y`, CCM starts capturing the logs from the selected services within the desired timeframe by using the selected format.
    - If you type `n`, CCM cancels the log capturing.

Result

- CCM allocates a sub folder for each set of logs for a service in the `/var/avaya/artifactCache` directory.
- When the log capture is complete, the script compresses these folders into a single tar file `AnalyticsLogCapture_<todays date>.tgz` in the `/var/avaya/artifactCache` directory.

Next steps

You can import this log bundle to your local machine to analyze the logs using any file transferring method, such as WinSCP.

Troubleshooting database issues

About this task

If there are issues with your reporting data, troubleshoot the typical database issues using the commands in this procedure. For example, if an ETL process stops running, data stops moving to External Data Mart (EDM).

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`

4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **Database** by pressing the corresponding number.

6. To check the running process table, select the **Check the running process table** option by entering the corresponding number.

   If a process is running, clear the running process.
7. Return to the previous page by entering `b`.
8. Quit the current page by entering `q`.
9. Return to the main menu by entering `m`.
10. To clear the running process table, select the **Clear the running process table** option by entering the corresponding number.

    This step restarts any running processes, and the ETL process gets enabled.
11. Return to the previous page by entering `b`.
12. Quit the current page by entering `q`.
13. Return to the main menu by entering `m`.
14. To check the status of the ETL process, select the **Check if scheduled ETL process is enabled** option by entering the corresponding number.

    If the ETL process is not enabled, data stops moving to EDM. You must check if there is a process running, and clear the running process by using the **Enable Scheduled ETL Process** option, if required.

    Alternatively, if you want to disable an ETL process, use the **Disable Scheduled ETL Process** option.
15. Return to the previous page by entering `b`.
16. Quit the current page by entering `q`.
17. Return to the main menu by entering `m`.
18. To check the Interval Data loads, select the **Check interval data loads** option by entering the corresponding number.

    The page displays the latest data that was sent to EDM for each metric. If the data is not updated for more than one interval, you must troubleshoot the issue by checking the ETL process.
19. Return to the previous page by entering `b`.
20. Quit the current page by entering `q`.
21. Return to the main menu by entering `m`.
22. To check the CDR data loads, select the **Check CDR data loads** option by entering the corresponding number.

    The page displays the latest data that was sent to EDM for each metric. If the data is not updated for more than one interval, you must troubleshoot the issue by checking the ETL process.

23. Return to the previous page by entering `b`.
24. Quit the current page by entering `q`.
25. Return to the main menu by entering `m`.
26. To check the Dimension data loads, select the **Check Dimension data loads** option by entering the corresponding number.

    The page displays the latest data that was sent to EDM for each metric. If the data is not updated for more than one interval, you must troubleshoot the issue by checking the ETL process.
27. To check disk space of pgdata, select the **Check disk space of pgdata** option by entering the corresponding number.

    Note:

    You can check the disk space only if the service crunchy-primary-service-orca-dbmgr-0 is in a ready or running state before continuing.
28. Return to the previous page by entering `b`.
29. Quit the current page by entering `q`.
30. Return to the main menu by entering `m`.
31. Select **Deployment** by pressing the corresponding number.
32. To select the **GDPR** option, enter the corresponding number.
33. To view the current settings for the Call Originator Redaction feature, select the **Show current GDPR configuration** option, enter the corresponding number.

    Depending on whether the Call Originator Redaction feature is enabled or disabled, the screen displays the following:

    - `Call Originator Redaction is currently enabled`: The Personally identifiable information (PII) is removed from the Analytics logs, which enables the Automatic encryption.
    - `Call Originator Redaction is currently disabled`:The Personally identifiable information (PII) is not removed from the Analytics logs, which disables the Automatic encryption.
34. Return to the previous page by entering `b`.
35. Quit the current page by entering `q`.
36. Return to the main menu by entering `m`.
37. To enable the Call Originator Redaction feature, select the **Enable GDPR configuration option**, enter the corresponding number.
38. In the **Proceed with disabling GDPR for Historical and Realtime reporting value** field, type `Y` or `N` to cancel the operation.

Warning:

Setting the GDPR value requires Input Adaptor pod(s) and CDR pod(s) restart.

39. To disable the Call Originator Redaction feature, select the **Disable GDPR configuration** option, enter the corresponding number.
40. In the **Proceed with disabling GDPR for Historical and Realtime realtime reporting value** field, type `y`.Entering `n` cancels the operation.

Warning:

Setting the GDPR value requires Input Adaptor pod(s) and CDR pod(s) restart.

41. Return to the previous page by entering `b`.
42. Quit the current page by entering `q`.
43. Return to the main menu by entering `m`.

Restarting Historical Reporting

About this task

If you cannot run historical reports using Avaya Analytics™ web and all other troubleshooting efforts are unsuccessful, use the following steps to restart the Historical Reporting server.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. To select the **Historical Reporting** option, enter the corresponding number.
6. To restart the Historical Reporting Server, select the **Restart the Historical Reporting Server** option by entering the corresponding number.
7. Return to the previous page by entering `b`.
8. Quit the current page by entering `q`.
9. Return to the main menu by entering `m`.

Unable to configure Historical Reporting with LDAP

Condition

Unable to configure Historical Reporting with Lightweight Directory Access Protocol (LDAP).

Cause

The LDAP server might be inaccessible from the environment where Historical Reporting is hosted.

Solution

1. To test the LDAP server accessibility from Avaya Analytics™, log in to the Cluster Control Manager (CCM) console.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. Select the **Historical Reporting** option by entering the corresponding number.
6. Select the **Test LDAP server accessibility** option by entering the corresponding number.
7. In the **Proceed with LDAP validation** field, type `y`.

   Typing `n` cancels the operation.
8. In the **Enter the FQDN of the LDAP server** field, type the FQDN of the LDAP server and press **Enter**.
9. In the **Enter the port of the LDAP server** field, type the port number of the LDAP server and press **Enter**.

   If the LDAP server accessibility test is successful, the CCM console displays the following message: `Successfully connected to LDAP server. Server is reachable.`

   If the LDAP server accessibility test fails, the CCM console displays the following message: `Unable to reach LDAP server using telnet.`
10. If the LDAP server accessibility test is successful and you are still not able to configure LDAP, verify that the LDAP password is correct.
11. If the LDAP server accessibility test fails, check the DNS server and the firewall settings.

The mstr_srv pod fails to start

Condition

The mstr_srv pod fails to start after you configure the Avaya Analytics™ email distribution services.

Cause

This issue occurs because of discrepancies in the configuration of Avaya Analytics™ email distribution services.

Solution

1. Reconfigure the email distribution settings and restart the mstr_srv pod.

   Ignore the following warning message:

   ```
   Maximum number of semaphore arrays 128 is low.
   ```

2. If the issue is not resolved, complete the following steps to capture debug information and share this information with the Avaya support team:
   a. Log in to the Cluster Control Manager (CCM) console as the customer user.
   b. To switch to the root user, type `su` and press **Enter**.
   c. Use the following command to edit the configmap:

      ```
      k edit cm -n mstr srv-configmap
      ```

      Use **i** for insert mode.
   d. Use the following command to update the property UPD_DEBUG from False to True:

      ```
      UPD_DEBUG: "false" →  UPD_DEBUG: "true"
      ```
   e. Note the podname for use in future sessions using the following command:

      ```
      export mstrsrvpod=$(k get pods -n mstr -o custom-columns=":metadata.name" |
      grep mstr-srv)
      ```
   f. Restart the mstr_srv pod using the following command:

      ```
      k delete pods -n mstr $mstrsrvpod
      ```
   g. Note the new podname for use in future sessions using the following command:

      ```
      export mstrsrvpod=$(k get pods -n mstr -o custom-columns=":metadata.name" |
      grep mstr-srv)
      ```
   h. Save a copy of the configmap using the following command:

      ```
      mkdir /var/avaya/artifactCache/mstr/$mstrsrvpod/
      k get cm -n mstr srv-configmap -o yaml > /var/avaya/artifactCache/mstr/
      $mstrsrvpod/srv-configmap.txt
      ```
   i. Copy the mstr_srv log directory to CCM after two to three minutes using the following command:

      ```
      export mstrsrvpod=$(k get pods -n mstr -o custom-columns=":metadata.name" |
      grep mstr-srv)
      k cp -n mstr $mstrsrvpod:/mnt/log/mstr/ /var/avaya/artifactCache/mstr/
      $mstrsrvpod/
      ```
   j. Ensure that the logs and configmap are available, and then turn off debug.

The mstr_srv pod fails to get to a running 1/1 state

Condition

The mstr_srv pod fails to reach a running 1/1 state with the exception - another user account or group uses the LDAP distinguished name.

Cause

The process imports the user and group names while configuring LDAP in Historical Reporting. This import can create duplicate groups in the Historical Reporting metadata. The issue occurs when these duplicate groups try to apply the unique LDAP DN, which is not allowed.

Solution

1. Take a backup of mstr-md pod and save it to a secure location on Cluster Control Manager (CCM).

2. Log in to the Cluster Control Manager (CCM) console as the customer user.
3. To switch to the root user, type `su` and press **Enter**.
4. To get a list of all the Historical Reporting pods, run the following command:

   ```
   kubectl get pods -n mstr
   ```

   Confirm the mstr-srv pod is in a running 1/1 state.
5. If the mstr-srv pod is in a running 0/1 state, use the following command to edit the mstr-srv configmap:

   ```
   kubectl edit cm -n mstr srv-configmap
   ```
6. Update the property UPD_GROUPS to False using the following command:

   ```
   UPD_GROUPS: "false"
   ```

   Use **i** to make the configmap editable.
7. To escape the edit mode, press the **Esc** button.
8. Type `:wq!` and press **Enter** to save the changes.
9. Restart the mstr-srv pod using the following command:

   ```
   kubectl delete pods -n mstr mstr-srv-<pod-id>
   ```
10. Confirm that the mstr-srv pod gets to a running state which can take more than 10 minutes.
11. Log in to the Historical Reporting server administrator page https://cluster.fqdn/AvayaAnalytics/servlet/AnalyticsServerAdmin using the administrator user account credentials.
12. To display the list of user groups, click **Server icon** > **User Manager**.

    Note:

    User groups are shown as a two-person icon. Users are shown as a single-person icon.
13. Confirm that the default Advanced, Basic, and Consumer groups are present.
14. Confirm if any additional groups have the customer's LDAP group name.
15. To confirm if there are LDAP users and LDAP user groups, click **Analytics Groups** > **LDAP Users**.
16. To view/edit a user group properties, right-click the user group and select **edit**.
17. On the **Project Access**, **Members**, and **Authentication** tabs, confirm the following for each user group:
    - **Project Access**: Security roles are selected.
    - **Members**: User group has local user members.
    - **Authentication**: LDAP login is the DN of the user group on the LDAP server.
18. Delete user groups that have no security roles selected and have the DN of an LDAP group in the Authentication tab.
19. To delete a user group, right-click the user group and select **delete**.

    Confirm deletion of the group.

Note:

- On completion, there are five user groups under **User Manager** (Advanced, Basic, Consumer, Analytics Groups, and Everyone).
- The **LDAP Users** under **Analytics Groups** should contain LDAP users and no LDAP groups.

20. Log out of the Historical Reporting server administrator page.
21. To edit the mstr-srv configmap, run the following command:

    ```
    kubectl edit cm -n mstr srv-configmap
    ```

22. Update the property UPD_GROUPS to True using the following command:

    ```
    UPD_GROUPS: "true"
    ```

    Use `i` to make the configmap editable.

23. To escape the edit mode, press the **Esc** button.
24. Type `:wq!` and press **Enter** to save the changes.
25. Restart the mstr-srv pod using the following command:

    ```
    kubectl delete pods -n mstr mstr-srv-<pod-id>
    ```

26. Confirm that the mstr-srv pod gets to a running state which can take more than 10 minutes.
27. Confirm the login of the LDAP user.
28. In the Historical Reporting server administrator UI, confirm the LDAP login under the authentication tab of each user group.
29. Take a new backup of the mstr-md metadata and export it to a secure location.

A Historical Reporting pod fails to start

Condition

A Historical Reporting pod fails to start.

Cause

Persistent volume (PV) fails to attach to a node because PV is already attached to another.

PV is a cluster-wide resource that you can use to store data to persists beyond the lifetime of a pod.

Solution

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. Restart a pod using the following command:

   `kubectl delete pod <pod-name>`, here pod name is the name of the Historical Reporting pod. For example: `kubectl delete pod -n mstr mstr-srv-6954bb949f-trdfg`

   Deleting the pod creates a new version of the pod.

4. To confirm the status of all pods, run the following command:

   kubectl get pods

   The pod STATUS column must display as `Running` and the READY column must display as `1/1`.

5. **(Optional)** To get status of the pods in the namespace mstr, run the following command:

   kubectl get pods --namespace=mstr

Data does not display in Agent by Routing Service report after a node restarts

Condition

If for some reason a node restarts, after you recover the node, the Agent by Routing Service historical report does not display data.

Solution

1. Restart both the pods for the Agent by Routing Service measure processor.
2. Verify whether you can now view data in the Agent by Routing Service report.

Multiple pods in Terminating state after node state changes to Not Ready

Condition

Multiple pods in a node display in `Terminating` or `Crashed` state after the node state changes to `Not Ready`.

Cause

This is a known issue. When Kubelet loses contact with the API server, the state of the node changes to `Not Ready`. When this happens, multiple pods go into `Terminating` or `Crashed` state.

Solution

1. Log in to the node that is in `Not Ready` state using the **cluster_ssh** console.
2. Type `sudo su` at the prompt and press **Enter**.
3. Enter the following command to restart the Kubelet service and restore all the pods:

   monit restart kubelet

   The terminated or crashed pods take 20 to 30 minutes to be restored to the normal state.

4. Enter the following command to verify whether all the components are working:

   ccm smoke-test

Troubleshooting general issues

Restarting a specific pod

About this task

You can restart an individual pod to troubleshoot all types of generic issues.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. Select the **General** option by entering the corresponding number.
6. To restart a specific pod, select the **Restart a specific pod** option by entering the corresponding number.
7. In the **Proceed to pod selection for restart** field, type `y`.

   Typing `n` cancels the operation.
8. In the **Select which pod to restart** field, type the line number of the specific pod and press **Enter**.

   The selected pod restarts.
9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.

Unable to access Historical Reporting web page

Condition

Unable to access the Historical Reporting web page.

Cause

The HTTP requests that are using cluster IP address or FQDN cannot reach the services in the cluster.

Solution

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`

4. Select **Troubleshooting** by pressing the corresponding number.
5. To select the **General** option, enter the corresponding number.
6. To select the **Restart keepalived-vip pods** option, enter the corresponding number.
7. In the **Proceed with Restarting Keepalived-vip pods** field, type `y`.

   Typing `n` cancels the operation.

   The Avaya Common Services keepalived-vip pod restarts.
8. Return to the previous page by entering `b`.
9. Quit the current page by entering `q`.
10. Return to the main menu by entering `m`.

Generating High Availability pod status report

About this task

Use this procedure to view the status of the pod in Avaya Analytics™ High Availability (HA). If you see any errors, contact Avaya support.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as a cust user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **General** by pressing the corresponding number.
6. Select the **HA pod status Report** option by entering the corresponding number.
7. In the **Proceed with generating HA Pod Status report** field, type `y` and wait for the report to generate.

   Typing `n` cancels the operation.

   The CCM console displays the detailed report.
8. Return to the previous page by entering `b`.
9. Quit the current page by entering `q`.
10. Return to the main menu by entering `m`.

Restarting standby pods and checking pod status

About this task

You must restart the standby pods after a node is shut down. Restarting the standby pods distributes the Avaya Analytics™ standby pods across three nodes instead of two for efficient load balancing of the Avaya Analytics™ pods.

Note:

This procedure is only applicable to Avaya Analytics™ High Availability (HA) deployments.

Before you begin

Check the status of the HA pods.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as a customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **General** by pressing the corresponding number.
6. To restart the standby pods, select the **Restart Standby Pods** option by entering the corresponding number.
7. In the **Proceed to Standby pods restart?** option, enter `y`.

   This step checks the logs of the standby pods for their HA status and displays a list of the standby pods or the services that the script restarts.
8. In the **Are you sure you want to restart the services listed above?** option, enter `y`.

   This step restarts all the services in a sequential order. After the completion of the process, CCM console displays `DONE`.
9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.
12. To check the status of all pods, in the CCM console, run the following command:

    `kubectl get pods | grep orca`
13. Ensure that the STATUS column of the pods displays `Running` and READY column displays `1/1`.
14. **(Optional)** To monitor the pods that you have restarted, run the following command:

    `kubectl get pods -w | grep orca`

Restarting measure processors

About this task

In Avaya Analytics™, on the completion of an engagement, the engagement gets added to the database. An engagement can have multiple events. When an event occurs, measures get incremental, which creates the following differences:

- The count of selected measures does not match the number of items when you drill-down in a report.
- If the engagement is still in process for a measure, the browser displays a `No Results` warning.

To resolve this issue, you must restart the measure processors, the Reliable Eventing Framework (REF) input adaptor, and the interval controller. The restart process also ensures efficient load balancing of the pods.

Warning:

You must restart the measure processors only during a maintenance window.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as a cust user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **General** by pressing the corresponding number.
6. To restart the measure processors, select the **Restart Measure Processors** option by entering the corresponding number.
7. In the **Proceed to Measure Processing Pods restart** option, type `y`.

   CCM console displays the following list of services that restart after you confirm to proceed:

   - The measure processors
   - The REF input adaptor
   - The interval controller
8. In the **Are you sure you wish to continue?** option, enter `y`.

   Note:

   This step resets the real-time data.

   This step restarts all the services in a sequential order. After the completion of the process, CCM console displays the message as `DONE`.
9. Return to the previous page by entering `b`.

10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.
12. To check the status of all pods, in the CCM console, run the following command as a customer user:

    `kubectl get pods | grep orca`

    The services that you restarted displays the STATUS column as `Running` and the READY column as `0/1`.
13. Wait for the READY column to display `1/1`.
14. To monitor the readiness of the pods that you restarted, run the following command:

    `kubectl get pods -w | grep orca`

    The services take about 15 minutes to restart. If the services do not restart after 15 mins, contact Avaya support.
15. Return to the previous page by entering `b`.
16. Quit the current page by entering `q`.
17. Return to the main menu by entering `m`.

Restarting all pods

About this task

Use this procedure to restart all pods for troubleshooting or after making any updates.

Warning:

Restarting the pods cause a total Avaya Analytics™ outage. You must perform this procedure only during a maintenance window. The pods are available again only after the restart.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as a cust user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **General** by pressing the corresponding number.
6. Select the **Restart ALL pods** option by entering the corresponding number.
7. In the **Proceed with Restarting ALL Pods** field, type `y`.

   Typing `n` cancels the operation.
8. At the prompt, confirm the restart by typing `y`.

Typing `n` cancels the operation.

Wait for the restart to complete.

9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.

Viewing Avaya Analytics™ certificate expiry date

About this task

Use this procedure to view the Avaya Analytics™ certificate expiry date:

Procedure

1. Log in to the Cluster Control Manager (CCM) console as a customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **General** by pressing the corresponding number.
6. Select the **Analytics Certificate expiry date** option by entering the corresponding number.

   The CCM console displays the expiry date of the Avaya Analytics™ certificate.
7. Return to the previous page by entering `b`.
8. Quit the current page by entering `q`.
9. Return to the main menu by entering `m`.

Viewing the list of all pods

About this task

Use this procedure to view the list of all the pods

Procedure

1. Log in to the Cluster Control Manager (CCM) console as a cust user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **General** by pressing the corresponding number.

6. Select the **List all pods for each node** option by entering the corresponding number.

   The CCM console displays the list of pods number with on the respective nodes.
7. Return to the previous page by entering `b`.
8. Quit the current page by entering `q`.
9. Return to the main menu by entering `m`.

Checking for missing Kafka topics

About this task

Use this procedure to check the list of available and missing Kafka topics. The script automatically creates the missing Kafka topic, if required.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as a customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **General** by pressing the corresponding number.
6. Select the **Check for missing Kafka Topics** option by entering the corresponding number.
7. In the **Proceed with checking for missing Kafka Topics** field, type `y`.

   Typing `n` cancels the operation.

   The CCM console displays the list of available Kafka topics.

   Note:

   - If all the Kafka topics are available, the CCM console returns to the main menu.
   - If any Kafka topic is missing, the CCM console displays the prompt to specify the deployment type, High Availability (HA) or non-High Availability (non-HA).
8. At the prompt, type the required deployment type.

   The CCM console displays that the missing Kafka topic is created with the required details.
9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.

Restarting Database Manager

About this task

Use this procedure to restart Database Manager.

Note:

This procedure is only applicable to Avaya Analytics™ High Availability (HA) deployments.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as a cust user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **General** by pressing the corresponding number.
6. To restart the Database Manager, select the **Restart Database Manager** option by entering the corresponding number.
7. In the **Proceed to Database Manager pod restart?** option, enter `y`.

   This step restarts the Database Manager.
8. Return to the previous page by entering `b`.
9. Quit the current page by entering `q`.
10. Return to the main menu by entering `m`.
11. To check the status of all pods, in the CCM console, run the following command as a customer user:

    `kubectl get pods | grep dbmgr`

    The dbmgr pod status is `0/1 Completed` once the job completes running.
12. **(Optional)** To monitor the pods that you restarted, run the following command:

    `kubectl get pods -w | grep dbmgr`

VMware reports consumption of thin provisioned storage on Avaya Analytics™ for Oceana® database

Condition

VMware reports consumption of thin provisioned storage on the Avaya Analytics™ for Oceana® database.

Cause

With thin provisioning, `vmdk` alarm is falsely triggered when you provision the disk space that is less than the required disk space in the footprint.

Solution

You must run the following script to determine the actual usage. You also need to monitor and increase the provisioned disk when actual usage approaches the currently provisioned limit.

```
Database troubleshooting: Check disk space of pgdata
```

Analytics 4.1.0.1 does not receive events from Oceana 3.7.0.1 when Async is enabled

Condition

When Avaya Analytics™ tries to subscribe to a UCM topic that does not exist in Oceana 3.7.0.1, it does not receive events and notifications from Avaya Oceana®.

Cause

The problem occurs because of incorrect settings in the deployment spreadsheet, when you select the correct Oceana version 3.7.0.1, but incorrectly set the UCM subscription notification value.

Solution

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. To switch to the root user, type `su` and press **Enter**.
3. To edit the input adapter configuration map, run the following command:

   `kubectl edit cm orca-ref-input-adaptor`
4. In `command` mode in a vi editor, enter the following command:

   `%s/SEND_NOTIFICATION/FORWARD_NOTIFICATION/g`
5. Run the following command to save the changes: **`wq`**
6. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
7. Select the **Restart Ref Input Adaptor** option by entering the corresponding number to restart the Input Adaptor pods.

   The restart might take several minutes to complete.
8. In the **Proceed with REF Inpur Adapter restart** field, type `y`.

   Typing `n` cancels the operation.
9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.

Deletion of ORCA product causes deleting the PV data

Cause

When you delete the product ORCA or any other product, the data related to PV also gets deleted. In the case of product ORCA where your data is stored, you need to reinstall Avaya Analytics™ for Oceana®.

Solution

Before deleting product ORCA you need to contact the Analytics design team for support on how to avoid deleting the PV and the data that it contains.

Changing Programmatic Account Password

Resetting Tomcat password on historical reporting

About this task

Use this procedure to reset Tomcat password on historical reporting.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as a cust user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **General** by pressing the corresponding number.
6. To select the **Change Programmatic Account Password** option, enter the corresponding number.
7. To reset the Tomcat password, select the **Reset Tomcat Password on Historical Reporting** option, enter the corresponding number.
8. In the **Continue with changing password?** option, enter `y`.
9. In the **Enter current Tomcat Password** option, enter your current Tomcat password.
10. In the **Enter new Tomcat Password** option, enter a new password of your choice.

    The Tomcat password is changed. The mstr-web pod is restarted.
11. Return to the previous page by entering `b`.
12. Quit the current page by entering `q`.
13. Return to the main menu by entering `m`.

Resetting EASG password on historical reporting

About this task

Use this procedure to reset EASG password on historical reporting.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as a cust user.
2. Switch to being the root user by entering the command `su`.

3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`

4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **General** by pressing the corresponding number.
6. To select the **Change Programmatic Account Password** option, enter the corresponding number.
7. To reset the EASG password, select the **Reset EASG Password on Historical Reporting** option, enter the corresponding number.
8. In the **Continue with changing password?** option, enter `y`.
9. In the **Enter current EASG Password** option, enter your current EASG password.
10. In the **Enter new EASG Password** option, enter a new password of your choice.

    The EASG password is changed. The mstr-srv pod and mstr-web pod are restarted.

11. Return to the previous page by entering `b`.
12. Quit the current page by entering `q`.
13. Return to the main menu by entering `m`.

Resetting postgres Password on the database

About this task

Use this procedure to reset postgres password on historical reporting database.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as a cust user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`

4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **General** by pressing the corresponding number.
6. To select the **Change Programmatic Account Password** option, enter the corresponding number.
7. To reset the postgres password, select the **Reset postgres Password on the Database** option, enter the corresponding number.
8. In the **Continue with changing password?** option, enter `y`.
9. In the **Enter current postgres Password** option, enter your current postgres password.
10. In the **Enter new postgres Password** option, enter a new password of your choice.

    The postgres password is changed. The mstr-web pod is restarted.

11. Return to the previous page by entering `b`.
12. Quit the current page by entering `q`.
13. Return to the main menu by entering `m`.

Resetting mduserservice Password on the database

About this task

Use this procedure to reset mduserservice password on historical reporting database.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as a cust user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **General** by pressing the corresponding number.
6. To select the **Change Programmatic Account Password** option, enter the corresponding number.
7. To reset the mduserservice password, select the **Reset mduserservice Password on the Database** option, enter the corresponding number.
8. In the **Continue with changing password?** option, enter `y`.
9. In the **Enter current mduserservice Password** option, enter your current mduserservice password.
10. In the **Enter new mduserservice Password** option, enter a new password of your choice.

    The mduserservice password is changed. The mstr-srv pod is restarted.
11. Return to the previous page by entering `b`.
12. Quit the current page by entering `q`.
13. Return to the main menu by entering `m`.

Resetting dbwriterservice Password on the database

About this task

Use this procedure to reset dbwriterservice password on the database.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as a cust user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

```
ccm release orca analytics
```

4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **General** by pressing the corresponding number.
6. To select the **Change Programmatic Account Password** option, enter the corresponding number.
7. To reset the tomcat password, select the **Reset dbwriterservice Password on the Database** option, enter the corresponding number.
8. In the **Continue with changing password?** option, enter `y`.
9. In the **Enter current dbwriterservice Password** option, enter your current Tomcat password.
10. In the **Enter new dbwriterservice Password** option, enter a new password of your choice.

    The dbwriterservice password is changed. All the services are restarted, which takes up to 15 mins.
11. Return to the previous page by entering `b`.
12. Quit the current page by entering `q`.
13. Return to the main menu by entering `m`.

Restarting Avaya Analytics™ after an Avaya Oceana® restart

About this task

If Avaya Oceana® restarts for any reason, do the following steps to ensure that Avaya Analytics™ continues to be functional.

For the detailed steps, see the *Deploying Avaya Analytics™ for Avaya Oceana®* document.

Before you begin

- Delete the Reliable Eventing group.
- Recreate the Reliable Eventing group.
- Restart Avaya Analytics™ Orca pods.

Warning:

You must restart the measure processors only during a maintenance window.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as a cust user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   ```
   ccm release orca analytics
   ```

4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **General** by pressing the corresponding number.
6. To restart the measure processors, select the **Restart Measure Processors** option by entering the corresponding number.
7. In the **Proceed to Measure Processing Pods restart** option, type `y`.

   CCM console displays the following list of services that restart after you confirm to proceed:

   - The measure processors
   - The REF input adaptor
   - The interval controller
8. In the **Are you sure you wish to continue?** option, enter `y`.

   Note:

   This step resets the real-time data.

   This step restarts all the services in a sequential order. After the completion of the process, CCM console displays the message as `DONE`.
9. Return to the previous page by entering `b`.
10. Quit the current page by entering `q`.
11. Return to the main menu by entering `m`.
12. To check the status of all pods, in the CCM console, run the following command as a customer user:

    `kubectl get pods | grep orca`

    The services that you restarted displays the STATUS column as `Running` and the READY column as `0/1`.
13. Wait for the READY column to display `1/1`.
14. To monitor the readiness of the pods that you restarted, run the following command:

    `kubectl get pods -w | grep orca`

    The services take about 15 minutes to restart. If the services do not restart after 15 mins, contact Avaya support.
15. Return to the previous page by entering `b`.
16. Quit the current page by entering `q`.
17. Return to the main menu by entering `m`.

Backup files gets recreated after deleting from Cluster Control Manager

Condition

Backup files gets recreated automatically even after you delete the files from Cluster Control Manager (CCM). The `/var/Avaya/artifcateCache/crunchybkp` folder displays the previous backup files with the new backup files

Solution

1. Log in to the Cluster Control Manager (CCM) console as a cust user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Database** by pressing the corresponding number.
5. Select the **Database backup** option by entering the corresponding number.
6. Select the **Delete all full backups** option by entering the corresponding number.
7. In the **Proceed deleting all backups** field, enter `y`.

   Entering `n` cancels the operation.
8. In the **Are you sure you want to delete all full backup files from directories** enter `y`.

   Entering `n` cancels the operation.

   This step clears:

   - The `/var/Avaya/artifcateCache/crunchybkp` directory on CCM.
   - The `/pgdata/full_backup/` directory on POD.

Error message seen in backup logs for Incremental backups

Condition

Below error message seen in backup logs for Incremental backups:

```
time="2022-07-07T15:24:41Z" level=info msg="stderr=[ERROR: [064]: unable
to write to remote-0 process on 'analyticsdb-node-0-8sl2-0.analyticsdb-
pods.default.svc.cluster.local.' write: [32] Broken pipe\nWARN: unable
to write to remote-0 process on 'analyticsdb-node-0-8sl2-0.analyticsdb-
pods.default.svc.cluster.local.' write: [32] Broken pipe\n]"
time="2022-07-07T15:24:41Z" level=fatal msg="command terminated with
exit code 64"
```

Cause

One of the database pods is not in a **Running** state. For Incremental Backups to complete successfully, all database pods needs to be in **Running** state with no errors.

Solution

1. pgBackRest re-runs an Incremental Backup, if it fails at the first time.
2. If a database pod is restarting when the backup is taken, you can monitor if there are new backup jobs created. Run the following command as a root user:

```
kubectl get po | grep analytics
analyticsdb-backup-4m92--1-7qn2b 1/1 Running 0 3m31s
analyticsdb-backup-4m92--1-b2fxm 0/1 Error 0 3m41s
analyticsdb-backup-4m92--1-lswnz 0/1 Error 0 4m20s
```

3. You can see 2 failed backup jobs with the latest job running. To view its log, you can run the following command as a root user:

```
kubectl get logs analyticsdb-backup-4m92--1-7qn2b
```

4. To confirm the incremental backup is successful, do the following:
   a. Run the `ccm release orca analytics` script.
   b. Select **List Incremental Backups** option.

Unable to correctly restore backups after changing Postgres database password

Condition

When you change the Postgres database password for a user, any backups that the user took before the password change cannot be restored correctly.

Solution

1. Restore any old backups before changing the Postgres database password for the user.
2. Perform a new full backup after you change the Postgres database password.

Avaya Analytics DR Monitoring tool for replication and failover

Avaya Analytics™ Disaster Recovery monitoring tool

Avaya Analytics™ offers a Disaster Recovery (DR) monitoring tool that enables administrators to check the replication status between DC1 and DC2 pods. The administrator can view the replication status on DC1 and troubleshoot the system to determine if it is in a good state before and after configuring Geo-Replication on DC2.

The Avaya Analytics™ DR monitoring tool helps in troubleshooting the following issues:

- Database split-brain: Identify if two pods in a cluster run simultaneously in a master or active mode.
- Network targets: Identify if the hosts or IP addresses the components communicate with are of the expected value.
- Local Cluster Status: Identify the status of the current cluster, such as Primary or Geo-Standby.
- NFS target: Identify the cluster of NFS servers used for Geo-Replication.
- Remote Cluster Status: Identify the status of the remote cluster.

- Replication delays: Know the update timestamp of the receipt of the WAL file. Use this information to determine if there is a lag in replication due to a slow network or if the hardware cannot process all the replicated data in a reasonable time.

Troubleshooting using DR monitoring tool

About this task

Use this procedure to check the status of replication between primary and replica pods on DC1 and between DC1 and DC2.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To run the `Analytics Administration` script, use the following command:

   `ccm release orca analytics`
4. Select **Troubleshooting** by pressing the corresponding number.
5. Select **Database** by pressing the corresponding number.
6. To check Crunchy pod replication status, enter the corresponding number.

   The output from the script shows different information for each pod depending on whether a pod is a primary, a standby replica, or a Geo replica. The expectations are as follows:

   - The primary pod **crunchy-primary-service-orca-dbmgr-0** must be listed as `master list` of replicas.
   - The `master list` displays a list of replica pods listening to the primary pod.

Chapter 5: Troubleshooting third-party certificates issues

Specifying the identity certificates used for generating CSRs

About this task

You can generate the Certificate Signing Request (CSR) using:

- Default set of identity certificates for externally facing interfaces.
- Specific identity certificates for external and internal interfaces.

  If you are using specific identity certificates for external and internal interfaces, then perform this procedure before generating CSRs.

  Note:

  Do not use this procedure if you are using default set of identity certificates for externally facing interfaces.

Tip:

Contact Avaya support personnel for assistance if you want to override a certificate's keySize or subjectAltName values.

Procedure

Create a file called `csr-serviceId-list-file`. Copy the service IDs listed below into this file and save it.

`ccm-identity.pfx-ccm-identity.pfx`

rbac-service-certificate-default-rbac-service-idcert

async-file-transfer-filetransferkeystore

prometheus-node-exporter-certificate-default-common-services-prometheus-node-exporter-idcert

fluentd-certificate-default-fluentd-idcert

platform-kubernetes--service-account.pem-platform-kubernetes--service-account.pem

egressgateway-certificate-default-egressgateway-idcert

prometheus-operator-certificate-default-common-services-prometheus-operator-idcert

eventing-kafka-cp-kafka-connect-kafkaconnectidcert

kibana-certificate-default-kibana-logging-service-mtls-idcert

orca-database-rest-databaserestkeystore

alarming-service-certificate-default-alarming-service-mtls-idcert

mstrsrv-service-certificate-mstrsrv-idcert1

egressgateway-mesh-default-egress-mesh-idcert

eventing-kafka-cp-schema-registry-kafkaconnectidcert

common-services-auth-http-service-certificate-default-common-services-auth-http-idcert

alarming-service-certificate-alarmingdb-idcert-primary

common-services-authdb-certificate-authdb-primary-idcert

orca-jrnl-agent-login-logout-measure-proc-arbackeystore

eventing-kafka-eventing-operator-eventingoperatoridcert

eventing-kafka-cp-kafka-kafkaexternalidcert

kube-state-metrics-certificate-default-common-services-kube-state-metrics-idcert

orca-itd-agent-measures-proc-arbackeystore

orca-open-interface-kafka-interface-orca-open-interface-kafka-interface-key-store

prometheus-alertmanager-certificate-default-common-services-prometheus-alertmanager-idcert

async-oceana-adapter-asyncoceanakeystore

ingressgateway-certificate-default-ingressgateway-idcert

cmonitor-service-certificate-cmonitor-default

prometheus-certificate-default-common-services-prometheus-idcert

certmgmt-agent-certificate-document-idcert-cms

orca-cdr-measure-proc-arbackeystore

platform-kubernetes--admin.pem-platform-kubernetes--admin.pem

certmgmt-agent-certificate-document-idcert-db

orca-itd-agent-by-account-measure-proc-arbackeystore

grafana-certificate-default-common-services-grafana-idcert

certmgmt-agent-certificate-document-idcert-cma

logelasticsearch-certificate-default-logelasticsearch-mtls-idcert

clusterhc2-certificate-clusterhc-default

eventing-kafka-topic-operator-kafkaconnectidcert

eventing-istio-sidecar-eventing-istio-sidecar-idcert

ingressgateway-mesh-default-ingress-mesh-idcert

eventing-operator-istio-sidecar-eventing-operator-istio-sidecar-idcert

orca-itd-agent-by-routing-service-measure-proc-arbackeystore

orca-ref-input-adaptor-arbackeystore

eventing-kafka-cp-zookeeper-kafkaconnectidcert

orca-interval-controller-arbackeystore

orca-streams-rest-streamsrestkeystore

common-services-authdb-certificate-authdb-replica-idcert

orca-breeze-authentication-streamsrestkeystore

orca-itd-routing-service-group-measure-proc-arbackeystore

orca-vdn-measure-proc-arbackeystore

orca-itd-routing-service-measure-proc-arbackeystore

orca-itd-agent-group-measure-proc-arbackeystore

orca-trace-measure-proc-arbackeystore

pgo-license-certificate-idcert1.default.license

platform-kubernetes--kube-scheduler.pem-platform-kubernetes--kube-scheduler.pem

logelasticsearch-certificate-admin-default-logelasticsearch-mtls-idcert

orca-itd-agent-not-ready-reason-code-measure-proc-arbackeystore

async-aggregator-interface-asyncaggregatorkeystore

platform-kubernetes--kubernetes.pem-platform-kubernetes--kubernetes.pem

platform-kubernetes--kube-proxy.pem-platform-kubernetes--kube-proxy.pem

kibana-certificate-logging-service.default-cli-idcert

common-services-jwks-store-jwksstoreidcert

mstrweb-service-certificate-mstrweb-idcert1

alarming-service-certificate-alarmingdb-idcert-replication

orca-itd-site-measure-proc-arbackeystore

eventing-kafka-cp-kafka-kafkaidcert

orca-admin-data-service-arbackeystore

platform-kubernetes--kube-controller-manager.pem-platform-kubernetes--kube-controller-manager.pem

default-mesh-default-mesh-idcert

pilot-mesh-default-pilot-mesh-idcert

orca-streams-data-publisher-orcapublisherkeystore

Next steps

Generate CSRs. If you manually specified identity certificates for external and internal interfaces, run the following command:

```
ccm release cert-manager third-party-certs --generate-service-csr --list-file csr-serviceId-list-file --output-dir <output-directory>
```

Specifying trust stores for adding a third-party CA certificate

About this task

For the third-party CA certificate, you can use:

- Default set of trust stores for externally facing interfaces.
- Specific trust stores for external and internal interfaces.

  If you are using specific trust stores for external and internal interfaces, then perform this procedure before importing the third-party CA certificate. Avaya recommends using this procedure to specify trust stores if you manually specified identity certificates for external and internal interfaces, before generating CSRs.

  Note:

  Do not use this procedure if you are using default set of trust stores for externally facing interfaces.

Procedure

Create a file called `trust-store-serviceId-list-file`. Copy the trust store service IDs listed below into this file and save it.

certmgmt-agent-certificate-document-trustedcert-cms

alarming-service-certificate-alarmingdb-trustedcert-primary

orca-itd-agent-by-routing-service-measure-proc-arbactruststore

platform-kubernetes--kubernetes.pem-platform-kubernetes--kubernetes.pem_trustedcert

mstrsrv-service-certificate-mstr-mstrsrv-trustedcert

rbac-service-certificate-default-rbac-service-trustedcert

pgo-license-certificate-cacert1

ingressgateway-certificate-default-ingressgateway-trustedcert

async-aggregator-interface-asyncaggregatortrustore

platform-kubernetes--kube-scheduler.pem-platform-kubernetes--kube-scheduler.pem_trustedcert

common-services-authdb-certificate-authdb-replica-truststore

platform-kubernetes--admin.pem-platform-kubernetes--admin.pem_trustedcert

orca-database-rest-databaseresttruststore

logelasticsearch-certificate-default-logelasticsearch-mtls-trustedcert

orca-itd-agent-not-ready-reason-code-measure-proc-arbactruststore

orca-streams-data-publisher-orcapublishertruststore

prometheus-operator-certificate-default-common-services-prometheus-operator-trustedcert

clusterhc2-certificate-default-common-services-clusterhc-trustedcert

prometheus-node-exporter-certificate-default-common-services-prometheus-node-exporter-trustedcert

eventing-kafka-cp-zookeeper-kafkaconnecttruststore

orca-streams-rest-streamsresttruststore

orca-cdr-measure-proc-arbactruststore

orca-itd-agent-by-account-measure-proc-arbactruststore

orca-itd-routing-service-group-measure-proc-arbactruststore

cmonitor-service-certificate-default-common-services-cmonitor-trustedcert

alarming-service-certificate-alarmingdb-trustedcert-replication

orca-itd-agent-measures-proc-arbactruststore

platform-kubernetes--kube-controller-manager.pem-platform-kubernetes--kube-controller-manager.pem_trustedcert

prometheus-alertmanager-certificate-default-common-services-prometheus-alertmanager-trustedcert

fluentd-certificate-default-Fluentd-trustedcert

orca-vdn-measure-proc-arbactruststore

orca-breeze-authentication-streamsresttruststore

async-oceana-adapter-asyncoceanatrustore

mstrweb-service-certificate-mstr-mstrweb-trustedcert

eventing-istio-sidecar-eventing-istio-sidecar-trustcerts

orca-itd-site-measure-proc-arbactruststore

eventing-kafka-cp-kafka-connect-kafkaconnecttruststore

platform-kubernetes--service-account.pem-platform-kubernetes--service-account.pem_trustedcert

ingressgateway-mesh-default-ingress-mesh-trustedcert

common-services-auth-http-service-certificate-default-common-services-auth-http-trustedcert

eventing-kafka-cp-kafka-kafkatruststore

alarming-service-certificate-default-alarming-service-mtls-trustedcert

orca-open-interface-kafka-interface-orca-open-interface-kafka-interface-trust-store

certmgmt-agent-certificate-document-trustedcert-db

eventing-kafka-cp-schema-registry-kafkaconnecttruststore

orca-interval-controller-arbactruststore

orca-ref-input-adaptor-arbactruststore

eventing-kafka-eventing-operator-eventingoperatortruststore

common-services-authdb-certificate-authdb-primary-truststore

orca-itd-agent-by-routing-service-measure-proc-arbactruststore

orca-trace-measure-proc-arbactruststore

egressgateway-certificate-default-egressgateway-trustedcert

pilot-mesh-default-pilot-mesh-trustedcert

eventing-kafka-topic-operator-kafkaconnecttruststore

prometheus-certificate-default-common-services-prometheus-trustedcert

eventing-operator-istio-sidecar-eventing-istio-sidecar-trustcerts

kibana-certificate-default-kibana-logging-service-mtls-trustedcert

orca-jrnl-agent-login-logout-measure-proc-arbactruststore

egressgateway-mesh-default-egress-mesh-trustedcert

orca-itd-agent-group-measure-proc-arbactruststore

ccm-identity.pfx-ccm-identity.pfx_trustedcert

async-file-transfer-filetransfertrustore

default-mesh-default-mesh-trustedcert

certmgmt-agent-certificate-document-trustedcert-cma

orca-admin-data-service-arbactruststore

platform-kubernetes--kube-proxy.pem-platform-kubernetes--kube-proxy.pem_trustedcert

kube-state-metrics-certificate-default-common-services-kube-state-metrics-trustedcert

common-services-jwks-store-jwksstoretruststore

grafana-certificate-default-common-services-grafana-trustedcert

Next steps

Import third-party CA certificates and identity certificates simultaneously, run the following command:

```
ccm release cert-manager third-party-certs --add-certs --list-file \
trust-store-serviceId-list-file --ca-cert-file <third-party-CA_PEM_filename> \
--id-cert-dir /home/<customer_account>/id-cert-files
```

Alternatively, import third-party CA certificates separately, run the following command:

```
ccm release cert-manager third-party-certs --add-trustcert --list-file \
trust-store-serviceId-list-file --ca-cert-file <third-party-CA_PEM_filename>
```

Ingress-gateway is not displayed in the list of identity certificates

Condition

Sometimes, the command **ccmcertmgr --identity-certs** does not return the ingress-gateway details.

Solution

1. Connect to the *cert-manager-certmgmt-agent* pod by running the following command:

   kubectl exec -it cert-manager-certmgmt-agent-f9bb9dc68-vtf65 /bin/bash

2. Delete the `crdResourceVersion.properties` file in the *cert-manager-certmgmt-agent* pod by running the following command:

   rm -f /opt/avaya/certmgmt/agent/data/crdResourceVersion.properties

3. After a five-minute interval, run the following command:

   ccm release cert-manager crtmgr --identity-certs

List of trust stores in Common Service Platform

Solution

To list the trust stores in Common Service Platform (CSP), run the following command:

ccm release cert-manager getcerts --trusted-stores

List of identity certificates in Common Service Platform

Solution

To list the identity certificates in Common Service Platform (CSP), run the following command:

```
ccm release cert-manager getcerts --identity-certs
```

List of secrets

Solution

To list the secrets, run the following command:

```
k get secrets
```

View the contents of a secret

Solution

To view the contents of a secret, run the following command:

```
k get secret eventing-kafka-cp-kafka-identity-cert-secret -n avaya-kafka
-o yaml
```

Decode the content of a secret

Solution

To decode the content of a secret, run the following command:

```
echo 'MWYyZDFlMmU2N2Rm' | base64 --decode
```

View the contents of a PEM file

Solution

To view the contents of a `.pem` file, run the following command:

```
openssl x509 -in certificate.pem -text
```

View the contents of a CSR file

Solution

To view the contents of a `.csr` file, run the following command:

`openssl req -in mycsr.csr -noout -text`

Services fail to consume certificates renewed by Certificate Manager service

Condition

The Certificate Manager service generates certificates for various services in the cluster that expire in 2 years from initial installation date. Certificate Manager renews these certificates before it expires and make these certificates available to the services. It is up to the services to consume these newly updated certificates. However, some services in the cluster fails to consume these newly updated certificates.

Solution: Renew certificates manually

1. Login to the Cluster Control Manager as cust after installing the patch.
2. Identify expiration of currently installed certificates, run the following command:

   ```
   renewServiceCertificates --checkExpiration
   ```

   If the **renewServiceCertificates** tool is not used yet:

   ```
   Begin checking Certificate Manager services' certificates
   expiration!

   Services' certificates were created when the cluster was installed
   on 04/28/2020

   These certificates will expire 2 years from then on 04/28/2022

   Finished checking Certificate Manager services' certificates
   expiration!
   ```

   If the **renewServiceCertificates** tool is used:

   ```
   Begin checking Certificate Manager services' certificates
   expiration!

   Services' certificates were renewed on 05/02/2021

   These certificates will expire 2 years from then on 05/02/2023

   Finished checking Certificate Manager services' certificates
   expiration!
   ```

3. Optional: Renew service certificates before performing an upgrade/planned restart, run the following command:

```
renewServiceCertificates
```

Note:

Run this command in maintenance window only before performing an upgrade.

When prompted to confirm that an upgrade/restart will take place after running this command, enter **y**.

4. Optional: Renew service certificates without performing an upgrade, run the following command:

```
renewServiceCertificates –restartAllServices
```

Note:

Run this command in maintenance window.

When prompted, enter **y** to restart all services.

All services are restarted. Services starts using renewed certificates.

5. Check if the services are restarted, run the following command:

```
ccm smoke-test
```

Note:

If any of the service pods are not in a running state after an hour, contact Avaya support.

Chapter 6: Troubleshooting Avaya Analytics™ on Avaya Common Services

Cluster is in an unusable state (clients cannot connect to the cluster)

Condition

Clients are not able to connect to the cluster because Kubernetes certificates have expired.

To confirm that certificates have expired, enter the `ccm release common-services alarmctl -l alarmEvents` command and look for a `certficate has expired` error message.

```
Unable to HelmReleaseListingParser::refresh in allotted time exited
with: Error: Kubernetes cluster unreachable: Get "https://<cluster
FQDN>:8443/version?timeout=32s": x509: certificate has expired or
is not yet valid: current time 2022-11-03T20:32:31Z is after
2022-11-02T20:58:00Z
```

Cause

Kubernetes certificates have expired.

⚠ Caution:

You must plan a maintenance window to perform this task. The `ccm rotate-cluster-certificates` command is service affecting.

Solution

1. Log in to Cluster Control Manager with your customer account.
2. If you have not already done so, start a **`screen`** session.
3. Run the `ccm rotate-cluster-certificates` command.

   This command can take more than 60 minutes to complete.

Cluster Failed to Install Due to Invalid Cluster Configuration

Condition

After starting the installation from the command line of an SSH session, the installation aborts with the error "Invalid data in Cluster/Services configuration" . From the terminal output or in the log `/var/log/avaya/ccm-main.log` look for log entries that indicate invalid data in Cluster/ Services configuration.

Cause

Invalid configuration data is one of the main causes for Cluster Control Manager installations to fail. During the initial installation Cluster Control Manager validates against certain cluster configuration fields. Ensure that the solution spreadsheet contains correct data.

Solution

If the installation fails early before installing VMs , modify the solutions spreadsheet to include the correct data and run the command **`ccm install <cluster-config.xlsx>`**.

Service failure during an installation or upgrade

Condition

While installing or upgrading the solution, a service fails. You must delete and reinstall the service that failed before you can deploy another service.

Solution

1. Run the command `ccm delete <service-name>` to delete the service that failed.

   Only delete the failed service. Do not delete any other services.

2. Run the command `ccm upgrade spec <solution spreadsheet name>.xlsx` to reinstall the service.

   Use the same solution spreadsheet that you originally used for the installation or upgrade.

Upgrade process stalls

Condition

During an infra upgrade, the process stops and seems to have stalled.

Solution

1. Wait to confirm that the upgrade operation has stalled, then cancel the upgrade.

2. Run `ccm upgrade hard-reset` to clear the state.
3. Retry the upgrade.

Terminal Shell Timed Out

Condition

During an installation you are disconnected from your terminal.

Cause

The session timed out. Use the Linux "screen" utility to run the installation in the background.

Solution

1. To resume the installation after losing the SSH session, run:

   ```
   $ screen
   [screen 0] $ ccm install resume
   <CTRL-a> d
   ```

   <CTRL-a> d detaches from the screen session and allows the installation to continue in the background.
2. To reattach to the installation screen session:
   a. At the CCM prompt, type **`screen -ls`** to retrieve the screen id of the session.
   b. Type **`screen -r <screen id>`** to reattach to the installation screen session.

DRS anti-affinity rule error during cluster installation

Condition

Cluster deployment failed because the user elected to provide vCenter credentials but the required vCenter node DRS anti-affinity rules were not correct for the High Availability check.

Cause

- The High Availability cluster does not have all VMs with the same anti-affinity rule.

Solution

1. In vCenter, ensure that all cluster node VMs have the same anti-affinity rule applied.

   For more information, see Creating a cluster node anti-affinity rule on page 178.
2. In Cluster Control Manager, run the **`ccm install resume`** command to resume the installation.

Related links

Creating a cluster node anti-affinity rule on page 178

Creating a cluster node anti-affinity rule

About this task

If you are deploying with high availability enabled, you must create an anti-affinity rule in vCenter for the cluster nodes.

Create the anti-affinity rule before powering on the VMs after initial OVA deployment to prevent a hot migration.

Procedure

1. Navigate to your vCenter cluster.
2. In the **Configure** tab, select **Configuration** > **VM/Host rules** > **Create VM/Host Rule**.
3. Create a name for the anti-affinity rule.
4. Select **enable rule**.
5. Type `Separate Virtual Machines`.
6. Add all the cluster nodes to the list.
7. Click **OK**.

Related links

DRS anti-affinity rule error during cluster installation on page 177

Third-party Certificates Are Not Being Used

Condition

Customer supplied third-party certificates for the solution are not being used by the solution.

Solution

1. Consult your solution documentation to verify that correct procedures were followed during installation for using third-party certificates.
2. If correct procedures were followed, contact Avaya support.

License node error message displayed at login

Condition

When you log in to Cluster Control Manager, the node license state is a value other than Normal.

Important:

After the 30-day licensing grace period elapses, the Common Services cluster is uninstalled. Product data is not preserved.

Error messages

The following are examples of error messages that might be displayed when the licensing state is Grace Period.

- If no Common Services node license is available, you will see a message similar to the following:

  ```
  Node License Error - Grace Period!
  The cluster has 29 days left in CSP Node License violation grace
  period!
  ```

- If only some of the required Common Services node licenses are available, you will see a message similar to the following:

  ```
  Node License Error - Grace Period!
  The cluster is only able to acquire 1 license for the 3 deployed
  nodes. Make sure the license server has an installed Common
  Services license with a sufficient node count for your deployment.
  Note, 1 node(s) will be deleted from the cluster to ensure the
  deployed node count matches the acquired license count.
  ```

When the licensing grace period ends, you might see an error message similar to the following:

```
Node License Error - RESTRICTED!!!
The 30-day licensing grace period has ended. All cluster nodes
have been shutdown because the cluster could not acquire licenses
for the 3 deployed nodes. Before powering up cluster nodes, make
sure the licensing server is reachable and has an installed Common
Services license with a sufficient node count for your deployment.
For more information about powering up the cluster, see your solutions
Maintaining and Troubleshooting documentation.
```

Solution

1. To resolve the licensing error, verify that the license is loaded on the WebLM server or on System Manager.
2. If you did not define the correct license server during deployment, run the `updateLicenseService <license_service>` command.

   `<license_service>` is the System Manager FQDN or IP address, or the WebLM service FQDN or IP address and port running inside the customer network. For example:

   - `smgr.example.com`
   - `weblm.example.com:52233`

   After running the **`updateLicenseService`** command, wait up to ten minutes for the license audit to refresh and license state to update.

Reported Alarms Not Seen on NMS

Condition

The Alarming service uses alarm definitions defined by other individual services to determine when to raise alarms upon reception of alarm events generated by the individual service. The alarms generated are expected to be received by the Network Management System (NMS).

Cause

The NMS might not have received the alarms because:

- The Alarming or Prometheus pods are not in a running state.
- The Alarming service is not properly configured.

Solution if Alarming or Prometheus pods are not in running state

1. Run the command **`ccm status --pod-details`**.
2. If the command shows the Alarming or Prometheus pod in a non-running state, contact Avaya support.

Solution if the Alarming service is not properly configured

1. Verify that the Alarming service has generated alarms by running the command **`ccm release common-services alarmctl -l alarmEvents`**.
2. Verify that the Alarming service destination is not empty by running the command **`ccm release common-services alarmctl -l destinations`**.
3. If the listing of the alarm request and alarm destination succeeds, but the alarm destination list is empty, you must configure the alarm destination. Run the command **`ccm release common-services alarmctl -a destinations --address <ip-address> --type <TRAP|KAFKA> --port <port-num> [--username <string>] [--password <string>] [--auth <SHA>] [--priv <AES>] [--authPwd <string>] [--privPwd <string>] [ --url <url-string> ]`**

   Refer to your solution documentation for additional information about using this command and setting the alarm destination.

Recovering a deleted virtual machine

Condition

You accidentally deleted a virtual machine in vCenter.

Solution

To recover the deleted virtual machine, contact Avaya support personnel.

In the future, do not delete virtual machines in vCenter.

Pods and services do not recover after a suspended VM is resumed

Condition

After a cluster node VM is suspended, some pods and services are not able to fully recover when the VM is resumed (powered on).

A restart of the cluster node enables the pods and services to recover.

Cause

Suspending and resuming a virtual machine in VMware is not supported.

Solution

1. Run the command **`ccm install cancel`**.
2. To recover, restart the cluster node.

Pod crash alarms

Condition

A KubePodCrash alarm is generated. If pods continuously restart, reports can have incorrect or missing data, or reports fail to run.

Cause

- Causes for this type of alarm include:
  - Service configuration error
  - Container ran out of memory
  - Application logic error
  - A resource that is expected to be running is currently not available
  - The liveness probe is failing

Solution

If the alarm persists for more than 1 hour, contact Avaya support.

Pod remains in init state

Condition

A pod remains in initialization (init) state after deployment and displays the following message:

```
Failed create pod sandbox.
```

Solution

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To restart the pod, run the following command:

   `kubectl delete pod <pod-name>`, where pod name is the name of the pod in init state. For example, **`kubectl delete pod alarming-db-common-services-7d64b69879-7drdq`**
4. To confirm that the pod is functional, run the following commands:

   `kubetl get pods`

Logging in and out of Kibana

About this task

Use this procedure to log in and log out of Kibana (OpenSearch) with your Keycloak user account.

Before you begin

Create a Keycloak user account.

Procedure

1. Go to the Kibana login page.

   To access the correct login page, go to `https://<cluster-FQDN>/logging`.
2. Enter your credentials.

   Enter the Keycloak user name and password.

   You are logged in to the OpenSearch Dashboards website on Cluster Control Manager.
3. To log out, select **Log out** from the user menu and close the browser tab.

   To avoid a known issue, close the browser tab after logging out. When you need to log in again, go to `https://<cluster-FQDN>/logging` using a new browser tab.

Error when logging in to Kibana

Condition

Unable to log in to Kibana after logging out.

Cause

Kibana does not correctly redirect to the login screen after a user logs out.

Solution

1. Go to the Kibana login page.

   In a new browser tab, access the login page at `https://<cluster-FQDN>/logging`.
2. Enter your credentials.

   Enter the Keycloak user name and password.

Restarting Kafka pod after multiple K8s node restarts

About this task

Use this procedure to restart Kafka pod after multiple K8s Node Restarts.

Procedure

1. To restart the pod, delete the following command: **`kubectl delete pod -n avaya-kafka eventing-kafka-cp-kafka-0(/1/2)`**
2. Wait for a while and if the pod still not coming to ready state do the following:

a. Scale down Kafka service to 0

`kubectl scale sts -n avaya-kafka eventing-kafka-cp-kafka --replicas=0`

b. From the pod describe output (done in the issue identification step) get the PV name.

c. k exec --namespace=piraeus deployment/piraeus-op-piraeus-operator-cs-controller -- linstor r l -r pvc-d2a17871-27b1-46e7-b64b-196547ff51ca

- Replace the PV Name in the above command with the pod describe output
- From the output find the Node on which the PVC is "inuse"

```
+-------------------------------------------------------------------------------
-----------------------------------------------+

| ResourceName                              | Node                          | Port |
Usage  | Conns |    State | CreatedOn           |

|
=================================================================================
================================================|

| pvc-d2a17871-27b1-46e7-b64b-196547ff51ca | node5336.puneccq.avaya.com | 7024 |
Unused | Ok    | UpToDate | 2022-07-16 05:53:09 |

| pvc-d2a17871-27b1-46e7-b64b-196547ff51ca | node5337.puneccq.avaya.com | 7024 |
Unused | Ok    | UpToDate | 2022-07-16 05:52:59 |

| pvc-d2a17871-27b1-46e7-b64b-196547ff51ca | node5338.puneccq.avaya.com | 7024 |
InUse  | Ok    | Diskless | 2022-07-16 05:53:03 |
```

3. Get the name of the satellite pod running on the node identified in previous step

```
k get pods -n piraeus -o wide | grep ns
```

```
piraeus-op-piraeus-operator-ns-node-chdh9                      2/2     Running
4 (2d13h ago)   6d8h    10.133.53.37    node5337.puneccq.avaya.com
<none>           <none>

piraeus-op-piraeus-operator-ns-node-sk5dh                      2/2     Running
2 (3d5h ago)    6d8h    10.133.53.38    node5338.puneccq.avaya.com
<none>           <none>

piraeus-op-piraeus-operator-ns-node-vh59k                      2/2     Running
4 (2d12h ago)   6d8h    10.133.53.36    node5336.puneccq.avaya.com
<none>           <none>
```

4. Disconnect the InUse resource using the drdbadmin cli

   kubectl exec -n piraeus <satellite pod name> -c linstor-satellite -- drbdadm disconnect <pv-name> –force

5. Check the status of the resource again and it should be in Unused state (excuse the command in step 2.c. and the usage should be "Unused" for all 3).
6. Connect the resource back

   kubectl exec -n piraeus <satellite pod name> -c linstor-satellite -- drbdadm connect <pv-name>

7. Scale up Kafka

   kubectl scale sts -n avaya-kafka eventing-kafka-cp-kafka --replicas=3

8. Wait for few minutes and check whether all the Kafka pod are in ready state

   kubcetl get pods -n avaya-kafka -o wide

Cluster Control Manager commands

Cluster Control Manager core commands

Some of the commands in the following table require the release name. You can obtain the release name from the Release column of the chart that is displayed when you run the `ccm status` command.

Command	Description
`ccm`	Provides command line syntax help for the **`ccm`** command level. The same as issuing the **`ccm help`** command.
`ccm archive help`	Displays `ccm archive` commands, including options for the `ccm archive config` command. For more information about these command options, see ccm archive command options on page 192.

Table continues…

Command	Description
`ccm backup`	Performs a backup. For information about the options available with this command, see ccm backup and ccm backup schedule command options on page 194.
`ccm config`	Provides command line syntax help for the `ccm config` command level. The same as issuing the `ccm config help` command.
`ccm config help`	Provides command line syntax help for the **`ccm config`** command level.
`ccm config show <release-name>`	Reads the configuration from the specified, running, helm release and displays it on the console in YAML format. Redirection of standard out can be used to capture the configuration data to a file. . For example: `ccm config show <example-helm-release > example-helm-release-config.yaml.`
`ccm config update <release-name> <config-file>`	Updates the specified, running, helm release with the configuration provided by the specified YAML configuration file. Redirection of standard out can be used to capture the configuration data to a file. For example: `ccm config update example-helm-release example-helm-release-config.yaml.`
`ccm delete <release> [--yes]`	Deletes the specified product chart after providing a confirmation prompt. If the optional argument `--yes` is used, the prompt is not provided. Only run `ccm delete <service name>` to delete a failed service in case of an installation or upgrade failure. For more information, see Service failure during an installation or upgrade on page 176.
`ccm help`	Provides command line syntax help for the **`ccm`** command level.
`ccm infra`	Provides command line syntax help for the **`ccm infra`** command level. Same as issuing the command **`ccm infra help`**.
`ccm infra help`	Provides command line syntax help for the **`ccm infra`** command level.
`ccm infra update-vcenter-creds`	Updates the vCenter credentials used by Cluster Control Manager to manage a deployed cluster.

Table continues…

Command	Description
`ccm install <cluster-config-xlsx-file> [--post-clean] [--force-download]`	Installs the cluster infrastructure and services as defined in the provided Excel workbook. The file must end in the .xlsx or .xlsm extention. The user must accept, the Avaya EULA before proceeding. This variation of the install command auto-downloads all artifacts as part of the installation process. If a given artifact is already in the `/var/avaya/artifactCache` directory, its download is skipped. If --post-clean is specified, product chart docker images that are downloaded as part of the services staging process are deleted from the local CCM cache after they are staged on the system. If --force-download is specified the download of artifacts is forced (copies in `/var/avaya/artifactCache` are first deleted).
`ccm install <resume | cancel> [--force-download]`	Resumes or cancels the installation of the cluster infrastructure. If --force-download is specified, the download of artifacts if forced (copies in `/var/avaya/artifactCache` are first deleted).
`ccm product`	Shows a list of staged products to which product-specific commands can be directed. Same as issuing the `ccm product help` command.
`ccm product <staged-product-name>`	Provides command line syntax help for the staged product identified in `<staged-product-name>`. Same as issuing the `ccm product <staged-product-name> help` command.
`ccm registry`	Provides command line syntax help for the `ccm registry` command level. Same as issuing the command `ccm registry help`.
`ccm registry purge <image_name> <image_tag>`	Deletes the docker container specified in `<image_name> <image_tag>` from the cluster's registry. **Note:** The `ccm registry purge <image_name> <image_tag>` command requires root access. Run `su - root` to access the root account.

Table continues…

Command	Description
`ccm registry repopulate <--use-local-images>`	Repopulates the docker image registry with docker images. When running this command: • If you do not set `--use-local-images`, the command pulls all images, which are associated with the charts staged on Cluster Control Manager, from the external (harbor) docker registry. It then pushes these images to the docker registry running in the cluster. • If you set `--use-local-images`, the command only uploads the docker images that are relevant for the charts deployed in the cluster. This command obtains the images from Cluster Control Manager and uploads them to the docker registry running in the cluster.
`ccm registry repos [--tags]`	Lists the docker repositories (containers) in the cluster's registry. If --tags is specified, each <image_name>:<tag> is printed.
`ccm registry tags <repository>`	Lists all the <image>:<tag> entries for the specified docker repository (image name) at the cluster registry.
`ccm release`	Shows a list of running helm releases (and the staged products they were installed from) to which release-specific commands can be directed. Same as issuing the `ccm release help` command.
`ccm release <helm-release-name>`	Provides command line syntax help for the help release identified by <helm-release-name>. Same as issuing the `ccm release <helm-release-name> help` command.
`ccm restore`	Performs the restore operation. For information about the options available with this command, see ccm restore command options on page 196.

Table continues…

Command	Description
`ccm report`	Enables the automatic collection of cluster and application-related information. The information collected is useful for troubleshooting purposes. For more information about this command, see ccm report command on page 197. **Note:** This command requires common-services chart version 1.1.0.0.x or later. On Cluster Control Manager, run `ccm status` to see which version of the chart is deployed.
`ccm smoke-test`	Runs a smoke test on the installed products. This command verifies that pods and contains are running as expected. It then pings the nodes and virtual IPs for proper operation. The exit status of the command is not zero if the test fails.

Table continues…

Command	Description
`ccm status [--health] [--pod-details] [--ps] [<product-name>]`	Shows general status information about the services deployed in the cluster. If the cluster is not in the installed state, information about the state of the cluster (installing, uninstalling, services upgrading) is displayed. If `--health` is specified, the program checks to make sure that all pods have a status of Running or Completed. **Note:** The command does not verify the state of the containers within the pod. If `--pod-details` is specified, the name, container count, restart count, and uptime for each pod is displayed. If the product name is specified, only pod details for that product are displayed. If `--ps` is specified, the status of the following Kubernetes processes on each cluster node is displayed: • `sdsetcd.service` • `etcd.service` • `kube-apiserver.service` • `kube-controller-manager.service` • `kube-scheduler.service` • `keepalived.service` • `flanneld.service` • `kubelet.service` • `kube-proxy.service` • `containerd.service`
`ccm system`	Provides command line syntax help for the `ccm system` command level. Same as issuing the `ccm system help` command.
`ccm system config`	Displays information about the configuration of Cluster Control Manager.
`ccm system log`	Provides command line syntax help for the **`ccm system log`** command level. Same as issuing the `ccm system log help` command.

Table continues…

Command	Description
`ccm system log config`	Shows the current configuration of Cluster Control Manager logging. Each row represents a "logger" (a log type) that is embedded in the Cluster Control Manager software, along with the logging filter level used by the logger. The filter level filters the logs that are injected into the logging system by the logger. Each column represents a log handler that receives logs from the logging system and writes them to the output device that it manages. The configuration for each log handler includes: • The name of its output file, if applicable • Whether each log line is prefixed • The log rotation parameters (size-in-MB, max-count) • The logging filter level used to filter the logs that the handler writes to the output device it manages • 'X' indicates that the logs injected by a logger for the given row are routed by the logging system to the handler for the given column The filter levels are NOTSET (no filtering), and in order of increasing priority, DEBUG3, DEBUG2, DEBUG1, INFO, WARNING, ERROR, and CRITICAL. Filtering is performed based on the priority of the generated log and the filter level. For example, a filter level of INFO will allow logs at levels INFO, WARNING, ERROR, and CRITICAL, but not at DEBUG3, DEBUG2, or DEBUG1. Use the command `ccm system log preset default` to set the log settings to the standard system default.
`ccm system log getlogs`	Gathers the CCM log files into a zip file placing the zip file into the user's current working directory. The output file is named `ccm-logs-<date>-<time>-<hostname>.zip`.
`ccm system log help`	Provides command line syntax help for the `ccm system log` command level.

Table continues…

Command	Description
`ccm system log listpresets`	Lists the available set of logging presets. A preset is a specific logging configuration. Available presets are: • Default: This preset refers to the factory default settings. This configuration is used on initial installation. • Debug: This preset is used to adjust settings to gather an additional level of detail. • Developer: This preset is useful for developers to analyze logs during software development.
`ccm system log preset preset-name`	Sets the logging configuration to the given preset.
`ccm uninstall resume`	Resumes the uninstallation of the cluster infrastructure (from a previous error exit).
`ccm uninstall [--force]`	Uninstalls the cluster infrastructure. Exits with an error if there are any helm releases running in the cluster, unless `--force` is specified, in which case the cluster uninstallation proceeds. **Tip:** Unless specified otherwise, always include `--force` when you run this command. The `--force` option deletes the deployed solution components along with the cluster virtual machines.
`ccm upgrade`	Provides command line syntax help for the `ccm upgrade help` command level. Same as issuing the `ccm upgrade help` command.
`ccm upgrade cancel [--no-prompt]`	Cancels a paused upgrade of services (product helm charts). A confirmation prompt is first given unless --no-prompt is specified.
`ccm upgrade hard-reset [--force]`	Performs a hard reset of the Cluster Control Manager services (helm chart) upgrade state machine. This command is used if the state machine experiences an unrecoverable fault. It is used in exception conditions only. An attempt is first made to have the state machine process the reset. If the state machine fails to process the reset, state machine tracking data is cleared and the system returns to the Installed state. If `--force` is specified, state machine data is immediately cleared (the state machine is bypassed).

Table continues…

Command	Description
`ccm upgrade system upgrade-config.yaml`	Upgrades the Cluster Control Manager machine. The system reboots after the upgrade is applied.
`ccm upgrade spec <complete path>/ <solution spreadsheet name>.xlsx --infra`	Upgrades the cluster. In `<complete path>`, provide the complete path to the solution configuration spreadsheet you are using for the upgrade.
`ccm upgrade spec <complete path>/ <solution spreadsheet name>.xlsx --products`	Upgrades the services in the solution. In `<complete path>`, provide the complete path to the solution configuration spreadsheet you are using for the upgrade.
`ccm version [-k] [-p <product_name>] [--runtime] [ --details] [-s] [-c]`	Displays the version(s) of software that is staged or installed on Cluster Control Manager and the Kubernetes cluster. Optional arguments: `-k`: Prints the version of Red Hat, Kubernetes, and dockers running on each of the Kubernetes nodes. `-p <product>`: Prints all versions of software of a specific product installed on Cluster Control Manager. `-s`: Shows the software versions of all the products deployed in the cluster. All other command options are ignored. `-c`: Prints the software versions of individual components (for example, charts) of all or specific products installed/staged on Cluster Control Manager. `--details`: Prints docker-image or binary versions associated with an individual component (charts/packages) of all or specific products. `--runtime`: Prints the runtime status of the software that is staged/installed on Cluster Control Manager.

ccm archive command options

You can configure archive destination settings using:

- The OVF template when deploying Cluster Control Manager
- The commands described in this section

All solution backup files are stored in the configured archive destination. Application and Cluster Control Manager data is stored in this location.

Command option	Description

`ccm archive help`	This command displays the available `ccm archive` command options. The following is an example of the output for this command:
	`config`
	`config [--help|-h]`
	`config remote [-h|--host <FQDN/IP> -p|--port -u|--user <username> -d| --directory <directory>]`
	`config local`
	`config sync`
	`config test`
	`help`
	`list`

`ccm archive list`	This command displays archive information. It includes archives from `ccmClusterBackup/*` and `cspLogs/*`. If the archive destination is set to local, only local archives are displayed. If the archive destination is remote, both local and remote archives are displayed. The output for this command includes the following information: • Location, which is either “Local” or the hostname (FQDN or IP address) of the remote server • Directory • File name • Last modified time • Size
`ccm archive config`	This command displays the configured archive destination. If the archive destination is set to remote, this command also displays the configured remote server details, including the FQDN or IP address, port number, user name, and base directory.
`ccm archive config local`	Run this command to set the archive destination to local.
`ccm archive config remote`	Run this command to set the archive destination to remote. You must also specify the details of your remote server. For example, run this command as follows: `ccm archive config remote -h <FQDN/IP> [-p <port>] -u <username> -d <directory path>`
`ccm archive config sync`	Run this command to synchronize the archive destination configuration to the cluster.

Table continues…

Command option	Description
`ccm archive config test`	If the archive destination is set to remote, you can use this command to test the connection to the remote server. This command confirms whether the remote server is accessible.

ccm backup and ccm backup schedule command options

You can run the `ccm backup` command on its own or with various options.

The following table lists the `ccm backup` command options.

Command option	Description
`ccm backup`	Use this command to perform a backup. If you run this command on its own, the backup file is stored in either the local or remote directory, depending on the configured archive destination.
`ccm backup setpassword`	Run this command to reconfigure the backup file password.
`ccm backup --local`	Use this command to store the backup file in the local directory at `/var/avaya/artifactCache/ccmClusterBackup`.
`ccm backup --remote-server`	Use this command to store the backup file on your remote server. You must also specify the details of your remote server. For example, run this command as follows: `ccm backup --remote-server "<FQDN/IP> [-p <port>] -u <username> -d <directory path>"` **Note:** `ccm backup --remote-server` replaces the `ccm backup --sftp-destination` command.
`ccm backup --local --password` or `ccm backup --remote-server --password`	Use the `--password` option if you want to override the configured backup file password. This command prompts you for a new backup password.
`ccm backup schedule`	This command displays the backup schedule.
`ccm backup schedule --history`	This command displays the backup schedule history.
`ccm backup schedule --sync`	Run this command to synchronize the backup schedule configuration to the cluster.

Table continues…

Command option	Description
`ccm backup schedule -t <time>`	Run this command to reconfigure the time when the scheduled backup will run. Enter the time in the hh:mm 24-hour time format. For example, replace `<time>` with `19:00` if you want the backup to run at 7pm.
`ccm backup schedule -d`	Run this command to set the backup recurrence to daily.
`ccm backup schedule -d [-r <recurrence> -t <time> -e <True|False>]`	When setting the backup recurrence, you can optionally include the `-r`, `-t`, and `-e` parameters to reconfigure other scheduled backup settings. • - `-r` is the Backup Recurrence Multiple setting. Enter a number, which defines how often the backup will run. For example, if you enter `3`, the backup will occur every three days. - For more information about `-t <time>`, see the `ccm backup schedule -t <time>` command description above. - `-e` indicates whether the scheduled backup is enabled. Enter `True` or `False`.
`ccm backup schedule -w <day_of_week>`	Run this command to set the backup recurrence to weekly. You must also specify the day of the week. For example, you can run `ccm backup schedule -w monday` if you want the backup to run on Mondays.
`ccm backup schedule -w <day_of_week> [-r <recurrence> -t <time> -e <True|False>]`	When setting the backup recurrence, you can optionally include the `-r`, `-t`, and `-e` parameters to reconfigure other scheduled backup settings. • `-r` is the Backup Recurrence Multiple setting. Enter a number, which defines how often the backup will run. For example, if you enter `3`, the backup will occur every three weeks. • For more information about `-t <time>`, see the `ccm backup schedule -t <time>` command description above. • `-e` indicates whether the scheduled backup is enabled. Enter `True` or `False`.

Table continues…

Command option	Description
`ccm backup schedule -m <date>`	Run this command to set the backup recurrence to monthly. Specify the date when you want the backup to run. For example, replace `<date>` with `5` if you want the backup to run on the 5th of the month. If you enter `31` and the month does not have 31 days, the backup will run on the last day of the month.
`ccm backup schedule -m <date> [-r <recurrence> -t <time> -e <True|False>]`	When setting the backup recurrence, you can optionally include the `-r`, `-t`, and `-e` parameters to reconfigure other scheduled backup settings. • `-r` is the Backup Recurrence Multiple setting. Enter a number, which defines how often the backup will run. For example, if you enter `3`, the backup will occur every three months. • For more information about `-t <time>`, see the `ccm backup schedule -t <time>` command description above. • `-e` indicates whether the scheduled backup is enabled. Enter `True` or `False`.

Related links

Backing up Common Services on page 125

ccm restore command options

The following table lists the `ccm restore` command options.

Command option	Description
`ccm restore <path to backup file> all`	Use this command to restore Common Services from either the local or remote directory, depending on the configured archive destination.
`ccm restore <path to backup file> ccm`	Use this command to restore Cluster Control Manager from either the local or remote directory, depending on the configured archive destination.
`ccm restore --local`	This command performs the restore operation from the local `<path to backup file>` directory.

Table continues…

Command option	Description
`ccm restore --remote-server`	This command performs the restore operation from your remote server. You must specify the details of your remote server. For example, to restore Common Services, run this command as follows: `ccm restore --remote-server "<FQDN/IP> [-p <port>] -u <username>" <path to backup file> all` **Note:** `ccm restore --remote-server` replaces the `ccm restore -- sftp-destination` command.
`ccm restore --password`	Use the `--password` option if you want to override the configured backup file password. This command prompts you for a new backup password.

Related links

Restoring Avaya Common Services in an online deployment environment on page 127
Restoring Cluster Control Manager on page 132

ccm report command

The **`ccm report`** command enables cluster and application-related information to be collected automatically. The generated report is stored in the local or remote directory, depending on the configured archive destination.

Usage options

You can run this command as follows:

```
ccm report [--help|-h ]
    [--hours <hours> | --days <days> | --from <iso-8601> --to <iso-8601>]
    [--local]
    [--remote-server "<FQDN/IP> [-p <port>] -u <username> -d <directory path>"]
    [--target <"instanceId -n <namespace> [any list of target options here]">]
    [--filter <"[-p <product>] [-r <release>] [-c <component> ] [-s <service>] [-o
      <pod>]">]
    [--skip <"[ccm] [infra] [app]">
    [--fileIntegrity (only collects file integrity logs; no other reports are collected)
    This option is only valid by itself or with the --remote-server option]
```

Time constraint options

When no time options are specified, `ccm report` collects a subset of available log files on Cluster Control Manager and nodes in the cluster. It also collects seven days of logs from applications. You can increase or decrease the logs from applications by specifying time constraints:

- `--hours <hours>` option specifies the number of hours of application logs to collect. For example, `ccm report --hours 3` collects application logs from the past three hours. This option cannot be combined with any other time constraint options.

- `--days <days>` option specifies how many days of application logs to collect. For example, `ccm report --days 2` collects application logs from the past two days. This option cannot be combined with any other time constraint options.
- `--from` and `--to` options specify a range of application logs to collect. For example, `ccm report --from 2020-05-04T20:40:03 --to 2020-05-04T21:40:03` collects one hour's worth of application logs for the specified date. Use `--from` and `--to` together and do not combine them with the `--hours` or `--days` options.

All system logs are collected regardless of the time constraints specified in the command.

Local option

Run `ccm report --local` to upload the collected artifacts to the local directory at `/var/avaya/artifactCache/cspLogs/ccmReport-<ccm server name>-<date>.tgz`.

Remote server option

The `--remote-server` option replaces the `--sftp-destination` option.

The `--remote-server` option automatically uploads collected artifacts to a remote SFTP server. When Cluster Control Manager collects an artifact, that artifact is uploaded to the remote server and deleted from Cluster Control Manager. This option is useful when a lot of logs need to be collected and there is insufficient disk space on Cluster Control Manager. The `--remote-server` option must include the following:

- FQDN or IP address of the remote server
- Port number if you are not using the default port 22
- User name and password for the remote server
- Base directory path to the remote server

The following is an example of this command:

```
ccm report --remote-server "10.10.5.7 -u bruce -d /home/bruce/system1/"
```

In this example:

- No port is specified, so the default SSH port 22 is used.
- The report is transferred to `/home/bruce/system1/ccmReport-<ccm server name>-<date>/`

When you run this command, you are prompted for a password. If your credentials are invalid, the command exits with an error message.

After the upload, the following artifacts are collected on the remote server:

Artifact name	Artifact description	Condition
ccmLogFiles.tar.gz	The following Cluster Control Manager logs: • /var/log/avaya/ccm/* • /var/log/avaya/ccm/remoteDesktopSession/* • /var/log/avaya/common_os/* • /var/log/messages* • /var/log/audit/* • /var/log/cmd_history.log • /opt/avaya/config/ccm-config.yaml • /opt/avaya/config/archiveConfig.yaml • /opt/avaya/config/scheduledBackupConfig.yaml • journalctl	Cluster Control Manager must be deployed
swversion	Software version information.	Cluster Control Manager must be deployed.
poddetails	Pod details.	The cluster must be deployed.
master_*-systemLogs.tar.gz	Cluster master node logs.	The cluster must be deployed.
worker*-systemLogs.tar.gz	Cluster worker node logs.	The cluster must be deployed with the worker node.
ccmSystemInformation.tar.gz	The following system information and Kubernetes logs: • systemInformation/kubectl_get_nodes • systemInformation/kubectl_describe_nodes • systemInformation/name_license_status • systemInformation/kubectl_api_resources	The cluster must be deployed.
systemLogs.tar.gz	Cluster node manager logs.	The cluster must be deployed.

Table continues…

Artifact name	Artifact description	Condition
`applicationLogs/*`	Application logs, including target logs.	The cluster and common-services must be deployed.

For information about obtaining the `.tar` log files, see Obtaining .tar log files on page 202.

Target option

The `--target` option enables the invocation of product-specific scripts. These scripts collect additional information that is not retrieved through the general `ccm report` command. For example, a custom script might collect database information for a product. When you specify the `--target` option, no other application logs are downloaded to the `applicationLogs/` directory.

The `--target` option must include the instance ID. For example, `ocra`. It can also include:

- `-n <namespace>`, which is used if the instance ID is not deployed in the default namespace. For example, if the instance ID is deployed in the acme namespace, the command requires `n acme`.
- A list of options passed directly to the custom script, which are not used by the `ccm report` command.

For example, `ccm report --target "ocra -n acme -o option1 option2 option3"` invokes the ocra script in the acme namespace with three options.

Elasticsearch logs

Elasticsearch data is gathered when you run `ccm report` without the `--target` option. This data is placed in the `applicationLogs/*` directory. For information about uploading these logs to view in Kibana or another Elasticsearch server, see Viewing application logs in an Elasticsearch server on page 202.

Filter option

By default, `ccm report` collects log files for all applications running in the cluster. To reduce time and disk space consumption, you can use the `--filter` option to specify a specific application or group of applications from the collection. You cannot combine the `--filter` option with `--target`.

When using `--filter`, include one or more of the following options within double quotes:

- `-p <product>`: Replace `<product>` with the product name. Examples of product names are CommonServices, Analytics, or mstr.
- `-r <release>`: Replace `<release>` with the release name or instance ID. Release examples are CommonServices, orca, or mstr.
- `-c <component>`: Replace `<component>` with the component or chart name. Examples of components are eventing-kafka-1.2.100001090700, istio-1.5.1010100011030005, orca-5.1.56, or mstr-0.1.216.

  Only use `-c` with the `-p` or `-r` options.
- `-s <service>`: Replace `<service>` with the service name. Examples of service names are istio-ingressgateway, orca-redis, or mstr-web.

- `-o <pod>`: Replace `<pod>` with the pod ID. Pod examples are orca-redis-server-1 or cert-manager-certmgmt-service-c5899b664-g8qjm.

Tip:

You can run `ccm version -c` to determine which products, releases, and components are installed in your cluster.

You can add multiple filter options to a single command. Separate each option with a comma. For example:

```
ccm report --filter -p <product> -c <component>, -s <service1>, -s <service2>
```

The following table provides examples of how to use the `ccm report --filter` command:

Command example	Description
`ccm report --filter "-p CommonServices"`	Collects logs for all containers deployed on the CommonServices product.
`ccm report --filter "-r orca"`	Collects logs for all containers deployed on the orca release.
`ccm report --filter "-p CommonServices -c istio-1.5.1010100011030005"`	Collects logs for all containers deployed on the istio-1.5.1010100011030005 chart for the CommonServices product.
`ccm report --filter "-s orca-redis"`	Collects logs for all containers from the orca-redis service.
`ccm report --filter "-o orca-streams-rest-7599c79656-2pfgv"`	Collects logs for all containers in the orca-streams-rest-7599c79656-2pfgv pod.
`ccm report --filter "-p CommonServices -c istio-1.5.1010100011030005, -s orca-redis, -s eventing-client"`	Collects logs for all containers deployed on the istio-1.5.1010100011030005 chart for the CommonServices product, under the orca-redis and eventing-client services. This is an example of a command with multiple comma-separated `--filter` options.
`ccm report --filter "-s orca-redis, -o orca-streams-rest-7599c79656-2pfgv"`	Collects logs for all containers from the orca-redis service and in the orca-streams-rest-7599c79656-2pfgv pod. This is another example of a single command with multiple comma-separated `--filter` options.

Skip option

By default, `ccm report` collects log files on cluster infrastructure and application logs by running applications within the cluster. Use the `--skip` option to skip one or more log areas. When using `--skip`, include one or more of the following options within double quotes:

- `ccm`: To skip Cluster Control Manager logs.
- `infra`: To skip logs for nodes in the cluster.
- `app`: To skip application logs.

If you include two of these options, add a space between them.

For example, if you only want Cluster Control Manager logs, you can skip cluster infrastructure and application logs by running the following command:

```
ccm report --skip "infra app"
```

File integrity logs

When file integrity validation is enabled, run `ccm report --fileIntegrity` to collect AIDE logs.

Related links

Obtaining .tar log files on page 202
Viewing application logs in an Elasticsearch server on page 202
Collecting file integrity logs on page 204

ccm report log file operations

Obtaining .tar log files

About this task

Use this procedure to obtain `.tar` files for the `ccm report --remote-server` command option. For more information, see Remote server option on page 198.

Procedure

1. Log in to the SFTP server.
2. Run one of the following commands:
   - To access all logs, run `cd <directory path>/`
   - To access application and target logs, run `cd <directory path>/applicationLogs`

   For example, if you run `ccm report --remote-server "10.10.5.7 -u bruce -d /home/bruce/system1/"`, then `<directory path>` is `/home/bruce/system1/`.
3. Run `tar -czvf <file name>.tgz*` to navigate to a specific log file.

Viewing application logs in an Elasticsearch server

About this task

Use this procedure to upload logs to another Elasticsearch server for off-site viewing.

The `es_upload.sh` script provides a template to convert log files into a format, which can be inserted into Elasticsearch. The script uploads the files to the server. You must modify the script based on the URL, credentials, and other security considerations for the Elasticsearch server.

The script accepts the URL to the Elasticsearch server and the index used in the Elasticsearch server. The script leaves behind the converted files for debugging purposes.

Procedure

1. Collect logs using `ccm report` without the `--target` option.

2. Create an `es_upload.sh` script.

   See es_upload.sh script example on page 203 for an example of the script.

3. Run `chmod 755 es_upload.sh` to make the script executable.
4. Navigate to the `applicationLogs/` directory.
5. Run the script.

   For example: `es_upload.sh flex188-54.dr.avaya.com/elasticsearch example_company`

   In this example, an `example_company-<date>` index is created in Elasticsearch, containing the data. The date is in the format yyyy-mm-dd. For example, the name of the index might be `example_company-2020-08-17`.

6. **(Optional)** To view index logs for all dates in Elasticsearch, create an index without a date.

   For example: `example_company-*`

es_upload.sh script example

Create an `es_upload.sh` script, such as the following.

Note:

If the targeted Elasticsearch server is protected by a security mechanism, add the credentials or mechanism to this script.

```
#!/bin/bash
#set -x
function upload_to_es() {
    local _es_server=$1 # URL of Elasticsearch
    local _cust_name=$2 # E.g., example_company.  This will result in one or more
indicies of example_company-{date}

    echo "Uploading logs to $_es_server with index pre-fix $_cust_name..."

    _index_id=1
    _chunks=$(ls *log | sort)
    for _chunk in ${_chunks[@]}; do
        IFS=
        _es_index="$_cust_name"-$(echo $_chunk | awk -FT '{print $1}')
        echo ""
        echo "Converting ${_chunk} to ${_chunk}_es ..."
        while read -r entry; do
            echo "{\"index\":{\"_index\":\"$_es_index\", \"_id\":\"$_index_id\"}}";
echo $entry
            ((_index_id++))
        done < <(cat $_chunk | jq -c '.') > ${_chunk}_es

        echo "Uploading ${_chunk}_es to https://$_es_server/$_es_index ..."
        _rc=$(curl --insecure -s -o /dev/null --show-error --fail \
            -X PUT "https://$_es_server/$_es_index/_bulk?pretty" \
            -H 'Content-Type: application/x-ndjson' \
            --data-binary @${_chunk}_es)

        if [[ $_rc -ne 0 ]]; then
            echo "curl failed with ${_rc}...aborting upload"
        break
```

```
        fi
    done
}

#E.g., es_upload.sh flex188-54.dr.avaya.com/elasticsearch example_company
upload_to_es $1 $2
```

Collecting file integrity logs

About this task

When file integrity validation is enabled, you can use this procedure to collect and extract AIDE logs.

Before you begin

Run `clusterFileIntegrity enable` to enable file integrity validation.

Procedure

1. Log in to Cluster Control Manager using your customer account.
2. To collect logs, run the `ccm report --fileIntegrity` command.

   The following is an example of the output for this command:

   ```
   Retrieve CCM Report Begin
   Starting to collect Cluster Control Manager
   Collect Cluster Control Manager completed
   Starting to collect Cluster Node Manager logs
   Collect Cluster Node Manager completed
   Starting to collect Cluster Node logs
   Collect Cluster Node completed
   The report location is: /var/avaya/artifactCache/cspInformation/ccmReport-
   flex-220-202009101918.tgz. Once the report has been processed (e.g., copied
   to another location for off-line analysis), remove the report in /var/avaya/
   artifactCache/cspInformation directory.
   Retrieve CCM Report End
   ```

3. To extract the logs, run the `tar -xvzf <report_file_name>.tgz` command.

   For example, `tar -xvzf ccmReport-flex-220-202009101918.tgz`. The extracted logs are saved in the `/var/avaya/artifactCache/cspInformation/` directory.

   This command extracts individual component logs, such as the following:

Log type	Example of log file name
Cluster Control Manager AIDE logs.	`ccmLogFiles.tar.gz`
Cluster node AIDE logs for each node. The exact file name for each node varies.	`master_171e1956-f0b0-45a0-9249-1271695616b1-systemLogs.tar.gz` or `worker_993ce24c-8650-4483-af35-14b70093c944-systemLogs.tar.gz`

Certificate Manager commands

Command	Description
`ccmcertmgr --help`	Gives a list of all the sub-commands available for Certificate Manager operations.
`ccmcertmgr <-id|--identity-certs>`	Gives the list of all identity certificates.
`ccmcertmgr --service-identity-certs <serviceID>`	Gets identity certificate details of a given service. If you do not have the service ID, you can find it by running the `ccmcertmgr <-id|--identity-certs>` command.
`ccmcertmgr --renew-service-identity-cert <serviceID>`	Renews the identity certificate of a given service. Renewal is only allowed if a given certificate is issued by the Certificate Manager CA. Renewal retains all the attributes of a certificate. If you do not have the service ID, you can find it by running the `ccmcertmgr <-id|--identity-certs>` command.
`ccmcertmgr --replace-service-identity-cert <serviceID> <certInfoFile>`	Generates and replaces the identity certificate using the certificate attributes specified in the `certInfo` file. If you do not have the service ID, you can find it by running the `ccmcertmgr <-id|--identity-certs>` command. The following is a sample `certInfo` file: `{"commonName": "certmgmt-agent-certificate-document-idcert-cms", "keySize": "2048", "keyAlgorithm": "RSA", "subjectAltName": "dNSName=certmgmt-loadbalancer-service", "subject": "CN=certmgmt-loadbalancer-service, C=US, OU=MGMT, O=AVAYA"}`
`ccmcertmgr --generate-service-csr <serviceID> <certInfoFile>`	Generates a Certificate Signing Request (CSR) for a service. If you do not have the service ID, you can find it by running the `ccmcertmgr <-id|--identity-certs>` command. The following is a sample `certInfo` file: `{"commonName": "certmgmt-agent-certificate-document-idcert-cms", "keySize": "2048", "keyAlgorithm": "RSA", "subjectAltName": "dNSName=certmgmt-loadbalancer-service", "subject": "CN=certmgmt-loadbalancer-service, C=US, OU=MGMT, O=AVAYA"}`

Table continues…

Command	Description
`ccmcertmgr --import-pem-service-identity-cert <serviceID> <certFile>`	Imports the identity certificate of a service, which is created by signing the CSR generated by the `ccmcertmgr --generate-service-csr <serviceID> <certInfoFile>` command. If you do not have the service ID, you can find it by running the `ccmcertmgr <-id|--identity-certs>` command. `CertFilefromcsr.txt` should contain the PEM-formatted identity certificate chain signed by the external CA.
`ccmcertmgr --import-service-identity-cert <serviceID> <certFile>`	Imports an external CA signed identity certificate for a service. `base64 -w 0 weblmserver.p12` is an example of the command you can run to get the certificate file text. In this example command, `weblmserver` is the third-party certificate file in P12 format from the third-party CA. This command converts the certificate file from the P12 format to the Base64 format. If you do not have the service ID, you can find it by running the `ccmcertmgr <-id|--identity-certs>` command. The following is a sample certificate file: `{"base64PKCS12FileText": "<base64pkcs>", "storePassword": "password", "keyPassword": "password", "external": true}` In this sample, `"<base64pkcs>"` is the output of the `base64 -w 0 weblmserver.p12` example command described above.
`ccmcertmgr <-ts | --trusted-stores>`	Lists all trusted stores managed by Certificate Manager.
`ccmcertmgr <-ti | --trust-store-info> <serviceID>`	Gives trusted store details for a specified service ID. If you do not have the service ID, you can find it by running the `ccmcertmgr <-ts | --trusted-stores>` command.
`ccmcertmgr <-tc | --trusted-certs> <serviceID>`	Lists all CA certificates present in a trusted store. If you do not have the service ID, you can find it by running the `ccmcertmgr <-ts | --trusted-stores>` command.

Table continues…

Command	Description
`ccmcertmgr --add-trustcert <serviceID> <certFile>`	Adds a trusted certificate to the specified trust store. The certificate file (`<certFile>`) contains the certificate in the Base64 format. If you do not have the service ID, you can find it by running the `ccmcertmgr <-ts | --trusted-stores>` command. The certificate file contains a PEM-formatted certificate.
`ccmcertmgr --delete-trustcert <serviceID> <certificateId>`	Deletes a trusted certificate from the specified trusted store. If you do not have the service ID, you can find it by running the `ccmcertmgr <-ts | --trusted-stores>` command. You can find the certificate ID by running the `ccmcertmgr <-tc | --trusted-certs> <serviceID>` command.
`ccm release cert-manager third-party-certs`	Manages third-party certificates in the cluster. Various sub-command options are available for this command.

Related links

Certificate management command help outputs on page 223

Internal cluster certificate commands

The Common Services cluster contains internal Kubernetes (k8s) certificates that require manual rotation. These certificates are not rotated automatically or during a cluster upgrade. Rotate the certificates before they expire to prevent cluster failure.

Use these commands to rotate the certificates or to check when they expire.

Command	Description
`ccm rotate-cluster-certificates`	Rotate internal certificates before they expire. Run this command during a maintenance window. The rotation process can take up to two hours to complete.

Table continues…

Command	Description

`ccm release cert-manager getcerts -id --output-format short`	Check the expiration date for your internal cluster certificates. If any certificates in the list expire, your cluster could become unusable.
	The following is an example of the output for this command:
	`# ccm release cert-manager getcerts -id --output-format short`
	`| validFrom | validTo | Issuer | SUBJECT | SAN`
	`"Fri Jun 17 23:24:17 UTC 2022","Sat Jun 17 23:17:08 UTC 2023","O=Avaya, CN=Certificate Manager CA","CN=flex-81, C=US, OU=MGMT, O=AVAYA","dNSName=flex-81"`
	`"Sat Jun 18 00:00:43 UTC 2022","Sun Jun 18 00:00:42 UTC 2023","O=Avaya, CN=Certificate Manager CA","CN=rbac-service, C=US, OU=MGMT, O=AVAYA","uniformResourceIdentifier=spiffe://cluster.local/ns/default/sa/rbac-service, dNSName=rbac-service.default.svc"`
	`. . .`

Related links

Rotating cluster certificates manually on page 208

Rotating cluster certificates manually

Condition

The Common Services cluster contains internal Kubernetes (k8s) certificates that require manual rotation. These certificates are not rotated automatically or renewed during a cluster upgrade.

You must monitor certificate validity and rotate the certificates manually before they expire. If the certificates expire, the cluster becomes unstable or unusable and requires recovery or reinstallation.

- Warning alarm displays every 5 days if the certificates expire in 31 to 60 days.
- Major alarm displays every 5 days if the certificates expire in 15 to 30 days.
- Critical alarm displays every day if the certificates expire in 10 days or less.

Solution

1. **(Optional)** To check when the cluster certificates expire, run the `clusterNodeCertificateExpiryCheck` command.
2. To manually rotate the certificates before they expire, run `ccm rotate-cluster-certificates` during a maintenance window.

   This process can take up to two hours to complete.

   If the certificates have already expired, you must reinstall the cluster.

Common Services commands

Command	Description
`ccm release common-services update-config istio-ingressgateway <a.b.c.d>`	Changes the cluster ingress IP address, where a.b.c.d is the new IPv4 address. The cluster ingress IP address is the cluster IP assigned as the kube-keepalived-vip when the solution is deployed.
`ccm release common-services update-config cluster_fqdn <new fqdn>`	Changes the cluster FQDN. If you change the cluster ingress IP address, you must update the cluster FQDN or DNS entry to reflect the new ingress IP address.
`ccm release common-services exportPrometheusSnapshot`	Backs up solution metric (Prometheus) data.
`ccm release common-services alarmctl -l destinations`	• `-l destinations`: Lists all available destinations. • `-a destinations [destination-options]`: Creates a new alarm destination. The following are the available destination options: - `--address`: The IP address of the configured NMS destination. - `--type`: Either TRAP or KAFKA. - `--port`: The listening port of the configured NMS or KAFKA. - `--username`: The user name of the SNMPV3 user configured for the NMS. - `--password`: The password of the SNMPv3 user configured for the NMS. - `--auth`: The authentication protocol for the NMS. Either SHA or MD5. - `--priv`: The privacy protocol configured for the NMS. Either AES or DES. - `--url`: The KAFKA topic name. - `--authPwd`: The authentication password (optional). - `--privPwd`: The privcat password (optional). • `-d destinations --address <ip-address>`: To delete a destination that matches the provided IP address. • `-l alarmEvents [list-alarm-options]`: Lists the alarms on the cluster.
`ccm release common-services alarmctl -a destinations --address <ip-address> --type <TRAP|KAFKA> --port <port-num> [--username <string>] [--password <string>] [--auth <SHA>] [--priv <AES>] [--authPwd <string>] [--privPwd <string>] [--url <url-string>]`	
`ccm release common-services alarmctl -d destinations --address <ip-address>`	
`ccm release common-services alarmctl -l alarmEvents [--alarmId <number>] [--state <state>] [--notificationOid <oid>]`	

Table continues…

Command	Description
`ccm release common-services alarmctl --ClearAlarm --eventId <number> --productName <CSP> --productVersion <version> --serviceName <service-name>`	The available list-alarm option is `--alarmId <number>`. This is an optional parameter. • `--state <state>` (optional): Indicates the state of the alarm. The state can be ACTIVE or CLEARED. • `--notificationOid <oid>`(optional): Lists the alarm ID associated with the specified OID. • `--ClearAlarm [clear-alram-options]`: Clears the given raised alarm. clear-alarm options: - `--eventId`: The eventId for the alarm being cleared. - `--productName`: The product name for the alarm being cleared. - `--productVersion`: The product version for the alarm being cleared. - `--serviceName`: The service name for the alarm being cleared.

Eventing commands

Command	Description
`ccm release eventing-kafka kafka-util certificate-update-restart`	Restarts the eventing-kafka service pods. The restart causes the pods to use the newly imported certificates.

Related links

Eventing-kafka service command help outputs on page 226

Miscellaneous commands

Command	Description
`ccmHealthCheck`	Displays Cluster Control Manager information, NTP server information, and licensing information. This command invoked upon logging in to Cluster Control Manager. Same as **`ccmStatus`**.

Table continues…

Command	Description
`ccmNetSetup`	Changes the network-related configuration on Cluster Control Manager. For more information, see Updating the network configuration of Cluster Control Manager and cluster nodes on page 106. You can also run this command to configure outbound proxy settings. You will need to populate all Cluster Control Manager network attributes. For more information, see Using the ccmNetSetup command to configure proxy settings on page 119.
`ccmNetSetup --collect-proxy-only`	Running `ccmNetSetup` without any options prompts you to populate all Cluster Control Manager network attributes. Running `ccmNetSetup --collect-proxy-only` prompts you to update proxy-related network attributes only. For more information, see Using the ccmNetSetup command to configure proxy settings on page 119.
`checkInfra`	Verifies that your vCenter configuration is correct. You can run this command with additional options. For more information about the available options for this command, see checkInfra command help output on page 220.

Table continues…

Command	Description
`EASGManage`	Performs management tasks for the Avaya EASG authentication management system. • `-f --enableEASG` • `-f --disableEASG` • `--enableEASG` • `--disableEASG` • `--enable <user>` • `--disable <user>` • `--userStatus <user>` • `--listUsers` • `--printDisableWarning` • `--printEnableWarning` `-f`: Forces the enable or disable action to run without prompts. Acceptance of the terms is implied. `--enableEASG`: Enables EASG authentication. `--disableEASG`: Disables EASG authentication. `--enable`: Enables EASG authentication only for the user specified in <user>. If the main EASG enable/disable switch is disabled, no users will have access, no matter what this setting reflects for an individual user. `--disable`: Disables EASG authentication only for the user specified in <user>. `--userStatus`: Displays the enabled or disabled status for the user specified in <user>. The return code also indicates the user's status: 4 indicates that the user is not an EASG user, 3 indicates that the user is a DISABLED EASG user, and 2 indicates that the user is an ENABLED EASG user. On error, 1 is returned. `--listUsers`: Lists information about the available EASG users. `--printDisableWarning`: Displays the warning message for disabling EASG on the system. This is used when a web page needs to print the message to a user before actually doing the disable in non-interactive mode with the `-f` option. `--printEnableWarning`: Displays the warning message for enabling EASG on the system. This is

Table continues…

Command	Description
	used when a web page needs to print the message to a user before actually doing the enable in non-interactive mode with the `-f` option.

Table continues…

Command	Description
`EASGSiteCertManage`	Sets site certificates used by EASG. • `--add <pkcs7_file_path>` • `--add <pkcs7_file_path> --saf <SAF_string>` • `--list` • `--show <installed_pkcs7_name>` • `--delete all` • `--delete <installed_pkcs7_name>` `--add`: Adds the specified file to the site certificate repository for this machine. Without the --saf option the command will prompt the user to confirm the certificate details and enter a site authentication factor. `--saf`: Provides the site authentication factor in the `--add` option, which allows the command to run in a non-interactive mode. `--list`: Lists all the site certificates currently installed on this machine. `--show`: Displays detailed information about the specified site certificate. The names returned from the --list option are valid names for this option. `--delete`: Deletes the specified site certificate from the machine. The names returned from the --list option are valid names for this option. Additionally, 'all' can be specified to remove all site certificates. **Note:** • The `--saf` option expects an alphanumeric string of at least 10 characters, but no more than 20 characters. • The names returned by the `--list` option can be used for the `--show` and `--delete` options. These arguments should always end with '.p7b'. • Return codes are 0 for success and 1 for failure. Additional messages will be printed on `stdout` and `stderr`, as appropriate.
`EASGStatus`	Print the status of EASG.

Table continues…

Command	Description
`EASGProductCert`	`--lessThanDays <number_of_days>`: Determines if the certificate will expire within the number of days indicated by `<number_of_days>`. `--certInfo`: Displays information about the EASG Product Certificate.
`remoteDesktopSession`	Sets up a VNC Server and prints the connection URL to the console. The remote desktop will be terminated after one hour of inactivity.
`ccm swhistory`	Provides the history of software products and components installed on the system. This information is useful for troubleshooting purposes. For more information, see the detailed description for this command in ccm swhistory command on page 215.
`swversion`	Print the CCM Software version along with the version information for any deployed products.
`clusterFileIntegrity`	You can run `clusterFileIntegrity enable` to enable file integrity validation. To disable this feature, run `clusterFileIntegrity disable`. When file integrity validation is enabled, the server contains additional log files that consume up to 400 MB of additional disk space.
`updateLicenseService <license_service>`	Configure the network address to the licensing service. license_service parameter is System Manager FQDN/IPAddess or WebLM service FQDN/ IPAddress and port running inside the customer network. For example: smgr.example.com weblm.example.com:52233

Related links

Updating the network configuration of Cluster Control Manager and cluster nodes on page 106
Using the ccmNetSetup command to configure proxy settings on page 119
Enabling or disabling file integrity validation on page 120
ccm swhistory command on page 215

ccm swhistory command

Command description

The **`ccm swhistory`** command output displays the history of the results of the `ccm install`, `ccm upgrade`, `ccm delete`, `ccm uninstall`, `ccm restore`, and any patching commands that modify the system's software products. The output also includes versions of the software before and after modification.

When you execute a command that modifies the system's software products, a section is added to the output of the **`ccm swhistory`** command. Each section contains:

- An entry for the initialization (`INIT`) of the command on the affected products.
- A list of component entries, where each entry contains:
  - The initialization of the action performed on the component. The action can include install, upgrade, delete, uninstall, restore, and patch.
  - An entry for the success or failure of the action performed on the component.
- An entry for the success or failure of the command on the affected products.

Example

The following sample output contains the following sections:

- `CCM OVA DEPLOYMENT`
- `ccm install <spreadsheet-INSTALL.xlsx>`
- `ccm upgrade spec <spreadsheet-UPGRADE.xlsx>`

Each section is divided by a line of equal signs (======).

```
 Timestamp               User  Command            Product_Modification_Status    Product         product_from_version     product_to_version
-------------------------------------------------------------------------------------------------------------------------------------------
07/12/2020 22:12:45      -     CCM OVA DEPLOYMENT  INIT                          CommonServices  -                        -

Timestamp               Component  Component_Product  Component_Status  component_from_version   component_to_version
-------------------------------------------------------------------------------------------------------------------------
07/12/2020 22:13:15      CCM        Common Services    SUCCESS           -                        1.1.0.0.95748

Timestamp               User  Command             Product_Modification_Status    Product         product_from_version     product_to_version
-------------------------------------------------------------------------------------------------------------------------------------------
07/12/2020 22:13:15      -     CCM OVA DEPLOYMENT  SUCCESS                       CommonServices  -                        -

=========================================================================================================================================

Timestamp               User  Command                                                    Product_Modification_Status    Product
product_from_version     product_to_version
-------------------------------------------------------------------------------------------------------------------------------------------
--------------------
07/13/2020 16:57:15      cust  ccm install solution-deployment-spreadsheet-INSTALL.xslx  INIT                           CommonServices
-                        1.1.0.0.95001

Timestamp               Component          Component_Product  Component_Status  component_from_version   component_to_version
-------------------------------------------------------------------------------------------------------------------------------------------
------------
07/13/2020 16:57:15      Infra              CommonServices     INIT              -                        1.1.0.0.95028
07/13/2020 18:24:53      Infra              CommonServices     SUCCESS           -                        1.1.0.0.95028
07/13/2020 19:12:05      common-crds        CommonServices     INIT              -                        1.1.0.0.93020
07/13/2020 19:15:15      common-crds        CommonServices     SUCCESS           -                        1.1.0.0.93020
07/13/2020 19:15:15      cert-manager       CommonServices     INIT              -                        1.1.0.0.95382
07/13/2020 19:21:55      cert-manager       CommonServices     SUCCESS           -                        1.1.0.0.95382
07/13/2020 19:21:55      istio              CommonServices     INIT              -                        1.5.0.0.95184
07/13/2020 19:29:04      istio              CommonServices     SUCCESS           -                        1.5.0.0.95184
07/13/2020 19:29:04      postgres-operator  CommonServices     INIT              -                        1.1.0.0.95199
07/13/2020 19:31:05      postgres-operator  CommonServices     SUCCESS           -                        1.1.0.0.95199
07/13/2020 19:31:05      common-services    CommonServices     INIT              -                        1.1.0.0.95157
07/13/2020 19:35:04      common-services    CommonServices     SUCCESS           -                        1.1.0.0.95157
07/13/2020 19:35:05      eventing-kafka     CommonServices     INIT              -                        1.1.0.0.95555
07/13/2020 19:38:13      eventing-kafka     CommonServices     SUCCESS           -                        1.1.0.0.95555

Timestamp               User  Command                                                    Product_Modification_Status    Product
product_from_version     product_to_version
-------------------------------------------------------------------------------------------------------------------------------------------
--------------------
07/13/2020 19:38:13      cust  ccm install solution-deployment-spreadsheet-INSTALL.xslx  SUCCESS                        CommonServices
-                        1.1.0.0.95001

=========================================================================================================================================

Timestamp               User  Command                                                        Product_Modification_Status    Product
product_from_version     product_to_version
-------------------------------------------------------------------------------------------------------------------------------------------
--------------------------
07/14/2020 04:01:24      cust  ccm upgrade spec solution-deployment-spreadsheet-UPGRADE.xslx  INIT                           CommonServices
1.1.0.0.95001            1.1.0.1.99001

Timestamp               Component        Component_Product  Component_Status  component_from_version   component_to_version
```

```
-------------------------------------------------------------------------------------------------------------------------------------------------------------------
-------------------------------
07/14/2020 04:01:24      common-services  CommonServices     INIT             1.1.0.0.95157           1.1.0.1.99350
07/14/2020 04:04:06      common-services  CommonServices     SUCCESS          1.1.0.0.95157           1.1.0.1.99350
07/14/2020 04:04:07      eventing-kafka   CommonServices     INIT             1.1.0.0.95555           1.1.0.1.99777
07/14/2020 04:06:25      eventing-kafka   CommonServices     SUCCESS          1.1.0.0.95555           1.1.0.1.99777


Timestamp                User  Command                                                        Product_Modification_Status    Product
product_from_version     product_to_version
-------------------------------------------------------------------------------------------------------------------------------------------------------------------
-------------------------
07/14/2020 04:06:25      cust  ccm upgrade spec solution-deployment-spreadsheet-UPGRADE.xslx  SUCCESS                        CommonServices
1.1.0.0.95001            1.1.0.1.99001
```

pvcCleanup command

You can use the **pvcCleanup** command in a VMware vCenter deployment.

Command conditions

The **pvcCleanup** command finds PVCs that are in one of the following conditions, which prevent pods from launching.

- Kubelet cannot mount the PVC. The associated virtual machine disk is in the Locked state on the VMware datastore.
- Kubelet tries to create a PVC, but vCenter indicates that the virtual machine disk with the PVC name already exists.

Usage options

You can run this command as follows:

```
pvcCleanup --unlock-pvc | --delete-pvc | --recover-vmdk | --delete-failed-pvs | --
review | --help
```

The following table provides a description of each option:

Option	Description
`--unlock-pvc`	Unlocks all PVCs that vCenter sees as attached to a virtual machine, but the PVCs are not mounted on the virtual machine.
`--delete-pvc`	Deletes PVCs for which Kubernetes does not have a record, but that vCenter sees as present.
`--recover-vmdk`	Creates or recovers the VMDK. If your docker registry is lost, you can run `pvcCleanup --recover-vmdk` to create a new VMDK, but the VMDK will not have any docker images. Run the `ccm registry repopulate` command to repopulate the docker image registry with docker images.
`--delete-failed-pvs`	Deletes all PVs that are in a Failed state from the Kubernetes cluster.
`--review`	Lists PVCs that vCenter sees as attached to a virtual machine, but the PVCs are not mounted on the virtual machine.
`--help`	Displays help information.

ccm resizePVC command

The **resizePVC** command increases the size of a Kubernetes Persistent Volume Claim (PVC). This command creates a new PVC that is the size specified and copies the existing application data into the new PVC. Technical staff can resize the disk space allocated to a service after it is deployed while preserving the original data.

Command requirements and technical constraints

- You must know the names of the PVCs that are to be resized.
- Only File System type PVC volumes are supported.
- Resizing is only supported on VMware.
- Services using the specified PVCs are taken offline during resize operations.
- Do not run installation and upgrade tasks concurrently.

Disk space considerations

- The cluster storage provider must have adequate free space for each resized PVC.

  For example, when resizing 3 PVCs and the existing size is 25GB and the destination size is 40GB, the storage provider requires 120GB of free space. After the command completes, an additional 45GB is consumed.
- If multiple PVCs are specified for resize in a single command, they are executed concurrently. If insufficient storage is available to execute all commands concurrently, resize each PVC individually.

  In the example, 120GB of free space is required at the time of the command. If PVCs are resized individually, the required free space is 70GB.

Usage

After a service is deployed, you can run the following command to resize the disk space allocated to the service: `ccm resizePVC <NewSize(Mi|Gi)> <pvc-name-1> <pvc-name-2> <pvc-name-n>`

Example 1: `ccm resizePVC 30Gi datadir-0-eventing-kafka-cp-kafka-0 datadir-0-eventing-kafka-cp-kafka-1 datadir-0-eventing-kafka-cp-kafka-2`

Example 2: `ccm resizePVC 5Gi alarming-db-service-pgo-repo1`

For details about resizing a PVC using a spreadsheet, see Resize PVC using a spreadsheet on page 113.

Related links

Resizing a PVC on page 218

Resizing a PVC

About this task

You can use this procedure to resize the disk space allocated to a service after it is deployed or during the initial installation. This procedure preserves the original data of a deployed service.

The resize command stops any running containers from using the PVCs. This command creates destination PVCs and copies the existing PVC data to the new PVC. The resize command

deletes the original PVC and re-creates it, using the newly-created copy. After the re-creation, any services using the PVC are restarted.

If you have sufficient free space, specify multiple PVCs. This results in faster processing as the commands to scale down and scale up the deployments are only executed once.

Important:

- Resizing is supported only on VMware.
- Only Filesystem type PVC volumes are supported.
- Do not run installation or upgrade tasks concurrently.

Before you begin

Ensure you know the names of the PVCs that are to be restored.

To retrieve the list of PVCs, run the `kubectl get pvc -A` command.

Note:

Services using the specified PVCs are taken offline during resizing operations.

Procedure

1. Log in to Cluster Control Manager as a user customer.

   You require root privilege to run the `ccm resizePVC` command.

2. To retrieve the list of PVCs, run the `kubectl get pvc -A` command.

3. Run the `ccm resizePVC <NewSize(Mi|Gi)> <pvc-name-1> <pvc-name-2> <pvc-name-n>` command to resize the PVC.

   This command resizes the disk space allocated to a service after the service is deployed.

   For example: `ccm resizePVC 40Gi datadir-0-eventing-kafka-cp-kafka-0 datadir-0-eventing-kafka-cp-kafka-1 datadir-0-eventing-kafka-cp-kafka-2`

4. Run `ccm smoke-test` to gauge when the service is recovered from the requested procedure.

   If the smoke-test continuously fails for more than an hour, contact Avaya Support for assistance.

Related links

ccm resizePVC command on page 218

clusterVPN commands

Cluster VPN enables encryption between the cluster nodes for software-defined storage (SDS). Cluster VPN can only be set up on a cluster containing two or more nodes.

Cluster VPN has two states, either enabled or disabled. Cluster VPN is enabled when the VPN is set up and started or restarted. Cluster VPN is disabled when the VPN service is stopped or disabled.

During an upgrade, IPSec VPN is not affected. If migrating, IPSec VPN configures the new nodes. If adding a node, IPSec is reconfigured on all the nodes and adds the new node to the tunnel configuration.

The following table lists the **clusterVPN** command options.

Command option	Description
clusterVPN setup	Installs and configures the cluster VPN.
clusterVPN start	Starts the cluster VPN on all cluster nodes.
clusterVPN restart	Restarts the cluster VPN on all cluster nodes.
clusterVPN stop	Stops the cluster VPN on all cluster nodes
clusterVPN disable	Stops and disables the cluster VPN on all cluster nodes and deletes the configuration.
clusterVPN status	Provides the status of cluster VPN from all the cluster nodes.

Related links

Managing cluster VPN on page 112

Command help output examples

The following sections provide command help output examples for reference purposes. These output examples might be useful while you configure certificates and services, or perform maintenance and troubleshooting tasks.

checkInfra command help output

You can use the `checkInfra` command to verify your cluster runtime configuration and the optional vCenter DRS rule when credentials are provided.

```
Usage: checkInfra [-C] [ -n ] [ -t ] [ -ha [-vu -va -vd -vc] ] [ -s ] [ -h ] [-S] [-Sd]
[-p [-nl -nb -io] [-cn] ] ]
 -C|--connectivity : Perform a connectivity check between CCM and cluster nodes
 -cr|--credentials : Validate the vCenter user account in the CSP cluster
 -n|--ntpservers : Check to see if configured NTP servers are reachable and responsive
 -t|--timezone : Check to see if configured timezone is set on CCM and cluster nodes
 -l|--license : Check to see if configured license server is reachable. Also shows the
current license Status
 -d|--dns : Check to see if configured servers are reachable and responsive.
 -L|--ext-logserver : Check to see if configured external log server is reachable.
 -a|--archive : Check to see if configured archive server is reachable.
 -S|--storage : Show overall status of storage services of the cluster.
 -Sd|--storage-details : Show current status of each storage component.
 -s|--show : Show current CSP Cluster configuration. If this option is given, the rest
of the options are ignored
 -ha|--high-availability : Check to see if node anti-afinity is set.
 -vu|--vcenter-user : vCenter user ID created for CSP use
 -va|--vcenter-address : vCenter IP address or FQDN
 -vd|--vcenter-datacenter : vCenter datacenter name
```

```
 -vc|--vcenter-cluster : vCenter cluster name within the given datacenter
 -cc|--cluster-config : fullpatch to cluster-config.yaml file
 -p|--performance : Check to see if environment's IOPs, network laterncy, and network
bandwidth meet the minimum requirement
 -nl|--network-laterncy : Sub-option of '-p|--performance' for checking network laterncy
 -nb|--network-bandwidth : Sub-option of '-p|--performance' for checking network
bandwidth
 -io|--iops : Sub-option of '-p|--performance' for checking iops
 -cn|--cluster-nodes : Comma-seperated cluster node IPs or FQDNs
 -h|--help : Print this help message. If this option is given, the rest of the options
are ignored
Examples:
checkInfra -C
 checkInfra -C -n
 checkInfra
 Note
 If no specific options are given, it assumes all options are set except for -s and -h
```

Example uses of checkInfra

- Validate the stored HA audit credentials:

```
$ checkInfra -cr

credentialsCheck:
  vCenter credentials validation for user=[nbo-
local@vsphere.local]: Success
```

- Validate the HA DRS rule:

```
$ checkInfra -va 192.0.2.215 -ha -vu readOnly@vsphere.local -vd
testDatacenter -vc testCluster
  vCenter credentials validation for user=[readOnly@vsphere.local]:
Success
  Check if cluster has 3 masters: Success
  Check if all cluster VMs all have the same DRS anti-affinity
rule: Success
    DRS anti-afinity rule set for cluster nodes:
          ClusterNode245-1.3.0.0.144001: 240-rule
          ClusterNode244-1.3.0.0.144001: 240-rule
          ClusterNode246-1.3.0.0.144001: 240-rule
```

- Default when no options are specified:

```
$ checkInfra

connectivityCheck:
  Cluster Control Manager ---> ClusterNodes: Success
  Cluster Control Manager ---> Cluster K8APIServer: Success

ntpServersCheck:
  ClusterControlManager   ---> ntpServer[198.152.8.11]: Success
  ClusterControlManager   --->
ntpServer[timeServer.dr.example.com]: Success
timezoneCheck:
  ClusterNodes timezone [America/Denver]: Success

licenseCheck:
```

```
  ClusterNodes:           --->
licenseServer[licenseServer.dr.example.com]: Success

  Cluster License status: NORMAL

archiveCheck:
  Archive server is not configured
  Run 'ccm archive config help ' to find out how to configure
archive server

externalLogServerCheck:
  external log server is not configured

dnsCheck:
  ClusterControlManager dns server: Success
  ClusterNodes dns server configuration: Success

storageCheck:
Cluster Storage overall status: Success
highAvailabilityCheck:
  vCenter credentials validation for user=[readOnly@vsphere.local]:
Success
  Check if cluster has 3 masters: Success
  Check if all cluster VMs all have the same DRS anti-affinity
rule: Success
    DRS anti-afinity rule set for cluster nodes:
          ClusterNode245-1.3.0.0.144001: 240-rule
          ClusterNode244-1.3.0.0.144001: 240-rule
          ClusterNode246-1.3.0.0.144001: 240-rule
```

- Performance check where the cluster environment meets the requirements.

```
$ checkInfra -p -cn 192.0.2.200,192.0.2.201,192.0.2.202
performanceCheck:
  checkNodesReachable:
      192.0.2.200: Success
      192.0.2.201: Success
      192.0.2.202: Success
  networkLatencyCheck:
      Max RTT 192.0.2.200 to 192.0.2.201: 0.323 ms
      Max RTT 192.0.2.200 to 192.0.2.202: 1.190 ms
      Max RTT 192.0.2.201 to 192.0.2.202: 0.357 ms
    Max rtt is within the max allowed network latency of 25ms.
  networkBandwidthCheck:
      Network bandwidth 192.0.2.201 to 192.0.2.200: 27.1 Gbits/sec
      Network bandwidth 192.0.2.202 to 192.0.2.200: 26.0 Gbits/sec
      Network bandwidth 192.0.2.202 to 192.0.2.201: 17.4 Gbits/sec
    Network bandwidth meet the minimum required network bandwidth
of 1 Gbits/sec.
  iopsCheck:
    Checking node with SDS Disk[192.0.2.200]:
      This check will take at least 4 minutes....
          Testing Read IOPS...
          Testing Write IOPS...
```

```
          Random Read/Write IOPS: 41.1k/19.3k
    Checking node without SDS Disk[192.0.2.202]:
      This check will take at least 4 minutes....
          Testing Read IOPS...
          Testing Write IOPS...
          Random Read/Write IOPS: 38.2k/17.9k
    Random Read/Write IOPS meet the minimum required IOPS of 500.
Performance check completed.
```

Certificate management command help outputs

ccm release cert-manager third-party-certs output

```
third-party-certs:

This command can be used to manage third party certificates in the cluster.

Sub-commands:
  --generate-service-csr --list-file <filename> --output-dir <dirname>
     list-file: (Atypical optional parameter) Location of the file having a customer defined list of
                service Ids. Service Ids should be separated by newlines.
                If the param is absent default lists, as defined by CCM and currently deployed
                services, are used.
     output-dir: The directory in which the command generates the CSRs.
                 It generates files with the name <serviceId>.csr.

   --add-trustcert --list-file <filename> --ca-cert-file <filename>
      list-file: (Atypical optional parameter) Location of the file containing a list of service Ids
                 to add the trusted certificate. Service Ids should be separated by newlines.
                 In case the param is absent, the default list is used.
      ca-cert-file: The file containing the CA certificate.

   --import-identity-cert-pem --id-cert-dir <dirname>
       id-cert-dir: Directory location where certificate files with the name <serviceId.pem>
                    are placed.

   --add-certs --list-file <filename> --ca-cert-file <filename> --id-cert-dir <dirname>
      list-file: (Atypical optional parameter) Location of the file containing a list of service Ids
                  to add the trusted certificate. Service Ids should be separated by newlines.
                  In case the param is absent, the default list is used.
      ca-cert-file: The file containing the CA certificate.
      id-cert-dir: Directory location where certificate files with the name <serviceId.pem>
                   are placed.

   -h|--help: Print this help message
```

Alarming command help outputs

ccm release common-services alarmctl output

```
Usage:

For help regarding the use of alarmctl tool:
    ccm release common-services alarmctl -h

List all available destinations:
    ccm release common-services alarmctl -l destinations

Create a destination entry:
    ccm release common-services alarmctl -a destinations --address <1.1.1.1> --type <TRAP> --port <162> --username
<admin> --password <****> --auth <SHA> --priv <AES> --authPwd <***> --privPwd <***>
In case of destination type = TRAP, no need to supply --url
    ccm release common-services alarmctl -a destinations --address <1.1.1.1> --type <KAFKA> --port <162> --username
<admin> --password <****> --auth <SHA> --priv <AES> --url <url> --authPwd <***> --privPwd <***>

Note: Authentication protocol(authPwd) and Privacy Protocol(privPwd) are optional.

Delete a destination by supplying only address of the destination:
    ccm release common-services alarmctl -d destinations --address <1.1.1.1>
If the destination with that address is not present, then such message will be displayed -> Destination with given
address not found

ErrorEvents:
List All Error Events:
```

```
    ccm release common-services alarmctl -l errorEvents
List ErrorEvents
a. based on event Id:
    ccm release common-services alarmctl -l errorEvents --eventId <10010>
b. based on product name:
    ccm release common-services alarmctl -l errorEvents --productName <CommonServices>
c. based on product version:
    ccm release common-services alarmctl -l errorEvents --productVersion <1.2>
d. based on any of above combinations:
    ccm release common-services alarmctl -l errorEvents --productVersion <1.2> --productName <CommonServices>
    ccm release common-services alarmctl -l errorEvents --productVersion <1.0.0.0> --productName <CommonServices> --
eventId <10012>

View the raised Alarms:
List all alarms:
    ccm release common-services alarmctl -l alarmEvents
List alarms:
a. Based on alarm's event Id:
    ccm release common-services alarmctl -l alarmEvents --eventId <10012>
b. Based on Alarm state. whether raised or cleared:
    ccm release common-services alarmctl -l alarmEvents --eventId <10010> --state RAISED
c. Based on Alarm Notification OID, event Id and state:
    ccm release common-services alarmctl -l alarmEvents --eventId <10012> --state RAISED --notificationOid
<.1.3.6.1.4.1.6889.2.121.0.10012>
d. Based on Product name:
    ccm release common-services alarmctl -l alarmEvents --productName <CommonServicesPlatform>
e. Based on Product Name and Product Version:
    ccm release common-services alarmctl -l alarmEvents --productName <CommonServicesPlatform> --productVersion
<1.0.0.0>
f. Based on Product Name, Product Version and Service name:
    ccm release common-services alarmctl -l alarmEvents --productName <CommonServicesPlatform> --productVersion
<1.0.0.0> --serviceName <CMonitor>
g. Based on Alarm's Severity:
    ccm release common-services alarmctl -l alarmEvents --severity <CRITICAL>
h. Based on all criterias:
    ccm release common-services alarmctl -l alarmEvents --eventId <10012> --productName
<CommonServices> --productVersion <1.1.0.0.970011> --state <RAISED> --serviceName <CMonitor> --notificationOid
<.1.3.6.1.4.1.6889.2.121.0.10012> --severity <CRITICAL>
Send Clear ErrorEvent to clear the raised Alarm:
This will work when the event Id in raised alarm is used as event id for clear error event.
    ccm release common-services alarmctl --ClearAlarm --eventId <52703> --productName <CommonServicesPlatform> --
productVersion <1.0.0.0> --serviceName <CMonitor>

Note: Use --tabulate argument to view the alarmEvents in tabular format.
        ccm release common-services alarmctl -l alarmEvents --tabular

Test the Alarm:
This is used to test whether an alarm is getting raised or not.
    ccm release common-services alarmctl --TestAlarm --productName <CommonServices> --productVersion <1.0>

Generate Mib:
The command to generate MIB file for product name and version. It will generate MIB file with product name provided as
argument on '/tmp' location.
    ccm release common-services alarmctl --generateMib --productName <Product Name> --productVersion <Product version>

Clear Older records:
If the number of Error Events/ Alarm Events is greater than some threshold(currently 100), this command will delete the
older error events/alarm events.
Error Events:
    ccm release common-services alarmctl --clean errorEvents  --productName <ProductName> --productVersion
<ProductVersion> --serviceName <ServiceName>
Alarm Events:
    ccm release common-services alarmctl --clean alarmEvents --productName <ProductName> --productVersion
<ProductVersion> --serviceName <ServiceName>

Get Help associated with alarm :
    Based on notification OID:
        ccm release common-services alarmctl --ah --productName <productName> --productVersion <productVersion> --
notificationOid <nOID>
    Based on serviceName and eventId:
        ccm release common-services alarmctl --ah --productName <productName> --productVersion <productVersion> --
serviceName <servicename> --eventId <eventId>
Note: If serviceName contains space, please provide it in double quotes

Description : Utility for Alarming service operations

optional arguments:
  -h, --help            show this help message and exit
  -l  , --list          list the resource eg. destinations|errorEvents
  -d  , --delete        delete the resource eg. destinations|errorEvents
  -a  , --add           add the resource eg. destinations
  --ah                  returns help associated with the alarm event
  -c C                  ignore
  --ClearAlarm          clears raised alarm. Required: eventId, productName,
                        productVersion, serviceName of ErrorEvent
  --TestAlarm           Raises alarm. Required: productName,productVersion of
                        ErrorEvent
  --clean               Cleans error events if count exceeds predefined value.
                        Required: Event(errorEvents or
```

```
                      alarmEvents),productName,productVersion of ErrorEvent
  --threshold         update the threshold of AlarmDefinition
  --generateMib       generate Mib for AlarmDefinition

Alarm Event Options:
  --alarmId ALARMID   Event ID of Alarm Event
  --state STATE       state of the Alarm raised (CLEARED|ACTIVE)
  --notificationOid NOTIFICATIONOID
                      notificationOid of the Alarm raised
  --severity SEVERITY severity of the Alarm raised
  --tabular

Error Event Options:
  --eventId EVENTID   Event ID of Error Event
  --productName PRODUCTNAME
                      Product Name
  --productVersion PRODUCTVERSION
                      Product Version
  --serviceName SERVICENAME
                      Service Name

Destination Options:
  --address ADDRESS   ip address of the destination
  --type TYPE         type of the destination (TRAP|KAFKA)
  --port PORT         port number of the destination
  --username USERNAME username of the user
  --password PASSWORD password of the user
  --auth AUTH         authenticationProtocol (SHA|MD5)
  --priv PRIV         privacyProtocol (AES|DES)
  --url URL           url: to be provided only in case of KAFKA
  --authPwd AUTHPWD   Authentication password.Optional Parameter
  --privPwd PRIVPWD   Privacy password.Optional Parameter

generate Mib Options:
  --productName PRODUCTNAME
                      Product Name
  --productVersion PRODUCTVERSION
                      Product Version

Clear Older Records Options:
  --productName PRODUCTNAME
                      Product Name
  --productVersion PRODUCTVERSION
                      Product Version
  --serviceName SERVICENAME
                      Service Name
```

Cmonitor service command help outputs

You can use the Cmonitor service to monitor your cluster.

ccm release common-services cmonctl output

```
Usage: cmonctl <options> [ <resource> [ <healthObject> ] ]
 options:
     -p          : Print current state of all cmonitor resources and healthObject to the logs
     -D  <DEBUG> : Debug flag in decimal or hexadecimal number
     -D1 <DEBUG> : Second level debug flag in decimal or hexadecimal nunmber
     -lr         : List all cmonitor resources and their healthObjects or healthObjects associated
with a given resource
                 : The cmonitor service is configured with "servicemonitor" and "network" resources
     -le         : List all errors or errors associated with a specific resource or an healthObject
associated with a resource
     -ce         : Clear all errors or errors associated with a specific resource or an healthObject
associated with a resource
     -h|--help   : Print this help message and exit

Current "servicemonitor" resource healthObjects are:
    ccmMonitor          : Checks the health of Cluster Control Manager
    bosh                : Checks the health of bosh director
    prometheus          : Checks the health of prometheus service running in the cluster
    ingress             : Checks to see if external services can access services running inside the
cluster via the ingress gateway
    alerting            : Checks the health of alerting service running in the
cluster

The "network" resource healthObjects are IP address cmonitor service configured to ping periodically.
```

ccm release common-services updateIp output

```
Usage: updateIp
     Cmonitor service periodically checks if CCM is reachable. It is configured with CCM IP address
     during cluster deployment. This command updates the cmonitor service with new CCM IP Address
     if CCM IP address has changed. If the CCM IP address has not changed, no action is taken.
```

Eventing-kafka service command help outputs

ccm release eventing-kafka kafka-util output

```
Usage: kafka-util <topics|consumerGroups|replication-config|delete-replication-config|replication-connector-tasks>
  This command can be used to manage topics, consumerGroups, replication configuration and
  replication connector tasks. It supports the following sub commands:

  topics                                : Manage topics in kafka
     options:
      --list                            : List all the topics currently in kafka
      --describe [--topic <topic_names>] : Describe all or a specific topic
      --delete <topic_name>             : Delete a topic

  consumerGroups                        : Manage consumerGroups in kafka
     options:
      --list                            : List all consumer groups
      --describe   (--group <group_name> | --all-groups) [ --members| --state ]  : Describe all or a specific consumer
group given by the --group option.
      --reset-offsets  (--group <group_name> | --all-groups)  (--topic <topic_name> | --all-topics) (--by-duration
'PnDTnHnMnS'|--to-offset <offset-number> |--to-datetime 'YYYY-MM-DDTHH:mm:SS.sss'|--to-earliest|--to-current|--shift-by
<number-of-offsets> |--to-latest ) (--execute|--dry-run) : Reset offsets on consumer groups

     sub-options:
      --group <group-name>              : To act on a specific group. Used only with --reset-offsets and --describe
options
      --all-groups                      : To act on all groups. Used only with --reset-offsets and --describe options
      --members                         : Describe members of the group or all groups. Used only with --describe option
      --state                           : Describe state of the group or all groups. Used only with --describe option
      --all-topics                      : to act on all topics. Used only with --reset-offsets option
      --topic <topic_name>              : to act on a specific topic. Used only with  --reset-offsets option

  replication-config
     options: [ all options are required for this subcommand ]
      -ldc|--local-datacenter <datacenter_name> : local data center name. No spaces allowed in the name
      -lbs|--local-bootstrap-server <bootstrap_server> : local bootstrap server
      -lc |--local-cluster <FQDN/IP>    : local cluster FQDN or IP address
      -rdc|--remote-datacenter <datacenter_name> : remote data center name. No spaces allowed in the name
      -rbs|--remote-bootstrap-server <bootstrap_server> : remote bootstrap server
      -rc |--remote-cluster <FQDN/IP>   : remote cluster FQDN or IP address

  delete-replication-config             : sub commmand to delete replication configuration
     options:
      -lc|--local-cluster <FQDN/IP>     : local cluster FQDN or IP address (required option)

 replication-connector-tasks            : sub command to view and set max connector tasks
     options:
      --view                            : See the connector tasks details
      --set-max-tasks <max-tasks-value> : Sets the maximum tasks value for replication connector

 certificate-update-restart             : this sub command restarts kafka after certificate updates

 Example  commands
  kafka-util topics --list
  kafka-util topics --describe --topic my_topic
  kafka-util topics --delete my_topic

  kafka-util consumerGroups --list
  kafka-util consumerGroups --describe --group my_group --state
  kafka-util consumerGroups --reset-offsets --group my_group --topic my_topic --to-earliest --dry-run

  kafka-util replication-config -ldc my_localDC -lbs my_lbs_server -lc 10.129.1.20 -rdc my_remoteDC -rbs my_rbs_server
-rc 10.130.1.20
  kafka-util delete-replication-config -lc 10.129.1.20

  kafka-util replication-connector-tasks --view
  kafka-util replication-connector-tasks --set-max-tasks 100
```

Logging command help outputs

ccm release common-services getlogs output

```
Usage: getlogs [OPTIONS]
         Downloads logs from the log manager running in the cluster. The command requires a container
```

```
name,
         a time range or both to limit the number of log entries to be downloaded

   OPTIONS:
        -c <containerName>      : Get logs for a container. Can request logs for multiple containers
with multiple -c options.
                                  The container names in a pod can be found by running:
                                  "kubectl get pod <yourPodID> -o jsonpath={.spec.containers[*].name}"
        -p <numHours>           : Request logs for past number of hours
        -f <timestamp>          : From timestamp (iso-8601 formatted from timestamp e.g.
2018-05-04T20:40:03)
        -t <timestamp>          : To timestamp (iso-8601 formatted to timestamp e.g.
2018-05-04T21:40:03)
        -l <level>              : Filter on the given log level. Valid options TRACE/WARN/WARNING/
ERROR/DEBUG/FINE/FINER/FINEST/AUDIT-SEC
        -F <outputFormat>       : Choose the log output format: json or short. defaults to json
        -o <outFileName>        : Name of the output tar file. Defaults to log-`date`.tgz. If the file
name
                                  does not include path, the output file will be placed in /var/avaya/
artifactCache directory
        -D <outputDirectory>    : Name of the directory where the downloaded logs (chunks of 10k
records) will be placed.
                                  If this directory is specified the downloaded logs will not be
tarred, unless -o option is also specified.
                                  If the output file name is specified with -o option and the name
contains path to the destination file,
                                  the path of the file name is ignored and the directory specified via
this option is used.
                                  The outputDirectoy specified in this option should not exist.
       -h|--help                : Print this help message
```

Chapter 7: Troubleshooting Avaya Analytics™ web issues

Unable to access the Avaya Analytics™ webpage

Condition

When you enter `https://<hostname>/AvayaAnalytics/servlet/AnalyticsWeb` in the address bar of a selected browser, the browser displays the following error:

```
There are no projects connected to this web server. To configure
projects, go to the Web Administrator.
```

Cause

The server definition on the Avaya Analytics™ Intelligence Server might be corrupted.

Solution

Note:

Only an administrator with the required permissions can complete these steps.

1. Log in to Avaya Analytics™ WebAdmin server by using the following URL:

   `https://<hostname>/AvayaAnalytics/servlet/AnalyticsWebAdmin`.

   The Administrator page displays the connected servers information.
2. To disconnect from the server, click **Disconnect** in the Action column.
3. To reconnect to the server, click **Connect** in the Action column.
4. Log in to Avaya Analytics™ again with the correct credentials.

Unable to log in to the Avaya Analytics™ Historical Reporting

Condition

Logging in to the Avaya Analytics™ webpage failed.

Cause

Invalid username and password.

Solution

Check you user credentials and log in again with the correct values.

Recover / Reset default Historical Reporting Administrator account password

Condition

Use this procedure if the default Historical Reporting Administrator account password details have been lost or misplaced.

Cause

Default Historical Reporting Administrator account password details have been lost or misplaced.

Solution

1. A customer should make support teams aware that their Historical Reporting Administrator password has been lost or misplaced.
2. Support will need to raise a ticket with Historical Reporting development team to reset Historical Reporting Administrator password.

Unable to log in to Avaya Analytics™ using LDAP

Condition

Users with the supervisor role cannot log in to Avaya Analytics™ using LDAP.

Cause

This issue occurs when you store supervisors and supervisor groups above the search root path.

Solution

1. Create supervisors and supervisor groups on or below the search root path.
2. Store supervisors and supervisor groups at a location that is one or more levels below the search root folder.

Unable to run an Avaya Analytics™ historical report with selected interval range

Condition

Attempting to run an Avaya Analytics™ historical report fails. An error message displays in the Avaya Analytics™ UI, which states that the job execution limit exceeded.

Cause

The interval range, for example the date range, selected for the report is too large.

Solution

1. Enter a more restrictive date range.
2. Run the report again with a shorter date range.

Routing Services no longer appear on the Routing Service Group

Condition

If you delete Routing Services from a Routing Service Group, they no longer appear when you report on the Group. Avaya Analytics™ behaves this way whether you run the report for historical or current periods.

Cause

The Routing Service Group reporting operation is different from the Agent Group reporting.

If Supervisors can access to a Routing Service Group, they can run reports for historical and current periods on this Group. If Supervisors do not have access to a Routing Service Group, they cannot report on this Group for historical or for current periods.

Voice channel activity does not increase on the Routing Service Summary report

Condition

Agents are handling contacts using chat, email, SMS and voice calls. When a supervisor looks at the Routing Service Summary report (real-time and historic) the values for all types of channel increase appropriately over time except for voice contacts.

Cause

This occurs when Routing Service Groups are enabled but the service name is missing the *Location.Inhouse* attribute.

Deleted agent groups data appears in historical reports that supervisors generate

Condition

When supervisors run Agent Group reports, they can view deleted or historical data about the Agent Groups that they were previously assigned.

Cause

This is intentional so that supervisors who were previously assigned an Agent Group can still run reports if they want to view historical information about the Agent Group. This information includes deleted records too.

If a supervisor sets the time filter as `Yesterday`, then the report displays only those agents that were configured for that Agent Group the previous day, not the present set of agents. Supervisors can only view historical data from the time that they were added to the Group.

MSTR-SRV pods stuck in PodInitializing status

Condition

The MSTR-SRV pod remains in a PodInitializing status and fails to get to a running status.

Cause

The mstr-srv pod startup failed and stopped.

Solution

1. Check the **`scripts.log`** to see the exception that has caused the startup to fail.
2. Resolve the issue and restart the mstr pod using these instructions:

   ```
   k scale --replicas=0 deployment mstr-srv -n mstr
   ```

   Wait for the pod to get fully terminated. Use

   ```
   k get pods -n mstr |grep mstr-srv
   ```

   to verify the pod has terminated, then confirm the volume attachment to the pv has been released. This will not happen until after the pod fully terminates.

   ```
   k get pv |grep mstr
   ```

   Take a note of the mstr-srv pv name.

   ```
   k -request-timeout=15s get volumeattachement |grep <mstr-srv pv name>
   ```

   If nothing is returned, continue to next step. Otherwise, escalate to support for further analysis.

   ```
   k scale --replicas=1 deployment mstr-srv -n mstr
   ```

   Note:

   A common cause for this issue is incorrect credentials for the mstr administrator.

Copying analyticsdb-node secrets into mstr namespace

About this task

Use this procedure to copy analyticsdb-node secrets into the mstr namespace.

Procedure

1. Confirm mstruser and postgres secrets are available in the default namespace.

   ```
   kubectl get secrets | grep 'analyticsdb-pguser-mstruser\|
   analyticsdb-pguser-postgres'
   ```

2. Copy secrets to tmp file

   ```
   kubectl get secrets -n default analyticsdb-pguser-postgres -o
   json | jq '.metadata.namespace = "mstr"' > /tmp/analyticsdb-pguser-
   postgres.yaml
   ```

   ```
   kubectl get secrets -n default analyticsdb-pguser-mstruser -o
   json | jq '.metadata.namespace = "mstr"' > /tmp/analyticsdb-pguser-
   mstruser.yaml
   ```

3. Edit the yaml files.

   ```
   vi /tmp/analyticsdb-pguser-postgres.yaml
   ```

4. Delete lines below "namespace": "mstr", to "resourceVersion": "Id Number", and save.
5. Apply the secrets into the mstr namespace

   ```
   kubectl apply -n mstr -f /tmp/analyticsdb-pguser-postgres.yaml
   ```

   ```
   kubectl apply -n mstr -f /tmp/analyticsdb-pguser-mstruser.yaml
   ```

6. Confirm mstruser and postgres secrets are available in the default namespace.

   ```
   kubectl get secrets -n mstr | grep 'analyticsdb-pguser-mstruser\|
   analyticsdb-pguser-postgres'
   ```

7. Restart mstr-srv and mstr-web pods.

Chapter 8: Troubleshooting Avaya Analytics™ migration issues

Migration from same source to other target error

Condition

Once the migration is complete from source to target. Then when we try to migrate from same source to other target, erorr pops up.

Solution

1. Run **`ccm release orca-dbmgr migration2to5.sh`**
2. Select option 2 **cleanup failed attempt**.
3. Run command with option 1.

Chapter 9: Troubleshooting Messaging

Verifying the status of the Messaging channel

About this task

Use this procedure to check whether the Messaging channel is operational. A successful result ensures that the Cloud Provider can access the new Messaging APIs using your external load balancer, or using the Avaya Common Services (Common Services) Cluster FQDN directly.

Important:

- Do not run this process from the Cluster Control Manager (CCM) console. Run it from another PC that has cURL installed and is on the customer's network.

Procedure

1. To check the health status of the Messaging channel:

   Run the following command, replacing *{Common Services.FQDN}* with the FQDN of the customer's Common Services cluster.

   `curl -X GET http://{Common Services.FQDN}:31325/messaging/v1/health-check`

   Verify the following response:

   ```
   {
       "alive": true
   }
   ```

2. To check the readiness of the Messaging channel:

   Run the following command, replacing *{Common Services.FQDN}* with the FQDN of the customer's Common Services cluster.

   `curl -X GET http://{Common Services.FQDN}:31325/messaging/v1/ready`

   Verify the following response:

   ```
   {
       "5e46bd56cb011b001076c314": {
           "Aggregator": {
               "Name": "Smooch",
               "State": "200"
           },
           "ContactCenter": {
               "Name": "Oceana",
               "State": "READY"
   ```

```
            }
        }
}
```

Contact not getting created in Avaya Oceana® or messages not flowing into Avaya Oceana®

Condition

Contact is not getting created in Avaya Oceana®, or messages are not flowing into Avaya Oceana®.

Solution

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. To switch to the root user, type su - and press **Enter**.
3. Run the following command to view the last log entries for one of the microservices:

   ```
   for i in $(kubectl get pods | grep async-oceana-adapter | awk
   '{print $1}'); do echo "**** Logs for Pod $i *****"; echo " ";
   kubectl logs $i async-aggregator-interface | tail -n 100; done
   ```

4. If the logs display data flowing in from the customer, do the following:
   a. Verify that your Messaging endpoint is configured correctly according to your hardware deployment configuration, and is publicly accessible.
   b. If the endpoint is publicly accessible, contact the DevOps team and ensure the team has correctly configured the Messaging endpoint. The DevOps team can also view the logs of the Cloud Provider applications to check whether the Cloud Provider can make a request to your Avaya Common Services (Common Services) Cluster.
   c. If steps a and b are correct, then verify that the customer is configured to the correct Cloud Provider application ID.
5. If the logs do not display data flowing in from the customer, do the following:
   a. Verify that there are no errors, the tenant/application ID is correctly configured for the deployment, and the message is sent.
   b. If there are no errors in the logs, run the following command:

      ```
      for i in $(kubectl get pods | grep async-oceana-adapter | awk
      '{print $1}'); do echo "**** Logs for Pod $i *****"; echo " ";
      kubectl logs $i async-oceana-adapter | tail -n 100; done
      ```

6. If the logs display data flowing in from the customer, but there are errors while making a request to Avaya Oceana®, do the following:
   a. Verify that the Avaya Oceana® FQDN is correct and follows the standard FQDN specifications.
   b. Verify that Avaya Oceana® is operational.
   c. Verify that the System Manager CA and Identity certificate are configured for Messaging.

Note:

You can also view the logs using Kibana by accessing `https://<your.Common Services.cluster.fqdn>/app/kibana`.

The following KSQL queries help you isolate the issue:

- Get all logs for the Messaging Connector microservices:

  ```
  kubernetes.pod_name :*async* and not
  kubernetes.container_image :*istio*
  ```

- Get any errors occurred for the Messaging Connector microservices:

  ```
  kubernetes.pod_name :*async* and not
  kubernetes.container_image :*istio* and level :error
  ```

- Track a specific Message ID as it flows through the system:

  ```
  kubernetes.pod_name :*async* and not
  kubernetes.container_image :*istio* and
  message :*{the.message.id}*
  ```

Async messages not getting to the customer

Condition

Async messages are not getting to the customer.

Solution

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. To switch to the root user, type su - and press **Enter**.
3. Run the following command to view the last log entries for one of the microservices:

   ```
   for i in $(kubectl get pods | grep async-oceana-adapter | awk
   '{print $1}'); do echo "**** Logs for Pod $i *****"; echo " ";
   kubectl logs $i async-aggregator-interface | tail -n 100; done
   ```

4. If the logs do not display the messages flowing from Avaya Oceana®, verify that Avaya Oceana® is operational.
5. If the logs display the messages flowing from Avaya Oceana®, run the following command:

   ```
   for i in $(kubectl get pods | grep async-aggregator-interface |
   awk '{print $1}'); do echo "**** Logs for Pod $i *****"; echo " ";
   kubectl logs $i async-oceana-adapter | tail -n 100; done
   ```

6. If you observe any errors in calling Cloud Provider, do the following:
   a. Check whether the security details of Cloud Provider are correct.
   b. Check whether Cloud Provider in non-operational.
   c. Check whether there is an intermittent problem with any Network Interface Card (NIC) where the Avaya Common Services cluster is installed.

7. If you do not find any errors and the messages appear to be sent to Cloud Provider, it might take a few minutes for Cloud Provider to propagate the messages to a customer. You can contact the DevOps team to check if there is an outage in Cloud Provider.

Chapter 10: Troubleshooting Avaya Analytics™ upgrade issues

Historical and real-time reporting fails after an upgrade

Condition

An upgrade from Avaya Analytics™ Release 4.0.0.1 to 4.1 followed by an upgrade from Avaya Oceana® Release 3.7 to 3.8 results in the following issues:

- Historical and real-time reporting fails
- Unified Collaboration Model (UCM) pumpup fails to complete

  The UCM pumpup process provides Avaya Analytics™ the latest state of all agents.

Cause

The UCM notification type list changes.

Solution

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To take a back up of the input adaptor configuration map to a file, run the folowing command:

   `k describe cm orca-ref-input-adaptor > input-adaptor-cm.txt`
4. To edit the input adapter configuration map, run the following command:

   `k edit cm orca-ref-input-adaptor`
5. In `command` mode in a vi editor, enter the following command to replace the UCM notification type string:

   `%s/FORWARD_NOTIFICATION/SEND_NOTIFICATION/g`
6. Write and quite the editor.
7. Log in to the Cluster Control Manager (CCM) console as the customer user.
8. Switch to being the root user by entering the command `su`.
9. Select **Deployment** by pressing the corresponding number.
10. Select **Service Restart Options** by pressing the corresponding number.

11. To restart the Reliable Eventing Framework Input Adaptor service, select the **Restart Ref Input Adaptor** option by entering the corresponding number.
12. In the **Proceed with REF Input Adaptor restart** field, enter `y`.

    Entering `n` cancels the operation.

    Wait for the service to start running.
13. To view the connection to REF and the pumpup process statuses, check the input adapter logs.

    The pumpup process must display the message that the process is complete.

Upgrade roll back

You can roll back to a previous release of Avaya Analytics™ when an upgrade fails. An upgrade can fail due to an incomplete patch install.

Use the following procedures to complete the roll back process.

Rolling back the Avaya Analytics™ software

About this task

Use this procedure to roll back the Avaya Analytics™ software to a previous version after a upgrade.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Copy the `Avaya_Oceana_Application_Deployment_4.x.xlsm` spreadsheet that you earlier saved on a secured location to the CCM server.
3. From the directory on CCM that contains the `Avaya_Oceana_Application_Deployment_4.x.xlsm`, run the following command:

   `screen`
4. Run the following command:

   `cc upgrade spec Avaya_Oceana_Application_Deployment_4.x.xlsm`
5. At the prompt, do the following:
   a. Accept the EULA.
   b. Confirm the upgrade.

   The Avaya Analytics™ software starts downloading and installing.
6. To monitor the progress of the install, run the following command:

   `tail -f /var/log/ avaya/ccm/ccm-main.log`

7. To check if the installation is successful, run the following command on the CCM console:

   ```
   ccm status
   ```

Rolling back the custom reports

About this task

Use this procedure to roll back the custom reports if an Avaya Analytics™ upgrade fails. The custom reports are exported before an upgrade to a secured location.

Before you begin

Roll back the Avaya Analytics™ software.

Procedure

1. Find the analyticsdb-node pods that has the master role.

   ```
   kubectl get pods --selector=postgres-operator.crunchydata.com/
   role=master | grep analyticsdb-node- | head -n1 | awk '{print $1;}'
   ```

2. Connect to the analyticsdb-node pods using the following command:

   ```
   kubectl exec -it <analyticsdb-node-Id> -- /bin/bash
   ```

3. List local backup folders to confirm saved backups:

   ```
   ls -l pgdata/md_full_backup/
   ```

4. If backup folder does not exists create a new folder to store backup:

   ```
   mkdir -p /pgdata/md_full_backup
   ```

5. To exit the pod, type `exit`.

6. On Cluster Control Manager, change the directory to where the backup was exported by using the following command:

   ```
   cd /var/avaya/artifactCache/historical_md_backups/
   ```

7. Copy the Historical Reporting backup file on to the analyticsdb-node pods by using the following command:

   ```
   kubectl cp <analyticsdb-node-Id>:/pgdata/md_full_backup/<backup-
   name>
   ```

8. Connect to the analytics database pod using the following command:

   ```
   kubectl exec -it <analyticsdb-node-Id> -- /bin/bash
   ```

9. To restore the full backup, run the following command:

   ```
   psql -U postgres -f /pgdata/md_full_backup/<backup-name> -d
   avaya_analytics_md
   ```

Note:

The CCM console might display errors related to the postgres user. Ignore the error messages.

10. Connect to the PostgreSQL database using the following command:

    ```
    psql
    ```

11. To view the list of available databases, run the following command:

    ```
    postgres=> \l
    ```

12. Confirm the avaya_analytics_md database is available.
13. To quit the database, run the following command:

    ```
    postgres=> \q
    ```

14. To exit the pod, type `exit`.
15. Verify that the custom reports are restored.

Rolling back to the previous version of the database

About this task

Use this procedure to roll back the Avaya Analytics™ database if an upgrade fails.

Warning:

Do not restore a database backup that was taken at an earlier patch level to the current database.

Before you begin

- Roll back the software.
- Roll back the custom reports.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Copy the backup of the database taken before running the upgrade to CCM.
3. To remove the postgres user access to the database, do the following:
   a. As the root user, run the following command:

      ```
      kubectl exec -it crunchy-primary-service-orca-dbmgr-0 -- //bin/bash
      ```

   b. To open a connection to the PostgreSQL command line tool, run the following command:

      ```
      psql
      ```

   c. Run the following command:

      ```
      REVOKE CONNECT ON DATABASE analytics_db FROM PUBLIC, postgres;
      ```

4. To end the database connections, do the following:
   a. If you are not connected to the crunchy pod, switch to root user and run the following command:

      ```
      kubectl exec -it crunchy-primary-service-orca-dbmgr-0 -- //bin/bash
      ```

   b. To open a connection to the PostgreSQL command line tool, run the following command:

      ```
      psql
      ```

   c. To disconnect from the PostgreSQL database, run the following command:

      ```
      SELECT pg_terminate_backend(pg_stat_activity.pid) FROM pg_stat_activity WHERE pg_stat_activity.datname = 'analytics_db' AND pid <> pg_backend_pid()
      ```

5. To drop the database, do the following:
   a. As the root user, run the following command:

      ```
      kubectl exec -it crunchy-primary-service-orca-dbmgr-0 -- //bin/bash
      ```

   b. To open a connection to the PostgreSQL command line tool, run the following command:

      ```
      psql
      ```

   c. Run the following command:

      ```
      DROP DATABASE "analytics_db";
      ```

6. To restore the database, run the CCM control script command.
7. Verify that the system is restored to the previous state by completing the post installation tasks.
8. To take a back up of the database, use the CCM control script.

Avaya Common Services upgrade revert

You can do a complete upgrade revert and re-install the previous release. For details, see the steps in the next section.

Reverting a failed upgrade to the previous release

About this task

After a failed upgrade, use this procedure to downgrade or revert back to the previous release of Avaya Common Services.

Before you begin

Ensure that you have:

- The original solution configuration spreadsheet used to install the previous release to which you are reverting back.

Procedure

1. Run `ccm uninstall --force` to uninstall the cluster and services as required.

   If the cluster uninstall fails, contact Avaya support personnel for assistance.

2. In vCenter, revert Cluster Control Manager to the snapshot image taken before the upgrade.
3. Log in to the restored snapshot of Cluster Control Manager and run the `ccm uninstall-cluster --force` command.

   If the node removal fails, contact Avaya support personnel for assistance. Manual cleanup might be required.

4. On Cluster Control Manager, type `screen` to run the installation in the background.
5. Using the original solution configuration spreadsheet, which is compatible with the release you are reverting back to, run the `ccm install <solution spreadsheet name>.xlsx` command.
6. Restore application data using the information in your solution documentation.

Unable to view the full list of routing services after pumpup

Condition

You cannot view the full list of routing services, which are configured in Avaya Control Manager, in the Routing Service or Agent by Routing Service historical reports after pumpup.

Cause

Presently, Avaya Analytics™ publishes metrics only for routing services that were recently upgraded during pumpup.

If there has been no activity on a routing service or a set of routing services, these routing services do not display in a realtime view after pumpup. Only the routing services that have undergone any activity display in Routing Service or Agent by Routing Service reports after pumpup.

Chapter 11: Troubleshooting Avaya Analytics™ installation and post install script

Troubleshooting Avaya Analytics™ installation errors

About this task

Use this procedure to troubleshoot installation errors that might result in installation failure. The reasons can vary from entering incorrect password, an incorrect value in the deployment spreadsheet excel file, or network issues.

Warning:

Do not delete nodes or VMs from your VMware.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Run one of the following commands:
   a. To resume the command, run:

      `ccm install resume`

   b. To cancel the command, run:

      `ccm install cancel`

   If the CCM console displays any of the following messages, it indicates that the cluster is installed, but the services are not deployed:

   - `Staged product not found: resume`
   - `No staged products, no helm releases`
3. To deploy the services, run the following command:

   `ccm upgrade spec <deployment spreadsheet>.xlsm`. where deployment spreadsheet is deployment excel spreadsheet used for the installation of the Avaya Analytics™ version.

Avaya Workspaces real-time dashboard does not update after installation

Condition

Avaya Workspaces real-time dashboards are not loading.

Solution

1. On the Avaya Control Manager webpage, click **Configuration** > **Customer Engagement** > **Avaya Analytics**.
2. On the Avaya Analytics Server List page, double-click the configured Analytics server.
3. Select the **Streams Servers** tab.
4. On the Avaya Analytics Server Edit page, double-click the Avaya Analytics Server.
5. In the **Name** field, enter the name of the Analytics server.
6. In the **FQDN** field, enter the FQDN of your Analytics Cluster.
7. In the **Port** field, enter `443/orca-streams-rest`.
8. Select the **TLS Flag** check box.
9. Click **Save**.

The ccm release orca analytics command fails

Condition

When trying to run the post install scripts for from the CCM console, the **`ccm release orca analytics`** command does not display the maintenance and troubleshooting options.

Solution

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. To find the location of the Avaya Analytics™ script, run the following command:

   `# find / -name analytics`

   Note:

   The path can be different in your environment.

   The CCM console displays the following:

   ```
   /opt/avaya/flex/clusters/analytics-env7/staging/orca-5.0.265/orca/
   avaya-flex/bin/release/analytics
   ```

4. Run the following command:

```
python3 /opt/avaya/flex/clusters/analytics-env7/staging/
orca-5.0.265/orca/avaya-flex/bin/release/analytics
```

The CCM console displays the options for maintaining and troubleshooting Avaya Analytics™.

Troubleshooting Analytics online and offline deployment issues

Use of unknown CA certificate

Condition

Docker client is not able to validate the docker server certificate while downloading or uploading the Avaya Analytics™ chart and images.

Cause

The certificate was not correctly imported.

Solution

Import the certificate again.

Unauthorized or failed logins

You are not able to log in while downloading the Avaya Analytics™ chart and images

Cause

- The user credentials you have used are incorrect.
- Invalid certificate in Docker server.
- No certificate of Docker server installed in Docker desktop client.

Solution

Ensure your Avaya SSO credentials are correct and then try to log in again.

Avaya Analytics™ chart and images fail to download

Condition

During an air gap deployment using an outbound proxy, the Avaya Analytics™ chart and images fail to download when the user runs the **`agn download`** command in the ccm-agn-ctl container.

Cause

CA certificate of image repository server is not valid or updated.

Solution

1. In the ccm-agn-ctl container, copy the `.pem` or the `.cer` version of your proxy CA certificates to the following directory:

   `/etc/pki/ca-trust/source/anchors/`

2. Run the following command:

   `update-ca-trust`

3. Run the following command:

   `agn download`

Avaya Analytics™ chart and images fail to upload

Condition

During an air gap deployment using an outbound proxy, the Avaya Analytics™ chart and images fail to upload to Cluster Control Manager (CCM) when the user runs the **`agn download`** command in the ccm-agn-ctl container.

Cause

- The air gap network container is not configured on CCM.
- The air gap network container on CCM is not running.
- The CCM FQDN certificate is not present in the air gap network container.

Solution

Add the CCM certificates to Docker Desktop client.

In real-time reporting the Routing Service Group producer is not visible

Cause

Routing Service Group producer is not installed during Avaya Analytics™ deployment.

Solution

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. Check if the processor is deployed, run the following command:

   ```
   [root@ccm~]#kubectl get pods --all-namespaces | grep orca-itd-agent-group-measure-proc
   ```

4. **(Optional)** Open deployment spreadsheet, go to Orca tab > Routing Service Group Measure Processor field.

If the value of the field is set as `False`, then Routing Service Group Measure Processor is not deployed. To deploy this processor, see section *Preparing the deployment spreadsheet* in *Deploying Avaya Analytics™ for Avaya Oceana®* guide.

Historical data not available in the database and on Agent Trace reports in MSTR

Historical data is not available in the database and on canned or custom Agent Trace reports in MSTR. This data is of agents who are being traced while they are logged onto Avaya Workspaces and handling calls.

Cause

Orca-trace-measure-proc is not installed during Avaya Analytics™ deployment.

Solution

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. Switch to being the root user by entering the command `su`.
3. Check if the processor is deployed, run the following command:

   ```
   [root@ccm~]#kubectl get pods --all-namespaces | grep orca-trace-measure-proc
   ```

4. **(Optional)** Open deployment spreadsheet, go to **Orca** tab > **Customize Processor to Install** field. If the value of the field is set as `False`, then set it to `True`.

   Under Measure Processors, if the **Install** value for **Agent Trace Measure Processor** is set to `False`, then the Agent Trace Measure Processor is not deployed.

   To deploy this processor, see section *Preparing the deployment spreadsheet* in *Deploying Avaya Analytics™ for Avaya Oceana®* guide.

Error mapping user for SAML

Cause

This error can occur while mapping a SAML user to a Historical reporting local user. This error can occur if username does not exist or must be created before mapping for SAML.

Solution

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. To switch to the root user, type **su** and press **Enter**.
3. To run the Analytics Administration script, use the following command:

   ```
   ccm release orca analytics
   ```

4. To select the **Historical Reporting** option, enter the corresponding number.

5. To select the **SAML** option, enter the corresponding number.
6. To select the **Map SAML users to Historical reporting**, enter the corresponding number.
7. To **Map a single SAML user to Historical reporting**, enter the corresponding number.
   a. Enter the **Username** of the local or LDAP user.
   b. Enter **Name ID** of the SAML user from the SAML assertion.

      The entered SAML user is mapped with Historical reporting local user or LDAP user.

Avaya Workspaces does not load real-time data

Cause

Cross-Origin Resource Sharing (CORS) is not configured appropriately.

Solution

Check the deployment spreadsheet to see if following fields are configured:

config:orca-streams-data-publisher:virtualService:allowOriginsSingleExact

config:orca-streams-rest:virtualService:allowOriginsSingleExact

The field value should be taken from the origin field of the HTTP request to the orca-streams-rest and orca-streams-data-publisher. For example, Go to **WORKSPACES** > **Chrome** > **Developer tools** > **Network** > **XHR** > **any orca-streams-rest URL**.

For more information on configuring the deployment spreadsheet, see *Preparing the deployment spreadsheet* section in *Deploying Avaya Analytics™ for Avaya Oceana®* guide.

Chapter 12: Troubleshooting common issues

Macros of deployment spreadsheet are disabled by administrator

Condition

Unable to enable the Macros of the deployment spreadsheet, as it displays the following error:

```
Macros in this document have been disabled by your enterprise
administrator for security reasons
```

Cause

Files are not in a trusted location of your machine.

Solution

Place your files, including the deployment spreadsheet, in a trusted location.

a. Go to **File** > **Options** > **Trust Center** > **Trust Center Settings**.
b. On the left navigation pane, click **Trust Locations**.
c. If a trusted location exists, select the location where you want to move the files.
d. Select the deployment spreadsheet from the current location and click **Ok**.
e. If there is no trusted location, first create one and then move the files.

Issues in localized spreadsheet

Condition

Issues while using excel spreadsheet in a language other than English.

Solution

1. Go to **Control Panel** and search for **Region**.
2. On Region panel, click **Additional settings**.
3. In the **Decimal symbol** field, change the separator from `,` to `.`.

4. Click `Ok`.
5. Save the spreadsheet.

   For information on global number formatting in MS Excel, see https://docs.microsoft.com/en-us/globalization/locale/number-formatting.

Error getting initialization data for producer: undefined

Condition

Supervisor see this banner presented when they login to Avaya Workspaces.: `Error getting initialization data for producer: undefined`.

Cause

Pump-up not completed.

Solution

1. Log out from Avaya Workspaces.
2. Wait for pump-up to complete.
3. Log in to Avaya Workspaces.

Restarting Avaya Analytics™ after applying Avaya Oceana® patches

About this task

If the Avaya Oceana® restarts for any reason such as, applied Avaya Oceana® patches, it is necessary to restart Avaya Analytics™.

Before you begin

- Delete the Reliable Eventing group.
- Create the Reliable Eventing group.

Procedure

1. Log in to the Cluster Control Manager (CCM) console as the customer user.
2. To run the Avaya Analytics™ Administration script, use the following command:

   ```
   ccm release orca analytics
   ```

3. To select the **Deployment** option, enter the corresponding number.
4. To select the **Service Restart Options** option, enter the corresponding number.
5. To select the **Restart Ref Input Adaptor** option, enter the corresponding number.

The restart can take several minutes to complete.

6. In the **Proceed with REF Input Adapter restart** field, type `y`.

   Typing `n` cancels the operation.

7. To return to the previous page, type `b` and press **Enter**.
8. To quit from the current page, type `q` and press **Enter**.
9. To return to the main menu, type `m` and press **Enter**.
10. Log in to System Manager web console, go to **Elements** > **Avaya Breeze®** > **Reliable Eventing Administration** > **Destination Status**.
11. Scroll down to the **Analytics.IA_12345.QueueUCM.PUMPUP** line item. This is usually on the second page.
    a. If the number of En-queued Messages is greater than 1, skip the rest of this procedure.
    b. If the number of En-queued Messages is zero, restart Avaya Analytics™.

    For the detailed steps of restarting Avaya Analytics™, see the *Deploying Avaya Analytics™* document.

Realtime reports and Historical reports are inaccessible after node outage

Condition

Users are unable to access Realtime or Historical reports, and any URL using the Cluster FQDN is inaccessible.

Cause

IPVS routing table are stalled due to a down node.

No HTTP request are routed from clusterIP to ingressgateway endpoints.

Solution

For Avaya Analytics™ release 4.0.0.1 and 4.1.1.0, perform following steps:

1. Login to CCM as root
2. Get the names of the kube-keepalived-vip pods

   ```
   kubectl get pods | grep common-services-kube-keepalived-vip
   ```

3. Restart each of the kube-keepalived-vip pods

   ```
   kubectl delete <pod_1> <pod_2> <pod_3>
   ```

   Once the kube-keepalived-pods is restarted and running, this issue is resolved. Realtime Reports and Historical Reports should be accessible.

Deleting ORCA product results in PV and data deletion

Condition

On deleting the ORCA product or any other product, the related PV is deleted automatically. User data is stored on ORCA product, therefore, data is also deleted automatically.

Solution

- This is an expected behavior and is not an issue unless for some reason a rebuild of Avaya Analytics™ install is required.
- Before deleting ORCA product on a client site, please contact the Avaya support team on how to avoid deleting the PV contain data.

Restrictions to vCenter user account password used for Avaya Analytics™

Condition

During installation or upgrade Avaya Analytics™ fails with incorrect user name or password error.

The following list of special characters cannot be used in the vCenter user account password provided during ccm install or ccm upgrade spec --infra:

` $ ( ) \ | ; : ' " < >

Cause

Avaya Analytics™ installation or upgrade fails if the vCenter user account password consists special characters listed above. The vCenter account is disabled/locked due to multiple failed log in attempts in Avaya Analytics™.

An exception to the above restriction is using a single $ symbol at the end of a password is allowed.

Solution

- When installing Avaya Analytics™ or upgrading from Avaya Analytics™ 4.0.0.1 patch 8 to later release, you must update your vCenter password to comply with the above restriction.
- Run the following command to update the vCenter user account password that is using before installing or upgrading Avaya Analytics™ to the higher release:

  ```
  ccm infra update-vcentercreds
  ```

Chapter 13: Resources

Documentation

Title	Use this document to:	Audience
Overview		
Avaya Oceana® Solution Description	Use this guide to know about the tested product characteristics and capabilities, including product overview and feature descriptions, interoperability, performance specifications, security, and licensing requirements.	• Sales engineers • Business partners • Solution architects • Implementation engineers
Implementing		
Deploying Avaya Oceana®	Use this guide to know how to deploy Avaya Oceana® Solution on the customer environment.	• Sales engineers • Business partners • Solution architects • Implementation engineers
Avaya Oceana® and Avaya Analytics™ Disaster Recovery	Use this guide to know how to restore Avaya Oceana®, solution when there is a complete outage at the primary data center.	• Sales engineers • Business partners • Solution architects • Implementation engineers
Migrating Avaya Oceana®	Use this guide to know how to migrate Avaya Oceana® solution from the existing version.	• Sales engineers • Business partners • Solution architects • Implementation engineers
Deploying Avaya Analytics™	Deploy Avaya Analytics™.	• Sales engineers • Business partners • Solution architects • Implementation engineers
Administering		

Table continues…

Title	Use this document to:	Audience
Administering Avaya Oceana®	Administer Avaya Oceana®.	• System administrators • Supervisors
Using		
Using Avaya Workspaces for Avaya Oceana®	Use Avaya Workspaces for Avaya Oceana®.	• Agents • Supervisors
Using Avaya Analytics™	Use the features and capabilities of Avaya Analytics™.	• Supervisors • Administrators • Report designers
Avaya Analytics™ Data Dictionary	Use historical and real-time measures in custom reports.	• Administrators • Report designer
Maintaining and Troubleshooting		
Maintaining and Troubleshooting Avaya Oceana®	Perform maintenance and troubleshooting procedures for routine maintenance and troubleshooting of Avaya Oceana®.	• Support personnel • Implementation engineers • Administrators
Maintaining and Troubleshooting Avaya Analytics™	Perform common maintenance functions of Avaya Analytics™ and use tools and utilities for troubleshooting of Avaya Analytics™.	• Support personnel • Implementation engineers • Administrators
Avaya Oceana® Alarms	View details about Avaya Oceana® alarms.	• Support personnel • Administrators

Finding documents on the Avaya Support website

Procedure

1. Go to https://support.avaya.com.
2. To log in, click **Sign In** at the top of the screen and then enter your login credentials when prompted.
3. Click **Product Support** > **Documents**.
4. In **Search Product**, start typing the product name and then select the appropriate product from the list displayed.
5. In **Select Release**, select the appropriate release number.

   This field is not available if there is only one release for the product.
6. **(Optional)** In **Enter Keyword**, type keywords for your search.
7. From the **Select Content Type** list, select one or more content types.

For example, if you only want to see user guides, click **User Guides** in the **Select Content Type** list.

8. Click to display the search results.

Avaya Documentation Center navigation

For many programs, the latest customer documentation is available on the Avaya Documentation Center website at https://documentation.avaya.com. Some functionality is only available when you log in to the Avaya Documentation Center. The available functionality depends on your role.

Important:

If the documentation you are looking for is not available on the Avaya Documentation Center, you can find it on the Avaya Support website.

While navigating through the Documentation Center, you can click the **Avaya Documentation Center** logo at the top of the screen to return to the home page anytime. On the Avaya Documentation Center, you can do the following:

- Click **Avaya Links** in the top menu bar to access other Avaya websites, including the Avaya Support website.
- Click **Languages** ( ) in the top menu bar to change the display language and view localized documents.
- In the **Search Documentation** field, search for keywords and click **Filter** to filter by solution category, product, or user role.

  You can select multiple items in each filter category. For example, you can select a product and multiple user roles.
- Click **Library** in the top menu bar to access the complete library of documents. Use the filtering options to refine your results.
- After performing a search or accessing the library, you can sort content on the search results page. When you find the item you want to view, click it to open it.
- Use the table of contents in a document for navigation. You can also click **<** or **>** next to the document title to navigate to the previous topic or the next topic.
- Click **Share** ( ) to share a topic by email or copy the URL.
- Download a PDF of the current topic in a document, the topic and its subtopics, or the entire document.
- Print the section you are viewing.
- Add content to a collection by clicking **Add to My Topics** ( ). You can add the topic and its subtopics or add the entire publication.
- View the topics in your collections. To access your collections, click your name in the top menu bar and then click **My Topics**.

You can do the following:

- Create, rename, and delete a collection.
- Set a collection as the default or favorite collection.
- Save a PDF of the selected content in a collection and download it to your computer.
- Share content in a collection with others through email.
- Receive collections that others have shared with you.

- Click **Watch** ( ) to add a topic to your watchlist so you are notified when the content is updated or removed.
- View and manage your watchlist by clicking **Watchlist** from the top menu with your name.

  You can do the following:

  - Enable **Email notifications** to receive email alerts.
  - Unwatch the selected content or all topics.
- Send feedback for a topic.

Training

The following courses are available for the Avaya Oceana® program.

Table 1: Sales Credentials

Course code	Course title	Course duration in hours	Delivery type
APSS – 1202 Avaya OneCloud™ CCaaS Sales			
41511W	Selling Avaya OneCloud™ CCaaS Solutions	0.75	Web-based Training
41551T	Avaya OneCloud™ CCaaS Sales Specialized Test	1.0	Web-based Training
ALCC –2005 Avaya Multiexperience Solutions Sales (ALCC-2005)			
41710W	The Avaya OneCloud™ Contact Center Automated Story	0.50	Web-based Training
41411W	Selling Avaya Oceana®	0.75	Web-based Training
41401W	Selling Avaya Analytics™	0.50	Web-based Training
41481W	Avaya Oceana® ROI for Sales	0.50	Web-based Training
41770W	Avaya Experience Portal and Proactive Outreach Manager (POM) for Sales	0.25	Web-based Training

Table 2: Pre-Sales Design

Course code	Course title	Course duration in hours	Delivery type
ACDS – 3480 Avaya Oceana® Solution Design			
34211W	Avaya Oceana® Overview for Design	0.75	Web-based Training
34811W	Designing the Avaya Oceana Solution Part 1 of 3	1.0	Web-based Training
34821W	Designing the Avaya Oceana Solution Part 2 of 3	1.0	Web-based Training
34831W	Designing the Avaya Oceana Solution Part 3 of 3	1.0	Web-based Training
34801X	Avaya Oceana® Solution Design Exam	1.50	Exam
ALRI-7001 Avaya Oceana® Product Release Information Collection			
39001W	Avaya Oceana® R3.8 with Breeze Snap-ins Details for Pre-Sales	1.0	Portable Document Format (PDF)
39020W	Avaya Breeze® Snap-ins for Avaya Oceana Details for Pre-Sales	1.0	PDF

Table 3: Technical Services Partner Credentials

Course code	Course title	Course duration in hours	Delivery type
ACIS – 7495 Avaya Oceana® Solution Implement			
74150V	Integrating Avaya Oceana® Core and Workspaces	40.0	Virtual Instructor-Led Training
74950X	Avaya Oceana® Solution Integration Exam	1.50	Exam
ACSS-7497 Avaya Oceana®			
74550V	Supporting Avaya Oceana®	24	Virtual Instructor-Led Training
7497X	Avaya Oceana® Support Exam	1.75	Exam
74360W	Installing Avaya Analytics™ for Oceana®	1.5	Web-based Training

Table 4: Pre-requisite Courseware

Course code	Course title	Course duration in hours	Delivery type
77900W	Avaya Control Manager Training Bundle (5 courses 21900W, 77910W, 77920W, 77930W, 77940W)	5.50	Web-based Training
70160W	Avaya Breeze® Implementation and Support	30.0	Web-based Training

Table 5: End User, Programmer, Administration

Avaya Learning Center				
Course code	**Course title**	**Course duration in hours**	**Delivery type**	**Vanity Link for Attachment**
ALEU-5002 Avaya Oceana® End-User Training				
24020W	Using Avaya Workspaces for Avaya Oceana® - Agent	1.0	Web-based Training	https://www.avaya.com/oceana-agent
24040W	Using Avaya Workspaces for Avaya Oceana® - Supervisor	1.0	Web-based Training	https://www.avaya.com/oceana-supervisor
ALUC-4001 Avaya Breeze® Client SDK				
2410W	Customer Communications and Apps with Oceana® for Developers	3.0	Web-based Training	
ASDC-0010 Avaya Workspaces® Framework				
24150W	Customizing the Avaya Workspaces® Framework	3.0	Web-based Training	
24150T	Avaya Workspaces® Framework R3 Test	1.0	Online Test	
ASAC-0005 Avaya Oceana® Administration				
21160W	Avaya Oceana® Fundamentals	0.5	Web-based Training	
24300V	Administering Avaya Oceana® R3 Omnichannel	40.0	Virtual Instructor-Led Training	Attached with the sale
2430T	Administering Avaya Oceana® R3 Online Test	1.0	Online Test	
24320W	Administering Avaya Oceana® - Basic	2.5	Web-based Training	https://www.avaya.com/Oceana-admin

Table continues…

Avaya Learning Center				
Course code	**Course title**	**Course duration in hours**	**Delivery type**	**Vanity Link for Attachment**
ASAC-0031 Avaya Analytics™ R4 for Oceana® Administrator				
24380T	Administering Avaya Analytics1M R4 for Oceana8 Specialized Test	1.0	Online Test	

Table 6: Other Miscellaneous Courseware

Course code	Course title	Course duration in hours	Delivery type	Vanity Link for Attachment
ALCC-0001 Avaya Workforce Optimization Select Integration with Avaya Oceana® Workspaces				
7014W	Integrating Avaya Workforce Optimization Select with Avaya Oceana® Workspaces	3.0	Web-based Training	
7014A	Avaya Workforce Optimization Select with Avaya Oceana® Workspaces Integration Assessment	1.0	Assessment	
71610W	Integrating POM with Avaya Oceana®	1.0	Web-based Training	

Viewing Avaya Mentor videos

Avaya Mentor videos provide technical content on how to install, configure, and troubleshoot Avaya products.

About this task

Videos are available on the Avaya Support website, listed under the video document type, and on the Avaya-run channel on YouTube.

- To find videos on the Avaya Support website, go to https://support.avaya.com/ and do one of the following:
  - In **Search**, type `Avaya Mentor Videos`, click **Clear All** and select **Video** in the **Select Content Type**.
  - In **Search**, type the product name. On the Search Results page, click **Clear All** and select **Video** in the **Select Content Type**.

    The **Video** content type is displayed only when videos are available for that product.

  In the right pane, the page displays a list of available videos.

- To find the Avaya Mentor videos on YouTube, go to www.youtube.com/AvayaMentor and do one of the following:
  - Enter a keyword or keywords in the **Search Channel** to search for a specific product or topic.
  - Scroll down Playlists, and click a topic name to see the list of videos available. For example, Contact Centers.

Note:

Videos are not available for all products.

Support

Go to the Avaya Support website at https://support.avaya.com for the most up-to-date documentation, product notices, and knowledge articles. You can also search for release notes, downloads, and resolutions to issues. Use the online service request system to create a service request. Chat with live agents to get answers to questions, or request an agent to connect you to a support team if an issue requires additional expertise.

Using the Avaya InSite Knowledge Base

The Avaya InSite Knowledge Base is a web-based search engine that provides:

- Up-to-date troubleshooting procedures and technical tips.
- Information about service packs.
- Access to customer and technical documentation.
- Information about training and certification programs.
- Links to other pertinent information.

If you are an authorized Avaya Partner or a current Avaya customer with a support contract, you can access the Knowledge Base without extra cost. You must have a login account and a valid Sold-To number.

Use the Avaya InSite Knowledge Base for any potential solutions to problems.

1. Go to https://support.avaya.com.
2. To log in, click **Sign In** at the top of the screen and then enter your login credentials when prompted.
3. Click **Product Support** > **Products**.
4. In **Search Product**, start typing the product name and then select the appropriate product from the list displayed.

5. Select the release number, if applicable.
6. Click the **Technical Solutions** tab to view articles for resolving technical issues.

Index

Special Characters

_troubleshooting
    pods 152
.tar log files
    ccm report 202

A

access to Historical Reporting web page
    troubleshoot 148
adding
    local user to group 58
adding nodes 112
ADFS 86
After patches 251
agent by routing service
    data not showing 147
Agent Trace report 88
Agent tracing report 89, 90
AIDE logs 204
    viewing in Kibana 122
alarm
    pod crash 181
alarms not reported on NMS 180
Analytics
    backup 22
analytics 4.1.0.1 does not receive events from Oceana 3.7.0.1 with Async enabled 156
Analytics upgrade roll back 239
anti-affinity rule 178
assessing
    cluster state 128
audit rule updates 124
audit rules
    Linux 124
    Red Hat 124
Avaya Analytics
    restart 160
    troubleshoot 137
Avaya InSite Knowledge Base 261
Avaya support website 261

B

backing up 124
    metadata 72
backup 124
backup options 194
backups not restored correctly 163
breeze-security secret 48, 49

C

Call Originator Redaction
    configure 15
CCM
    backup 125
    restore 132
ccm archive command options 192
ccm archive config command options 192
ccm backup options 194
CCM OVF
    proxy settings 118
ccm release common-services alarmctl 223
ccm release common-services cmonctl 225
ccm release common-services getlogs 226
ccm release common-services updateIp 225
ccm release eventing-kafka kafka-util 226
ccm restore command options 196
ccmNetSetup
    proxy settings 119
certificate expiration alarming 46
certificate import
    specific trust stores 168
certificate rotation 208
certificates 175
certificates for secure LDAP connection
    configure 68
certification expiration alarms 46
changing
    default settings 12
check
    EASG availability 71
checkInfra
    usage 220
checking
    disk space usage 81
    high availability 83
    missing Kafka topics 154
close SSH session 134
cluster
    cannot connect 175
    certificates 175
    powering on 101
Cluster Control Manager
    configuring proxy settings 118
cluster node anti-affinity 178
cluster node restore 111, 135
cluster state 128
clusterVPN command options
    clusterVPN disable 112
    clusterVPN restart 112
    clusterVPN setup 112
    clusterVPN start 112

clusterVPN command options *(continued)*
clusterVPN status ... 112
clusterVPN stop ... 112
clusterVPN commands ... 219
collecting logs
file integrity validation ... 204
collection
delete ... 256
edit ... 256
generating PDF ... 256
sharing content ... 256
command
ccm report ... 197
ccm resizePVC command ... 218
ccm swhistory ... 215
checkInfra ... 220
pvcCleanup ... 217
command help output
checkInfra ... 220
command help outputs ... 220
alarming ... 223
certificates ... 223
Cmonitor service ... 225
eventing ... 226
eventing-kafka service ... 226
logging ... 226
monitoring ... 225
command options
ccm archive ... 192
ccm archive config ... 192
ccm restore ... 196
commands
certificate manager ... 205
cluster VPN ... 219
common services ... 209
core CCM ... 184
eventing ... 210
eventing-kafka ... 210
internal certificates ... 207
kafka ... 210
miscellaneous ... 210
Complete upgrade revert ... 242
configure ... 62, 84
configure Historical Reporting with LDAP
troubleshoot ... 142
configuring ... 47
Active Directory Federation Services ... 86
Analytics email distribution ... 62
Call Originator Redaction ... 15
certificates for secure LDAP connection ... 68
daily data roll-up ... 42
database backup ... 24
DataWarehouse Password ... 56
group filters ... 12
Historical Reporting ... 55
incremental database backup ... 32
LDAP authentication ... 66
configuring *(continued)*
Okta as IDP ... 84
SNMP alarm destinations ... 43
user management ... 19
Configuring dimensions ... 14
Configuring pgBackRest backup ... 32
configuring proxy settings
CCM OVF ... 118
ccmNetSetup command ... 119
Cluster Control Manager OVF ... 118
connection issues
cluster ... 175
content
publishing PDF output ... 256
searching ... 256
sharing ... 256
sort by last updated ... 256
watching for updates ... 256
copying
Historical Reporting logs ... 82
Copying analyticsdb-node secrets into mstr namespace ... 232
CPU ... 108, 109
creating
local user ... 57
metadata backups ... 72
read-only database user ... 19
creating index patterns ... 121
crunchy database report ... 96
csr file ... 173
CSRs
specifying identity certificates ... 165
CSV Producer
enable ... 16

D

data retention maximum limits ... 39
DataWarehouse Password
configure ... 56
Debug level logging ... 94
default Historical Reporting Administrator account
password details have been lost or misplaced ... 229
default settings
change ... 12
delete ORCA product ... 253
deleted agent group data in historical reports ... 231
deleting
Historical Reporting pods ... 83
local user from a group ... 61
deleting a failed service ... 176
deleting VM remnants ... 135
deletion of ORCA ... 156
deployment failed
DRS anti-affinity rule error ... 177
detaching
restore SSH session ... 134
disable EASG ... 70

disabling
routing service group ... 52
disk ... 109
Docker server certificate invalid ... 246
documentation center ... 256
finding content ... 256
navigation ... 256
documentation portal ... 256
DR monitoring tool ... 163, 164
dropping
database schemas ... 39

E

EASG
authentication ... 69–71
availability ... 69
Elasticsearch
AIDE logs ... 122
viewing in the Kibana interface ... 122
Elasticsearch logs ... 202
email distribution services ... 62
enable EASG ... 70
enabling
CSV Producer ... 16
routing service group ... 52
enabling or disabling zero row suppression parameter ... 97
Error getting initialization data for producer: undefined ... 251
Error in backup log ... 162
es_upload.sh script ... 203

F

failed login ... 246
failed upgrade
reverting ... 242
file integrity validation
collecting logs ... 204
disabling ... 120
enabling ... 120
finding content on documentation center ... 256
full metadata backup
schedule ... 75
full restore ... 127, 129

G

generating
CSV files for Agent by Account ... 16
pod status report ... 149
Routing Service Producer ... 16
generating CSRs
specific identity certificates ... 165
geo database replication for alarming ... 47
governing rule
update ... 64
grace period
license ... 178
Grafana reporting system ... 96
Grafana reports ... 95
group filters
configure ... 12
group settings
historical reporting ... 52

H

historical reporting
routing service group ... 52
user settings ... 54
Historical Reporting ... 84
configure ... 55
Historical Reporting logs ... 82
Historical Reporting logs in CCM ... 82
Historical Reporting secure LDAP ... 68
Historical Reporting with LDAP configuration
test ... 142, 148
HTTP proxy ... 118
HTTPS proxy ... 118

I

importing
migration packages ... 78
importing SAML users ... 92
importing third-party CA certificates
specifying trust stores ... 168
Inactivity Timeout Setting ... 10
incremental backup ... 162
Incremental backup ... 31
incremental database backup
schedule ... 33
index patterns
Kibana ... 121
logging service ... 121
install failure
invalid cluster configuration ... 176
terminal shell timed out ... 177
internal cluster certificates ... 207
introduction ... 10
IPD ... 86

K

KB
Support site ... 261
Kibana
log in ... 182
login error ... 182
viewing AIDE logs ... 122
Kubernetes
certificates ... 175

L

LDAP ... 69
LDAP authentication
  configure ... 66
license error message ... 178
local user
  create ... 57
localization excel sheet ... 250
localized spreadsheet ... 250
log in license error message ... 178
logging interface
  viewing AIDE logs ... 122
logs
  AIDE ... 204
  file integrity ... 204

M

Macros of deployment spreadsheet are disabled ... 250
Manage remote server ... 22
managing
  database user ... 19
  package migration ... 80
  scheduled remote backups ... 28
manually rotating certificates ... 208
manually specifying trust stores ... 168
Mapping SAML user ... 92
memory ... 108, 109
metadata
  back up ... 72
metadata backups
  create ... 72
MIB file ... 45
migrating
  packages ... 80
Migration from same source to other destination ... 233
monitoring
  certificates ... 71
monitoring crunchy database ... 95
monitoring system status ... 103
MSTR-SRV pod ... 231

N

network configuration
  altering ... 106
new reporting user ... 91
node ... 108, 109, 112
node anti-affinity ... 178

O

obtaining .tar files ... 202
obtaining log files ... 202
offline Analytics deployment
offline Analytics deployment *(continued)*
  chart and images fail to download ... 246, 247
offline environment
  restoring ... 129
open node port ... 95
Open Virtualization Format deployment
  proxy settings ... 118
outage
  Cluster Control Manager ... 132
outbound proxy configuration
  HTTP(S) ... 118
output
  checkInfra ... 220

P

package migration ... 80
packages
  migration ... 80
password restriction ... 253
pem file ... 172
PgBackRest backup ... 31
pods terminating ... 147
post install commands ... 12
post install scripts ... 12
postgres password change ... 163
proxy
  outbound ... 118
pumpup ... 243
PVC
  resize ... 114, 218
  resize by spreadsheet ... 113
PVC cleanup command ... 217

R

reattaching
  SSH session ... 134
recovering a deleted virtual machine ... 180
recovery
  full outage ... 127, 129
Ref Input Adaptor
  restart ... 160
registering ... 46
registration ... 115
related documentation ... 254
remnants in datastore
  deleting ... 135
remote backup ... 24, 26
Remote connection settings ... 22
remote desktop session ... 116
  ending the session ... 117
  help ... 118
  resetting the password ... 117
  terminating the session ... 117
remote server ... 24
remove

remove *(continued)*
high availability check ... 83
removing
local user from a group ... 59
local user from group ... 58
read-only database user ... 19
reopening the SSH session ... 134
replacing
authorization certificates ... 48
authorization certificates manually ... 49
migration package metadata ... 79
reports not accessible ... 252
resetting
dbwriterservice password ... 159
EASG password ... 157
local user from a group ... 60
mduserservice password ... 159
postgres password ... 158
Tomcat password ... 157
resizing a PVC ... 218
resizing a PVC using spreadsheet ... 114
Restart ... 251
restarting
a pod ... 148
agentbyaccount pod ... 16
Avaya Analytics ... 160
database manager ... 155
Historical Reporting ... 142
Historical Reporting with LDAP configuration ... 148
individual pod ... 148
measure processors ... 151
Ref Input Adaptor ... 160
routingservice pods ... 16
standby pods ... 149
restarting kafka pod ... 183
restore ... 124
restore command options ... 196
restoring ... 124
database ... 29, 36
database schemas ... 39
incremental database backup ... 37
restoring Cluster Control Manager ... 132
restoring common services ... 127
air gap ... 129
restoring from remote backup ... 36
restoring remote backup ... 29
resume VM issue ... 181
reverting
database version ... 241
reverting a failed upgrade ... 242
reverting to the previous release ... 242
roll back
previous database version ... 241
roll back custom reports ... 240
rolling back
Analytics upgrade ... 239
rotating certificates manually ... 208
routing service group ... 52, 53
Routing Service Group ... 230
Routing Service Group reporting ... 230

S

SAML ... 90, 94
SAML user mapping ... 91
scheduling
full database backup ... 26
full metadata backup ... 75
incremental database backup ... 33
metadata backups ... 74
script example
Elasticsearch ... 203
SDS ... 109
searching for content ... 256
secrets ... 172
security warning banner ... 124
Services fail to consume certificates renewed by Certificate Manager service ... 173
setting
data retention limits ... 40
high availability check interval ... 83
time zones ... 21
sharing content ... 256
sort documents ... 256
specifying identity certificates
CSRs ... 165
generating CSRs ... 165
specifying trust stores
importing certificates ... 168
Stop Agent Trace ... 90
stopping
traffic to datacenter ... 101
supervisor reports
deleted agent groups data ... 231
support ... 261
system
routing service group ... 53
system health check ... 103
system status
monitoring ... 103

T

testing
Historical Reporting with LDAP configuration ... 142
third-party certificates not used ... 178
trace processor ... 88, 89
training ... 257
troubleshoot real-time dashboard ... 245
troubleshooting
access to Historical Reporting web page ... 148
Analytics historical reporting ... 230
Analytics log in issues ... 228
Analytics login access issues using LDAP ... 229

troubleshooting *(continued)*
Analytics web access issues ... 228, 229
Avaya Analytics ... 137
backup files issues ... 162
certificate rotation ... 208
configure Historical Reporting with LDAP ... 142
database ... 139
deleted virtual machine ... 180
Docker server certificate error ... 246
failed service ... 176
historical pod fails to start ... 146
Historical Reporting ... 142
licensing error message ... 178
logs ... 137
mstr srv pod fails to start ... 143, 144
POD issue ... 181
post install script ... 245
recovering a deleted virtual machine ... 180
rotating cluster certificates ... 208
rotating internal cluster certificates ... 208
service failure ... 176
staggered upgrade ... 238
stall during upgrade process ... 176
unknown CA certificate error ... 246
upgrade stalls ... 176
Troubleshooting
Analytics roll back ... 239
Analytics upgrade ... 240
troubleshooting failover ... 164
troubleshooting installation ... 244
troubleshooting replication ... 164
Troubleshooting replication and failover ... 163

U

unauthorized login ... 246
updates to audit rules ... 124
updating
governing rule ... 64
user idle timeouts ... 65
user settings
historical reporting ... 54

V

vCenter user account password ... 253
vCPU ... 108, 109
verbose level ... 88, 89
videos ... 260
View
scheduled incremental backups ... 35
viewing
analytics certificate expiry date ... 153
database backup list ... 24
incremental backup list ... 32
list of data retention limits ... 39
list of imported migration packages ... 78
viewing *(continued)*
list of pods ... 153
Viewing
Email Distribution Service ... 63
viewing AIDE logs in Kibana ... 122
Viewing configuration for an Email Distribution Service ... 63
Viewing LDAP settings ... 69
viewing list
local user from a group ... 62
viewing logs
Elasticsearch ... 202
virtual machine remnants
deleting ... 135
vmware reports consumption of thin provisioned storage . 155

W

watchlist ... 256
web single sign-on ... 92