

Quest® Unified Communications Analytics
8.4.2
Deployment Guide



© 2017 Quest Software Inc.

ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.
Attn: LEGAL Dept.
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our website (www.quest.com) for regional and international office information.


Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at www.quest.com/legal.

Trademarks

Quest and the Quest logo are trademarks and registered trademarks of Quest Software Inc. in the U.S.A. and other countries. For a complete list of Quest Software trademarks, please visit our website at www.quest.com/legal. Microsoft, Active Directory, ActiveSync, Excel, Lync, and Skype are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. All other trademarks, servicemarks, registered trademarks, and registered servicemarks are the property of their respective owners.

Legend

 **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

 **IMPORTANT NOTE, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

Contents

Prerequisites for your installation	8
Overview	8
Supported environments	8
Supported browsers	9
Architecture	9
About the web site	9
About the Data Engine service	10
About the Query Engine service	10
About the Storage Engine service	10
Deployment options	12
Hardware minimum requirements	13
About NAS/SAN support	13
Suggested hardware configurations for different size environments	13
Software prerequisites	15
Prerequisites for the UC Analytics services	15
Web site prerequisites	17
Planning for deployment	17
Determining where to install services	17
What files should be excluded from anti-virus software?	18
What files need to communicate through software firewalls?	18
About UC Analytics configuration	19
Installing UC Analytics	20
Types of installations	20
Upgrading from 8.1 or later	20
Considerations before you install	20
About a pilot installation	21
About a production installation	21
Rules when performing a distributed installation	21
Accounts used during installation	22
Performing a pilot installation	22
Performing a production installation	23
Installing additional Storage Engine services after installation	25
Setting up the web site for HTTPS	26
Configuring UC Analytics	27
About administration settings	27
How do I set a user to be a product administrator?	28
About multi-tenant environments	28
Configuration process overview	28
Verifying the initial target environment	29
About target environments	29
About the authentication credential	30

Adding multiple Active Directory forests	30
Configuring UC Analytics for resource forests	31
Adding a target environment for native Office 365	32
Setting the time period for retaining data	33
Setting the start date for data collection	33
Adding and configuring data sources	33
Adding more than one instance of the same data source	35
Recommendations for collecting from Office 365	36
When would I use the Delete Data option?	36
How often do collections update the data?	36
Viewing the collection job status	37
Copying the job details information	38
Filtering job status results by state or type	38
Filtering job status results by date range	38
Forcing a data source collection to run now	39
Renaming a data source	39
Identifying your internal domains	40
Guidelines when specifying domains	40
Classifying domains for message traffic	40
Granting full access to admin settings	41
Adding a tenant administrator	41
Granting users access to data	41
Differences between aggregate and unrestricted access	42
To grant access to specific collected data	43
Setting display formats in UC Analytics	44
Changing formats for date, time, and digit separators	45
Overriding the time zone offset	45
Adding data sources for Active Directory or Azure Active Directory	46
Adding data sources for different target environments	46
Permissions needed to collect Active Directory data	46
Permissions needed for the Domain Controller data source	47
Permissions needed for the Office 365 user subscription configuration data source ..	47
Adding data sources for Active Directory / Office 365 (hybrid)	47
Creating an Domain Controller data source	47
Creating a data source for Office 365 user subscription configuration	48
Adding data sources for native Office 365	49
Adding data sources, chargeback costs, and thresholds for Exchange and Exchange Online	50
Permissions needed to collect Exchange on-premise or hybrid data	50
Permissions needed for the Exchange configuration data source	51
Permissions needed for the Exchange tracking logs data source	52
Permissions needed for the Exchange mailbox contents data source	52
Permissions needed for the Exchange IIS logs (ActiveSync) data source	53
Permissions needed for the Exchange public folders data source	53
Permissions needed for the Exchange Online hybrid user configuration data source ..	54

Permissions needed for the Exchange Online hybrid mailbox configuration data source	54
Permissions needed for Exchange Online mailbox contents data	54
Permissions needed to collect from native Exchange Online	55
Permissions needed for Exchange Online native user configuration data source	55
Permissions needed for Exchange Online native mailbox configuration	55
Creating a data source for Exchange configuration	56
Why should I specify target mailboxes?	56
Best practices for gathering performance	57
Using dynamic distribution groups to group target users	58
Can I enter the Domain Users group as the target for the data collection?	59
Troubleshooting the Exchange configuration collection	59
Creating a data source for Exchange tracking logs	60
About collecting historical tracking logs	60
Creating a data source for Exchange mailbox contents (EWS)	61
Tips for better performance for mailbox contents collection	62
Do I need both Exchange tracking log and mailbox contents (EWS) collections?	63
Creating a data source for Exchange IIS Logs	64
About the IIS log file locations	64
Creating a data source for Exchange public folders	65
Adding Exchange Online hybrid data sources for hybrid Office 365	66
About AD synchronization methods for hybrid Exchange Online	66
Creating a data source for Exchange Online hybrid user configuration	67
Creating a data source for Exchange Online hybrid mailbox configuration	68
Adding Exchange Online data sources for native Office 365	69
Creating a data source for Exchange Online native user configuration	69
Creating a data source for the Exchange Online native mailbox configuration data source	70
Creating a data source for Exchange Online mailbox contents (EWS) data source	71
Setting chargeback costs for Exchange	72
Setting thresholds for Exchange metrics	72
Omitting words when filtering by subject or body	74
Adding data sources, chargeback costs, and thresholds for Skype for Business/Lync	75
Permissions needed to collect Skype for Business/Lync data	75
Permissions needed for Skype for Business/Lync configuration data source	75
Permissions needed for Skype for Business/Lync CDR data source	76
Permissions needed for Skype for Business/Lync QoE data source	76
Creating a data source for Skype for Business/Lync configuration	76
Creating a data source for Skype for Business/Lync CDR	78
Creating a data source for Skype for Business/Lync QoE	79
Setting call classifications for Skype for Business/Lync	79
About specifying Enterprise Voice call classifications	80
Setting chargeback costs for Skype for Business/Lync	80
Setting thresholds for Skype for Business/Lync metrics	81
About the default Skype for Business/Lync quality metrics	82

Adding new threshold classifications	83
Adding data sources, chargeback, and thresholds for Cisco	86
Permissions needed to collect Cisco data	86
Permissions needed for the Cisco configuration data source	86
Permissions needed for the Cisco CDR logs data source	87
Creating a data source for Cisco configuration	87
Creating a data source for Cisco CDR logs	88
Setting call classifications for Cisco	88
About specifying Cisco call classifications	89
Setting chargeback costs for Cisco	89
Setting thresholds for Cisco metrics	90
Configuring subscription global settings	91
What are insight subscriptions?	91
Configuring settings for subscriptions	91
Prerequisites	91
Making changes to your deployment	93
Options available in the Deployment Manager	93
About the customer feedback program	94
Deploying additional Storage Engines	94
Using the Tools menu for support activities	95
Appendix A:	
Configuring Impersonation	97
Setting impersonation for Exchange 2013 or Exchange 2016	97
Setting impersonation for Office 365	98
Appendix B:	
Configuring the Skype for Business or Lync Server	99
Configuring Lync Server 2010	99
1. Adding the SQL store for monitoring	99
2. Installing the Monitoring role	99
3. Enabling Call Detail Recording (CDR)	100
4. Starting the monitoring services	100
Configuring Lync Server 2013 or Skype for Business 2015	100
Prerequisite	101
1. Associating the store with the Front-End pool	101
2. Updating the Lync or Skype for Business Server	101
3. Enabling and configuring monitoring	102
4. Starting the monitoring services	102
Appendix C:	
Configuring IIS Log Files to capture ActiveSync events	103
Configuring IIS Logging on the Exchange CAS and Mailbox servers	103
Configuring IIS Logging if Exchange is hosted on Windows Server 2008 or later	104
What ActiveSync events are collected and displayed in the insights?	105

Appendix D:	
PowerShell cmdlets used by data sources	107
Exchange configuration data source	107
Exchange public folder data source	108
Skype for Business/Lync configuration data source	108
Exchange Online hybrid and native user configuration data sources	108
Exchange Online hybrid and native mailbox configuration data sources	109
Exchange Online mailbox contents data source	109
Office 365 user subscription configuration data source	109
Appendix E:	
Custom configurations and backup and recovery options	110
Setting up a multi-forest environment with a one-way trust	110
Firewall configuration: ports for data collection	111
Ports used for service-to-service communication in a distributed installation	111
Backing up and restoring your data using scripts	112
Supported scenarios	112
About the backup and restore batch files	113
Step 1: Edit the backup.bat file parameters for your installation	113
Step 2: Run the backup batch file	114
Step 3: Edit the restore batch file	114
Step 4a: Restore the storage data in an existing installation	115
Step 4b: Restore the storage data to a new installation	116
Scheduling the backup batch file to run automatically	116
Performing a manual backup of the storage folder before upgrade	117
Moving your storage location	117
Recommendations for disaster recovery	118
Setting a custom title page for exported or subscription insights	119
Changing the interval time before job status is purged	120
Appendix F:	
Questions and answers about data source configuration	122
Introduction	122
How often do the data collections actually gather data and when do they run?	122
Why did an insight show no data for a 30-day range though I initially set the data source to collect 30 days back?	123
If I collect both Exchange Tracking Logs and Exchange Mailbox Contents, are there duplicate items?	124
If I collect message data only from the Exchange tracking logs, is the message Send Date and delivery time available?	125
What are the differences between the Exchange Mailbox Contents and Exchange Tracking Logs data sources?	125
Why do I get an error when collecting Exchange configuration from multiple Exchange versions?	127
Why do I have to specify domain name when doing a multi-forest collection?	127
About us	129
We are more than just a name	129
Our brand, our vision. Together.	129

Contacting Quest	129
Technical support resources	129

Prerequisites for your installation

- [Overview](#)
- [Supported environments](#)
- [Architecture](#)
- [Deployment options](#)
- [Hardware minimum requirements](#)
- [Software prerequisites](#)
- [Planning for deployment](#)
- [Determining where to install services](#)
- [What files should be excluded from anti-virus software?](#)
- [What files need to communicate through software firewalls?](#)
- [About UC Analytics configuration](#)

Overview

Quest Unified Communications Analytics provides dynamic and up-to-date analytics about your Exchange on-premise messaging environments, about your Skype for Business and Lync on-premise environments, about your Cisco environments, and about Office 365 Exchange Online hybrid or native environments. UC Analytics unlocks the business intelligence in your corporate messaging system to enable better business decisions.

Predefined insights provide a simple way to view Exchange messaging volumes and details between individuals, groups, and external entities, and to see peer-to-peer session and conference information for Skype for Business/Lync or Cisco. You can also view the quality of experience (QoE) information about all Skype for Business or Lync enterprise voice calls, conferences, and peer-to-peer sessions.

You can select predefined insights from the insight library and add them to your personal dashboard.

Once UC Analytics is installed, access the UC Analytics home page, at the following location:

<http://<ServerName>/Analytics/>

Supported environments

You can use UC Analytics to collect data from the following environments:

- Microsoft Exchange 2007 (SP1 and later)
- Microsoft Exchange 2010 (SP1 and later)
- Microsoft Exchange 2013
- Microsoft Exchange 2016
- Exchange Online (Office 365 hybrid with on-premise Exchange)

- Exchange Online (Office 365 native)
- Microsoft Lync 2010
- Microsoft Lync 2013
- Microsoft Skype for Business 2015
- Cisco Unified Communications Manager (CUCM) 9.x to 11.0 and 11.5 with Active Directory integration

Supported browsers

Once you have installed UC Analytics, you can access the web site from anywhere using one of the following browsers:

- Microsoft Edge
- Internet Explorer 10 or later
- Mozilla Firefox
- Google Chrome
- Apple Safari

Since UC Analytics uses Windows Integrated Authentication to connect to the web site, if you are on a different domain than the web site or are using an iPad, you will be prompted for your user name and password when you connect to the web site.

Architecture

The following are the main services that comprise UC Analytics:

- Web site (IIS)
- UC Analytics Data Engine
- UC Analytics Query Engine
- UC Analytics Storage Engine

You can install all the services on a single computer or distribute the services across two or more computers. For a distributed installation, the Data Engine and the Query Engine services can be divided into two different roles:

- collector role
- insights role

About the web site

All user interaction with the product is done through the UC Analytics web site. You configure your target environments and create your data collections through the web site using the Admin Settings. You also create data classifications and grant users access to specific types of data using the Admin Settings.

You view the collected data through insights and use filters to customize insights to include only the data that you want. You can also export insights and set up subscriptions to automatically export and email specific insights on a regular schedule.

About the Data Engine service

The Data Engine contains all the business logic for the product. It determines how UC Analytics stores data and performs the data collections. The Data Engine includes the custom logic used for insights, licensing, tiles; in short, all the logic that defines UC Analytics as a product.

The Data Engine can be split into two roles:

- collector role
- insights role

The collector role is responsible for all the data collections from your environments such as Active Directory, Exchange and Exchange Online, Skype for Business, Lync, and Cisco. It contains the logic associated with retrieving the data.

The insights role is involved with the user interactions with the web site such as launching insights and entering configuration settings.

About the Query Engine service

The Query Engine processes all the data that is stored by the Storage Engine. Processing activities can include finding data, counting data values, and aggregating data.

The Query Engine can be split into two roles:

- collector role
- insights role

The collector role is responsible for storing data to the Storage Engine, storing it efficiently and merging it with existing stored data. The collector role can be a heavy user of resources such as RAM. For larger environments, spikes in collector role activity could affect the responsiveness of insights if both the collector role and insights role are installed on the same computer.

The insights role is responsible for interacting with the Data Engine (insights role) to handle user activity with the web site such as launching insights and setting configuration. The insights role can also be a heavy user of RAM.

About the Storage Engine service

The Storage Engine is the repository for all the data that you collect and for the settings that you specified when you configured the product. You can install multiple storage engine services to distribute the data storage load across several computers. If you install multiple Storage Engines, the data set is divided across the Storage Engines. There is no redundancy since only one copy of a specific piece of data is stored.

Information stored by the Storage Engine

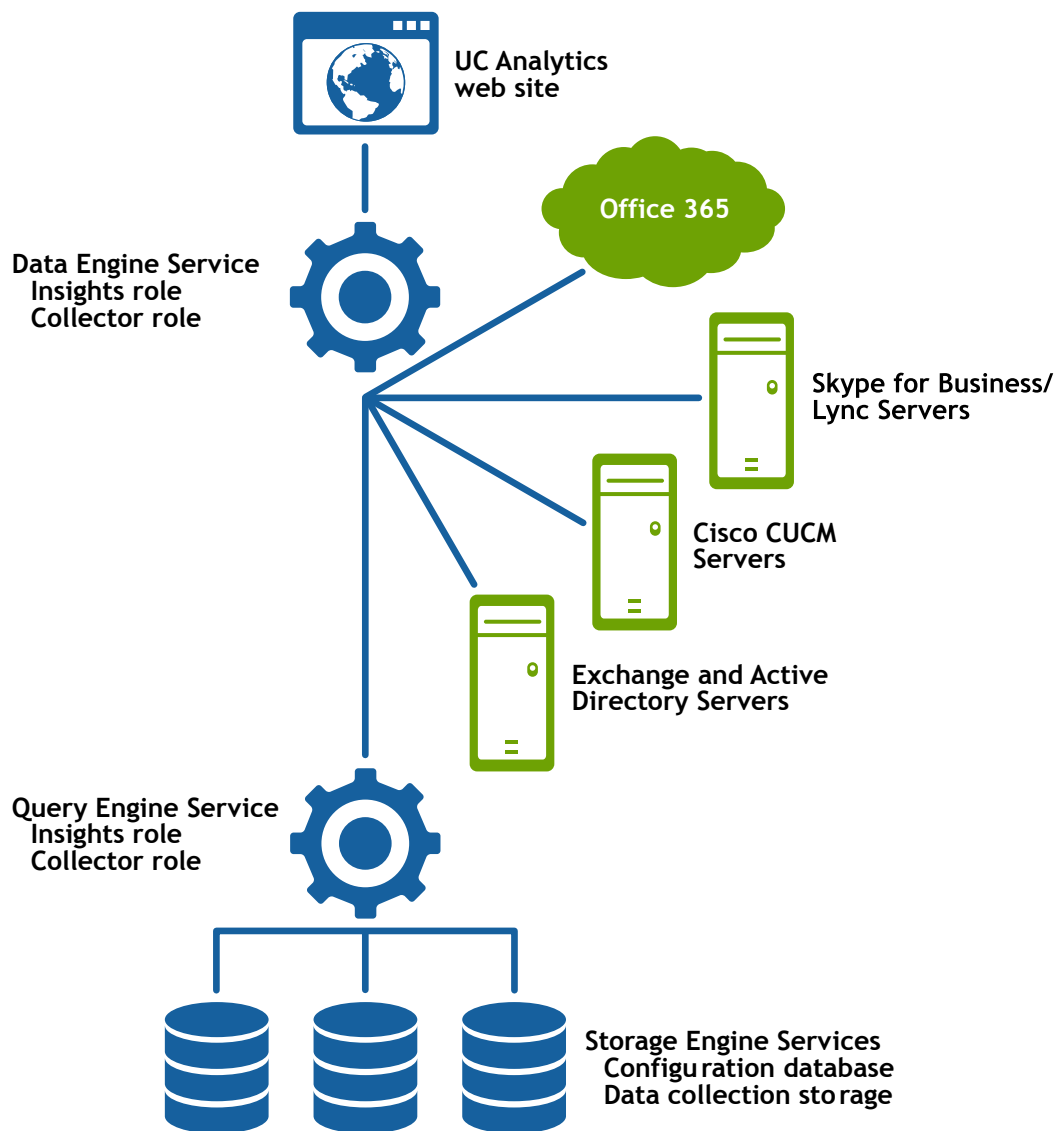
Data is gathered from a variety of sources. You create data source collections to gather data from Active Directory, from Exchange mailboxes through EWS (Exchange Web Services) and from Exchange tracking logs.

You can also collect from Exchange ActiveSync logs, from Exchange Online through PowerShell, from Skype for Business/Lync CDR and QoE databases, and from Cisco Unified Communications Manager server and from Cisco CDR logs.

The Storage Engine service is used to store the collected information:

- User, group, and Data Loss Prevention (DLP) policy information from Active Directory.
- Configuration including organization, server, DAG, database status, database copy, mobile device configuration, personal archive and mailbox statistics from the Exchange server.
- Message traffic information gathered from Exchange message tracking logs.
- Message information gathered through EWS (Exchange Web Services).
- Data Loss Prevention (DLP) policy rule matches collected from Exchange message tracking logs.
- Mobile device information from Exchange IIS ActiveSync logs.
- Public folder information for legacy public folders (Exchange 2007 and Exchange 2010) and new public folders (Exchange 2013 and later).
- Exchange Online mailbox statistics such as mailbox size, permissions, and mobile devices from hybrid (on-premise and Office 365) and from native Office 365 environments.
- Exchange Online user and group configuration, and mailbox content collected from a native Office 365 deployment.
- Office 365 user subscription configuration information including licenses and subscribed services such as Exchange Online, Skype for Business Online, and SharePoint Online.
- Server, pool, and user policy configuration data from the Skype for Business/Lync server
- Peer-to-peer session and conference data collected from the Skype for Business/Lync Call Detail Recording (CDR) database.
- QoE (Quality of Experience) information from the Skype for Business/Lync QoE database.
- Configuration end-user data from the Cisco Unified Communications Manager (Call Manager) server and user data from Active Directory (LDAP).
- Peer-to-peer and conference data collected from the Cisco Call Detail Recording (CDR) log files.

Figure 1. UC Analytics services and relationships to your on-premise environments.



Deployment options

There are two main options for installing UC Analytics:

- Pilot deployment—all services on a single computer, simple configuration.
- Production deployment—services can be distributed across several computers:
 - Web site & Data Engine (insights role)
 - Data Engine (collector role)
 - Query Engine (collector role and insights role)
 - multiple instances of the Storage Engine service.

If you install multiple Storage Engines, the data set is divided across the Storage Engines. There is no redundancy since only one copy of a specific piece of data is stored; no multiple copies of data are stored.

You install the Data Engine (insights role) and the web site on the same server for authentication reasons.

Hardware minimum requirements

If you are installing UC Analytics, the computer must meet the following minimum requirements:

Table 1. Minimum requirements for hardware.

Type	Minimum
Processor	Minimum: Quad-core 64-bit computer.
RAM	Minimum: 16 GB.
Disk	<ul style="list-style-type: none">• 5 GB for the application• X GB where X is the required disk space for collected data <p>For estimates regarding the required disk space for the collected data for different numbers of users for all the platforms, see the <i>UC Analytics Deployment Sizing.xlsx</i> spreadsheet which can be found under Documentation in the autorun.exe.</p>

About NAS/SAN support

Generally, using a NAS/SAN device for the Storage Engine is not supported. UC Analytics might support specific NAS/SAN devices but Quest would require full testing with the devices, or device simulators, to support them. Since each NAS/SAN device, depending on manufacturer, is unique, Quest will work with NAS/SAN vendors to certify or qualify a device but the NAS/SAN vendor must be willing to assist. To engage a NAS/SAN vendor and initiate the qualification process, you would send an email to your NAS/SAN vendor and contact Quest Support.

Suggested hardware configurations for different size environments

Generally you can install UC Analytics on a single computer. For some large installations, you might have a distributed installation on two or more computers. The following examples show the estimated requirements for a single server installation. The estimated requirements assume that you are collecting from all data sources, including Exchange public folders and Exchange IIS Logs (ActiveSync events).

IMPORTANT: These estimates provided here are guidelines only. You might require more resources, depending on your environment and how you configure UC Analytics. Foremost, if you collect email message body, the required RAM and disk space will be greater than the estimates. If you configure many data sources, the required resources can increase.

For distributed installations, you can use the *UC Analytics Deployment Sizing.xlsx* spreadsheet to see different options that can meet your environment's needs. The spreadsheet can be found under Documentation in the autorun.exe.

Table 2. Sample minimum requirements.

Minimum Requirements for	5,000 users	20,000 users	50,000 users	100,000 users
Minimum RAM	16 GB	36 GB	84 GB	144 GB
Disk Space	721 GB	2.9 TB	7.2 TB	14.4 TB

Table 2. Sample minimum requirements.

Minimum Requirements for	5,000 users	20,000 users	50,000 users	100,000 users
Processors	Quad core	Quad core	Eight core	Eight or 16 Core
Data retention	1 year	1 year	1 year	1 year

Example: Small single server installation

This example shows the suggested hardware requirements for a smaller deployment with all the service roles installed on a single computer. The environment from which data is to be collected has the following characteristics:

- # of AD users = 5000
- # of Exchange mailboxes = 5000
- # of Exchange public folders = 5000
- # of Skype/Lync users = 5000
- # of Cisco users = 5000
- Retention period = 365 days

Table 3. Shows the minimum requirements for a small single server installation.

Processor	RAM	Required Disk Space
Quad-core 64-bit	Minimum: 16 GB	721 GB

Example: Medium single server installation

This example shows the suggested hardware requirements for a medium installation with all the service roles installed on a single server. The environment from which data is to be collected has the following characteristics:

- # of AD users = 20000
- # of Exchange mailboxes = 20000
- # of Exchange public folders = 20000
- # of Skype/Lync users = 20000
- # of Cisco users = 20000
- Retention period = 365 days

Table 4. Shows the minimum requirements for a medium single server installation.

Processor	RAM	Required Disk Space
Quad-core 64-bit	36 GB	2.9 TB

Example: Larger single server installation

This example shows the suggested hardware requirements for a larger installation with all the service roles installed on a single server. The environment from which data is to be collected has the following characteristics:

- # of AD users = 50000
- # of Exchange mailboxes = 50000
- # of Exchange public folders = 50000
- # of Skype/Lync users = 50000

- # of Cisco users = 50000
- Retention period = 365 days

Table 5. Shows the minimum requirements for a larger single server installation.

Processor	RAM	Required Disk Space
Eight-core 64-bit	84 GB	7.2 TB

If you install the Storage Engine on a server with two disks, you should specify a directory on the non-operating system drive as the Storage Directory for the Storage Engine during installation.

Software prerequisites

This section lists the prerequisites for the individual UC Analytics components (services and service roles). If you want to install more than one component on a single server, you must ensure all the prerequisites for the components are met.

The server on which UC Analytics is installed must meet following minimum requirements:

Table 6. Software requirements

Server	Minimum Requirements
Operating Systems	<p>One of the following:</p> <ul style="list-style-type: none"> • Windows 7 (64 bit version) or later • Windows Server 2008 R2 (Service Pack 1) or later • Windows Server 2012 and later • Windows Server 2016 <p>Any server on which a UC Analytics service or service role is installed must have a static IP address defined.</p> <p>For a distributed deployment, the Remote Registry Service must be running.</p>

Prerequisites for the UC Analytics services

UC Analytics includes several services that are used to perform different activities. The following prerequisites must be met to install the following services:

- Data Engine service
- Query Engine service
- Storage Engine service.

Table 7. Prerequisites for UC Analytics services.

Service Display Name	Minimum Requirements
UC Analytics Data Engine service	<ul style="list-style-type: none"> • Microsoft .NET Framework 4.6 (full version) or later • To collect Exchange 2007 configuration data, Exchange management tools must be installed on the computer which hosts the UC Analytics Data Engine service. • To collect Office 365 user subscription data, the following software prerequisites must be met: <ul style="list-style-type: none"> ▪ PowerShell 3.0 or later is required. ▪ Microsoft Online Services Sign-in Assistant must be installed. ▪ Windows Azure Active Directory Module for Windows PowerShell version 1.x must be installed. <p>Also, since the Microsoft Azure AD cmdlets use the proxy settings for Internet Explorer, ensure that the Internet Explorer proxy settings for the service account are set correctly.</p>
UC Analytics Query Engine service	<ul style="list-style-type: none"> • 64-bit Java Runtime Environment (JRE) 1.8 or later
UC Analytics Storage Engine service	<p>You can download the Java Runtime Environment (JRE) from the following web site: http://java.com/en/download/manual.jsp.</p>

Web site prerequisites

The Analytics web site requires IIS (Internet Information Services).

Under Internet Information Services (IIS), the World Wide Web Services must be installed. You install World Wide Web Services using Windows Features and enable the required services.

Table 8. World Wide Web Services (IIS) services that must be enabled.

Web server (IIS) services	Services that must be enabled
Application Development Features	<ul style="list-style-type: none">• ASP.NET Enable both ASP.NET 3.5 and ASP.NET 4.5 if available.• .NET Extensibility• ISAPI Extension• ISAPI Filters
Common HTTP Features	Default Document Static Content HTTP Errors
Security	Windows Authentication

Under Web Management Tools, the following service must be enabled:

Table 9. Web Management services that must be enabled.

Web Management tools	Services that must be enabled
IIS 6 Management Compatibility	IIS Metabase and IIS 6 configuration compatibility

For information about setting security for your UC Analytics web site, see [Setting up the web site for HTTPS](#) on page 26.

Planning for deployment

When you are planning your deployment, you must determine the number of servers across which you are deploying UC Analytics. You also must determine how much storage is required for the number of mailboxes from which you are collecting data.

For additional information see the *UC Analytics Deployment Sizing.xlsx* spreadsheet which can be found under Documentation in the autorun.exe.

Determining where to install services

It is **not** recommended that you install any UC Analytics services on the same server where core MessageStats is installed.

If you are installing the UC Analytics services on several computers, consider the following information:

- Data collection performance is better if the Data Engine service is located near the Exchange server and to the Skype/Lync CDR database server.
- Ideally, you would install UC Analytics on a computer that is in the same forest as the Exchange servers from which you are collecting data. However, you can install in a forest that is separate from the forest that contains your Exchange servers. For information about configuring multi-forest environments, see [Adding multiple Active Directory forests](#) on page 30.

- You can install UC Analytics on a virtual machine (VM) but, to obtain optimal performance for a production installation, it is recommended that you install on a physical machine.
- The Storage Engine service nodes automatically synchronize with each other. If you install multiple Storage Engine services, it is better to have them collocated when possible for performance reasons.
- It is best to have all the UC Analytics services located as close as possible. At a minimum they should all be within the same site.
- For security considerations, all services are recommended to be installed within the intranet. For information about configuring the web site for https, see [Setting up the web site for HTTPS](#) on page 26.
- For small to mid-size installations. UC Analytics can be installed on a single computer.
- For large installations, use one of the recommended deployments provided in the installer. For additional information see the *UC Analytics Deployment Sizing.xlsx* spreadsheet which can be found under Documentation in the autorun.exe.

What files should be excluded from anti-virus software?

You should exclude the UC Analytics storage folder from anti-virus scans. By default, the storage directory is located in the following path:

C:\Program Files\Quest\UC Analytics\Storage

If the storage folder is not excluded from anti-virus scans, the anti-virus software will lock files in the storage folder. During normal operation, when your configured UC Analytics data collections attempt to update the stored data, the file locks can cause problems. Also, if anti-virus scans are running against the storage folder, it can create issues when the installer attempts to upgrade your stored data.

What files need to communicate through software firewalls?

Software firewalls, such as Windows Firewall, can prevent the code in the UC Analytics processes from performing necessary network input and output. Such firewalls must be configured to allow the UC Analytics Data Engine to communicate with the Query Engine and to communicate with your Exchange, Active Directory, Cisco, and Skype for Business resources.


The files that can be affected by a software firewall are as follows:

- C:\Program Files\Quest\UC Analytics\Storage Engine\bin\prunsvr.exe
- C:\Program Files\Quest\UC Analytics\Query Engine\bin\prunsvr.exe
- C:\Program Files\Quest\UC Analytics\Data Engine\UC.Analytics.Insights.DataEngine.Service.exe
- C:\Program Files\Quest\UC Analytics\Data Engine\UC.Analytics.Insights.DataEngine.BulkDataExport.exe
- C:\Program Files\Quest\UC Analytics\DeploymentManager.exe

About UC Analytics configuration

After you install UC Analytics, you must configure Admin Settings. You can access Admin Settings from the UC Analytics home page located at:

`http://<ServerName>/Analytics.`

Click the gear icon  on the home page side bar to access Admin Settings.

For more information, see [Configuration process overview](#) on page 28.

Installing UC Analytics

- [Types of installations](#)
- [Considerations before you install](#)
- [Accounts used during installation](#)
- [Performing a pilot installation](#)
- [Performing a production installation](#)
- [Installing additional Storage Engine services after installation](#)
- [Setting up the web site for HTTPS](#)

Types of installations

When you install UC Analytics, you can perform one of two types of installation:

- a pilot installation (single computer)
- a production installation (single computer or distributed)

For performance reasons, it is recommended that you do not install UC Analytics on the same computer as MessageStats.

For information about the number of mailboxes from which you can collect data see [Hardware minimum requirements](#) on page 13.

If you have a native Office 365 environment with no on-premise users, you can install UC Analytics in a workgroup.

Upgrading from 8.1 or later

When you are upgrading from UCCS - Analytics 8.1 or later, perform an in-place upgrade to ensure you maintain your collected data. Do **not** uninstall the previous version. If you uninstall the previous version, you will lose all the collected data.

i | **NOTE:** If you previously customized a data collection by modifying a .config file (usually with the help of Quest Support), after you upgrade to a new version, check the same .config file and reapply all the changes.

Considerations before you install

The following limitations must be considered before you install UC Analytics:

- Do not install UC Analytics on the same server on which you have installed MessageStats.
- For evaluation scenarios, you can install UC Analytics on a virtual machine (VM). However, in a production environment, a physical computer provides better performance.

About a pilot installation

In a pilot installation, all the UC Analytics services are installed on a single computer. Though pilot installations are typically used in test installations, if you have a computer with enough resources you could also use a pilot installation for a production environment.

If you have a native Office 365 environment with no on-premise users, you can install UC Analytics in a workgroup.

About a production installation

In a production installation, you can distribute the UC Analytics services across several servers. Additionally, the Data Engine service and the Query Engine service can be distributed in two separate roles, the insights role and the collector role. In a UC Analytics installation you can distribute the following services as follows:

- Web site & Data Engine (insights role)
- Data Engine (collector role)
- Query Engine (collector role)
- Query Engine (insights role)
- multiple instances of the Storage Engine service.

For example, you might install the web site and the Data Engine (insights) and Query Engine (Insights) on one server and the Data Engine (collector) and Query Engine (collector) on a second computer. You might install instances of the Storage Engine service on both the servers and an additional Storage Engine on a third server.

Table 10. Example of a distributed installation

Server1	Server2	Server3
Web Site & Data Engine (insights)	Data Engine (collector)	Storage Engine
Query Engine (insights)	Query Engine (collector)	
Storage Engine	Storage Engine	

i | **IMPORTANT:** Due to the volume of communication between the Data Engine and Query Engine, it is recommended that these services be installed on the same server.

For information about storage requirements, see [Hardware minimum requirements](#) on page 13.

Also, you must ensure that the computer on which UC Analytics is to be installed meets the minimum software requirements. For information about the minimum requirements, see [Software prerequisites](#) on page 15.

Rules when performing a distributed installation

Generally you select one of the recommended installations and specify the servers on which services are to be installed. For larger, more complex environments, you can perform a production installation and specify the location for each service or service role.

If you are performing a production installation, be aware of the following rules:

- You install multiple instances of the Storage Engine but only one web site and only one instance of each role for the Data Engine and the Query Engine.
- For authentication reasons, the web site and the Data Engine (insights role) are always installed on the same computer.

- If you install the Data Engine (collector role) and Data Engine (insights role) on the same computer, you must also install the Query Engine (collector role) and Query Engine (insights role) on the same computer. UC Analytics does not support having the Data Engine roles installed on single computer while the Query Engine roles are installed on separate (multiple) computers.
- If you install the Data Engine (collector role) and Data Engine (insights role) on two different computers, you can install the Query Engine (collector role) and Query Engine (insights role) on two different computers.

Accounts used during installation

The account that you use to install UC Analytics must be a local administrator on the computer. If you are distributing the services, the account must have local administrator rights on each computer on which you are installing.

During installation you can specify an account that is used to run the Data Engine service. This account must be a domain account and must have local Admin rights on the computer on which the Data Engine service is installed.

By default, the account that is used to install UC Analytics is set as a product administrator, which means that the account can access the Admin Settings which are used to configure the product.

Performing a pilot installation

A pilot installation installs all services on a single computer. All servers on which UC Analytics services are installed must have a static IP address defined.


To perform a pilot installation of UC Analytics

- 1 Copy the **UC Analytics.exe** file to the computer on which you want to install and double-click the file.
- OR -
Double-click the **autorun.exe** file and select the **Install** tab.
- 2 Beside the heading for the UC Analytics 8.4.2 Installer, click **Install**.
- 3 Select **Pilot Install**.
- 4 Enter information about the environment from which you want to gather data such as the number of Active Directory users, Exchange mailboxes, Skype for Business/Lync users, and/or Cisco users and the time that you want to retain the collected data.

The UC Analytics installer verifies the disk space and physical memory (RAM) on the selected server and provides approximate estimated recommendations for your hardware configuration.
- 5 Enter the credentials that will be used to run the Data Engine service and click **Next**.

The account must be a domain account (not a local computer account) and must have local Admin rights to the computer on which the Data Engine service is being installed.
- 6 Verify that the software prerequisites are met and click **Next**.
- 7 Specify the users who will be product administrators and have access to Admin Settings and click **Next**.

By default, the account used to install UC Analytics is added to the accounts that can access the Admin Settings. You can add additional users by rerunning the installer and selecting to add additional product administrators.
- 8 Select the country in which you are installing the product.
 - a For some countries, you will be asked whether you want to participate in the customer feedback program. For more information, see [About the customer feedback program](#) on page 94.

- 9 Accept the license agreement and install the product.
- 10 To access Admin Settings, open the UC Analytics web site:
http://<ServerName>/Analytics
- 11 Click the gear icon  on the home page side bar.

Performing a production installation

When you perform a production installation, you can install different services on separate computers. You can run the installer centrally and remotely deploy the different services to different computers.

When planning a distributed installation, be aware of the following factors:

- All servers on which the UC Analytics services are installed must have a static IP address defined.
- The port numbers that are specified for the various services are the user port numbers (1024-49151). Any port number that you specify should not already be in use. If you have a firewall, ensure that the ports that you specify have access.
- To improve performance and distribute the data storage load, you can install multiple instances of the Storage Engine service. Ensure that you specify a storage directory in which you have enough available disk space; otherwise the installer will use the default installation location.

For information about rules for distributed installations, see [Rules when performing a distributed installation](#) on page 21.

To perform a production installation of UC Analytics

- 1 Copy the **UC Analytics.exe** file to the computer on which you want to install and double-click the file.
- OR -
Double-click the **autorun.exe** file and select the **Install** tab.
- 2 Beside the heading for the UC Analytics 8.4.2 Installer, click **Install**.
- 3 Select **Production Install**.
- 4 Enter information about the environment from which you want to gather data such as the number of Active Directory users, Exchange mailboxes, Skype for Business/Lync users, and/or Cisco users and the time that you want to retain the collected data and click **Next**.
- 5 Select either **Recommended Deployment** or **Advanced Deployment**.

Recommended Deployment

If you selected Recommended Deployment, the installer displays the recommended deployment options for one or more servers. It shows the recommended memory and storage for each configuration.

Table 11. Deployment options based on estimated RAM required.

Estimated RAM required	Number of options displayed
Less than 64 GB RAM	One option is displayed.
Greater than 64 GB RAM but less than 128 GB RAM	Two options are displayed.
Greater than 128 GB RAM	Three options are displayed

- 1 Select the option that you want.
The Configuration Deployment page shows all the service roles to be installed and shows a placeholder server (Server1, Server2, and so on) with the appropriate roles selected.
- 2 Select a server and click **Edit Properties**.
- 3 Enter the actual server name for the server on which the service roles will be installed.

- a If you want to install the service application binaries (program modules) on a drive other than C or to a different directory, enter the path for the location in the **Application Binaries Directory Path** field.
- 4 Repeat [Step 2](#) and [Step 3](#) for each server on which service roles will be installed.

Advanced Deployment

If you selected Advanced Deployment, the installer displays the Configuration Deployment page with all the service roles to be installed and lists the current server with all roles selected.

- 1 To remove any roles from the current server, click **Edit Properties** and remove the roles that you do not want installed.
- 2 To add additional servers and specify the roles that should be installed on each server, use the following steps:
 - a Click **Add Server**.
 - b Specify the server name and select the check boxes for the server roles that are to be installed.
 - c If you want to install the service application binaries (program modules) on a drive other than C or to a different directory, enter the path for the location in the **Application Binaries Directory Path** field.
 - d Repeat [Step a](#) and [Step b](#) for each server on which UC Analytics roles are to be installed.

For most installations, you can use the default values defined for each server role. However, you can change the default values if necessary.

- 6 To change the default values for a server role, click **Edit Properties**.
- 7 Click **Advanced**, enter any changes to the properties for the selected server role and click **OK**.

Table 12. Data Engine service properties.

Query Port Number	If the port is already used by another application, change the port number. This port is used by the web site to access the Data Engine service for query purposes.
Configuration Port Number	If the port is already used by another application, change the port number. This port is used by the web site to access the Data Engine for configuration purposes.

Table 13. Query Engine service properties.

Physical Memory (MB) Calculate automatically	Ensure the check box is selected to calculate whether sufficient memory is available. If you clear the check box, the value you enter in text box is assigned. NOTE: Physical memory is automatically assigned during installation. After you successfully install, if you display the properties dialog, the currently installed physical memory is displayed in the text box.
Query Port Number	If the port is already used by another application, change the port number. This port is used by the Data Engine service to access the Query Engine service.

Table 14. Storage Engine service properties.


Physical Memory (MB) Calculate automatically	Ensure the check box is selected to calculate whether sufficient memory is available. If you clear the check box, the value you enter in text box is assigned. NOTE: Physical memory is automatically assigned during installation. After you successfully install, if you display the properties dialog, the currently installed physical memory is displayed in the text box.
Query Port Number	If the port is already used by another application, change the port number. This port is used by the Query Engine service to access the Storage Engine service.

- 8 In the Storage Directory Path field, specify the directory path in which all the collected data is stored and click **Next**.

The UC Analytics installer verifies the hardware prerequisite for required disk space against actual disk space and physical memory (RAM) on the selected servers for the roles you have specified and provides recommendations for your hardware configuration.
- 9 Enter the credentials that will be used to run the Data Engine service and click **Next**.

The account must be a domain account (not a local computer account) and must have local Admin rights to the computer on which the Data Engine service is being installed.
- 10 Verify that the software prerequisites are met.
- 11 Specify the users who will be product administrators and have access the Admin Settings and click **Next**.

By default, the account used to install UC Analytics is added to the accounts that can access the Admin Settings. By rerunning the installer and selecting **Change Product Administrators**, you can add or remove product administrators.
- 12 Select the country in which you are installing the product.
 - a For some countries, you will be asked whether you want to participate in the customer feedback program. For more information, see [About the customer feedback program](#) on page 94.
- 13 Accept the license agreement and install the product.
- 14 To access Admin Settings, open the UC Analytics web site:

`http://<ServerName>/Analytics`
- 15 To begin product configuration, click the gear icon  on the home page side bar. For more information, see [Configuring UC Analytics](#) on page 27.

Installing additional Storage Engine services after installation

At some later date, you might decide you need to install additional Storage Engine services to distribute your stored data. For information about adding more Storage Engines to an existing installation, see [Deploying additional Storage Engines](#) on page 94.

Setting up the web site for HTTPS

Generally it is recommended that you set security for your UC Analytics web site. The steps for configuring Secure Sockets Layer (SSL) or Transport Layer Security (TLS) for a site generally include the following:

- 1 Get an appropriate certificate.
- 2 Create an HTTPS binding on a site.
- 3 Optionally, configure TLS/SSL options by making TLS/SSL a requirement.
- 4 Test by making a request to the site.

First you must obtain a web server certificate. For information about how to install an Internet Server Certificate (IIS 7.0) see the following article:

[http://technet.microsoft.com/en-us/library/cc731977\(v=ws.10\)](http://technet.microsoft.com/en-us/library/cc731977(v=ws.10))

After the certificate is obtained, you must install it. Now you can configure the https binding and the TLS/SSL options. Use the procedure that follows.

To add the https binding and configure the TLS/SSL options

- 1 Open IIS Manager.
- 2 In the Connections pane, expand the **Sites** node in the tree and select the **Default Web Site**.
- 3 In the Actions pane, click **Bindings**.
- 4 In the Site Bindings dialog box, click **Add**.
- 5 In the Add Site Binding dialog box, add the binding information including the certificate and click **OK**.
- 6 In Features View at the UC Analytics application level, double-click **TLS/SSL Settings**.
- 7 On the TSL/SSL Settings page, select **Require TLS/SSL**.
- 8 In the Client certificates area, select **Ignore**.
- 9 In the Actions pane, click **Apply**.
- 10 To open the web site, enter the following address in the browser:
`https://<ServerName>/Analytics/`


Configuring UC Analytics

- [About administration settings](#)
- [How do I set a user to be a product administrator?](#)
- [Configuration process overview](#)
- [Verifying the initial target environment](#)
- [Adding multiple Active Directory forests](#)
- [Configuring UC Analytics for resource forests](#)
- [Adding a target environment for native Office 365](#)
- [Setting the time period for retaining data](#)
- [Setting the start date for data collection](#)
- [Adding and configuring data sources](#)
- [How often do collections update the data?](#)
- [Viewing the collection job status](#)
- [Forcing a data source collection to run now](#)
- [Renaming a data source](#)
- [Identifying your internal domains](#)
- [Classifying domains for message traffic](#)
- [Granting full access to admin settings](#)
- [Granting users access to data](#)
- [Setting display formats in UC Analytics](#)

About administration settings

If you are a product administrator, you can configure UC Analytics at the following location:

<http://<ServerName>/Analytics/>

Click the gear icon  on the home page side bar to access Admin Settings.

Before you begin to use UC Analytics to gather and analyze information, you must specify settings that determine what data is collected (data sources) and when it is collected. Depending on the data sources that you configure, UC Analytics can collect different types of data.

You configure classifications to identify the domains that are internal to your organization, to configure call classifications for Skype for Business/Lync and Cisco, and to set thresholds for quality metrics for the different platforms.

In the Admin Settings, you also configure security to control which users can see data in which insights. For example, you can specify security settings which determine whether specific users can see detailed or aggregate data for the Exchange messages and Skype for Business/Lync data.

How do I set a user to be a product administrator?

When you install UC Analytics, you can add users as product administrators during installation. A product administrator has unrestricted access to all the configuration settings under the Admin Settings for all tenants (target environments).

After installation, to add additional administrators you can use the UC Analytics Deployment Manager from the Start menu.

To access the Deployment Manager

- 1 Click Start and select **Programs | Quest UC Analytics | UC Analytics**.
- 2 Select the **Change Product Administrators** option.


About multi-tenant environments

For implementations in which you have multiple tenants, you can add separate environments for each tenant. In this case, you could add tenant administration rights for a specific environment to an account. The account would be the tenant administrator only for the tenant environment.

For information about configuring an account with access to the admin settings for a specific environment (tenant), see [Adding a tenant administrator](#) on page 41.

Configuration process overview

If your user account is set as a product administrator or with access to Tenant Configuration (tenant administrator), you can configure UC Analytics using the Admin Settings. The following steps provide a sample workflow for an initial configuration and identify the tile (in parentheses) in which you enter the configuration.

- 1 Click the gear icon  on the home page side bar to access Admin Settings.
- 2 Review the target environment used to create the initial connection. (Target Environments)
- 3 Add any additional target environments such as another Active Directory forest or a native Office 365 deployment. (Target Environments)
- 4 Add and configure the data sources used to create the data collections that gather data. (Data Collection).

For information about configuring each type of data source, see the following sections:

- [Adding data sources for Active Directory or Azure Active Directory](#) on page 46
- [Adding data sources, chargeback costs, and thresholds for Exchange and Exchange Online](#) on page 50
- [Adding data sources, chargeback costs, and thresholds for Skype for Business/Lync](#) on page 75
- [Adding data sources, chargeback, and thresholds for Cisco](#) on page 86

- 5 Set a data aging time period for the data you collect (Data Collection).
- 6 Set the initial data collection start date (Data Collection).
- 7 Identify the domains that are internal to your environment (Classifications | Domain Classifications).
This information is used in insights to identify internal and external message traffic and Skype for Business/Lync activities.
- 8 Specify classifications that map to specific domains (Classifications | Domain Classifications).

The classifications are used to group the Exchange messaging, Skype for Business/Lync session and conference data, and Cisco session and conference in the insights.

- 9 Set security to set the levels of access to the Exchange, Skype for Business/Lync, and Cisco data that each user has when viewing insights (Security).

Other configuration tasks can include:

- Specifying words that should be omitted when filtering for subject keywords (Queries).
- Entering the call classifications for the different Skype for Business/Lync or Cisco call types (Classifications | Call Classifications).
- Entering the cost values to be used for chargeback insights (Chargeback).
- Configuring the global settings that are used by user-created subscriptions for insights (Subscriptions).

For more information about subscriptions, see the section titled “Automatically creating and sending insights” in the *UC Analytics User Guide*.

Verifying the initial target environment


When you install UC Analytics, the installer creates an initial target environment. By default, the target environment is the Active Directory forest (or workgroup) in which UC Analytics is installed. The target environment is used to collect data and to provide users with access to the UC Analytics web site.

Check the initial target environment to verify that the default credentials being used for authentication have the required permissions. For more information, see [About the authentication credential](#) on page 30.

To provide access for users or to collect data from outside the initial target environment, you can add a new target environment. You can add target environments to connect with other Active Directory forests or with a native Office 365 deployment:

- If you have a hybrid Office 365 environment with users in both Office 365 and in Active Directory, you would add another Active Directory forest / Office 365 (hybrid) environment.
- If you have a native Office 365 environment, you would add an Office 365 (native) target environment and allow your Office 365 users to be authenticated.

To review the initial target environment

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Target Environments**.
By default, an environment is configured for the Active Directory forest in which you installed UC Analytics.
- 3 Click the tile for the forest.
- 4 Review the information. The **Allow users from this target environment to log in to UC Analytics** check box is selected. This indicates that the environment is used for authenticating users.
- 5 Specify an Authentication Credential, if necessary.

About target environments

If you have configured additional forests or an Office 365 target environment (for a native Office 365 deployment), when you click Add Users to grant access to a specific type of data, you must select the target environment (an Active Directory forest or an Office 365 site) for the users.

For the users who are being granted access:

- You can grant aggregate access and/or unrestricted access.

- You can grant access to all users in all target environments or grant access to specific users in a specific target environment.

When granting access to specific users, you can enter either an individual user or a distribution group. If you selected Office 365 as the target environment, it is not recommended that you enter a dynamic distribution group or a distribution group with a large number of members due to performance issues.

About the authentication credential

If you have a single forest, you typically do not need to specify an Authentication Credential in the default target environment. By default, the credentials specified for the Data Engine service are used to verify user permissions.

However, in secure environments, you might need to restrict rights for the UC Analytics service account. You could install UC Analytics using a local computer account or using a domain account with restricted rights.

In this case, you specify an Authentication Credential in the default environment to allow users to access UC Analytics:

- The credential must have read rights to the Active Directory forest specified in the target environment.
- The credential must also have sufficient rights to browse users and groups and to resolve group memberships for all users and groups from this environment that are specified in the Security settings. For information about the Security settings, see [Granting users access to data](#) on page 41.

If you do not provide an authentication account or if the account has insufficient rights, the security settings to access data are not applied when user attempts to log in to UC Analytics. As a result, the user might be unable to log into UC Analytics or unable to access certain insights.


Adding multiple Active Directory forests

When you install UC Analytics, an initial target environment is created for the Active Directory forest in which you installed.

You can add target environments for other Active Directory forests/Office 365 hybrid environments as required. If you add target environments for additional Active Directory forests, there must be at two-way or a one-way trust between the forest from which you collect data and the forest in which UC Analytics is installed.

For detailed instructions about configuring target environments for a multi-forest environment with a one-way trust, see [Setting up a multi-forest environment with a one-way trust](#) on page 110.

To add an target environment for an additional Active Directory forest or Office 365 hybrid environment

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Target Environments**.
- 3 Click **+** beside Target Environments.
- 4 Select **Active Directory Forest / Office 365 (hybrid)**.
- 5 Under Name, click **Active Directory Forest / Office 365 (hybrid)**.
- 6 Enter a name to identify the forest and click **Apply**.
- 7 Under Forest Name, click the **<None specified>** field.
- 8 Enter a domain name associated with the forest
- 9 If necessary, enter the credential that is used to access Active Directory and resolve the forest.

10 Click **Resolve Forest Name**.

The name of the forest in which the domain resides is displayed.

11 Click **Apply**.

12 If this environment is to be used for authenticating users, select the check box beside **Allow users from this target environment to log in to UC Analytics**.

13 Specify the credential if the credential is to be used for authentication.

14 Click **Add**.

Configuring UC Analytics for resource forests

If you have resource forests, typically you have your active user accounts in one forest and the Exchange resources (such as mailboxes) in a different forest. In other cases, you may have a mix of user accounts and Exchange resources in a single forest.

You must configure UC Analytics to collect Active Directory information from the account forests and Exchange information from the resource forests.

For each forest that contains Active Directory accounts and/or Exchange resources, you must:

- add a target environment for each forest that hosts Exchange (pure Exchange resource forest)
- add a target environment for each forest that hosts the Active Directory accounts that access Exchange (pure account forest).
- add a target environment for each forest that hosts both the Exchange resources and the Active Directory accounts used to access Exchange (blended forest).

For each forest target environment, you must add a Domain Controller data source.

For each forest target environment that hosts Exchange resources (both pure Exchange resource forests and blended forests), you must add Exchange data sources (such as the Exchange Configuration and IIS Log Files data sources).

Resource forest configuration process

In Admin Settings, the configuration process is as follows:

1 In Target Environments, add a target environment for each forest that:

- hosts the Active Directory users that are used to access Exchange.
- hosts Exchange resources.

2 In Data Collection, add a Domain Controller data source in each of the target environments.

The data source collects the Active Directory objects of each forest. The Domain Controller data source is also needed in a resource forest to collect the security groups and Exchange distribution groups.

3 In Data Collection, add an Exchange Configuration data source in each target environment of a forest that hosts Exchange resources. Configure as usual with the following extra steps:

a In each Exchange Configuration data source, select the **Search additional forests for user accounts** check box.

The additional Account Forest LDAP Connection Parameters are displayed.

b Specify whether the domain controller for the account forest should be automatically discovered or enter a specific domain controller.

c Enter the additional LDAP connection parameters to the account forest that contains the master accounts for the Exchange-linked mailboxes.

The additional LDAP credentials are needed to resolve the master account SID of linked mailboxes to the distinguished name of the master account in the appropriate account forest.

Collecting properties from active and disabled user accounts


In a resource forest configuration, both the active (master) user account and the disabled (stand-in) user account can have values specified for the user properties. UC Analytics collects these properties as follows:

- If a property is multi-valued, such as email addresses, UC Analytics collects the values from both the master and stand-in accounts.
- If the property is single-valued, such as department, UC Analytics collects the values from both the master and stand-in accounts. When the value is present in both accounts, the stand-in account value is used with the following exceptions which are always collected from the master account:
 - Name
 - First Name
 - Initials
 - Last Name
 - SAM Account Name
 - OU
 - Object GUID
 - Object Type

Adding a target environment for native Office 365

If you have a native Office 365 environment, without on-premise Exchange or Active Directory, you can add a Office 365 target environment to connect directly to your native Office 365 environment.

To add a target environment for native Office 365

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Target Environments**.
- 3 Click **+** beside Target Environments.
- 4 Select **Office 365 (native)**.
- 5 Under Name, click **Office 365 (native)**.
- 6 Enter a descriptive name for the environment and click **Apply**.
- 7 To use this environment for authenticating users, select the check box beside **Allow users from this target environment to log in to UC Analytics**.

This setting only controls whether the users from this target environment can log in to UC Analytics. You control the types of data that users can view in the insights through the Admin Settings | Security section. For information, see [Granting users access to data](#) on page 41.

- 8 Click **Set credential** and enter the Office 365 credential to be used to resolve user security settings for users that log in with Office 365 credentials.
- 9 Click **Set**.
- 10 Click **Add**.

Setting the time period for retaining data

On the Data Collection page, you can set the amount of time that data is retained by the UC Analytics Storage Engine. This feature allows you to automatically age data from the Storage Engine database to manage the amount of storage required.

The data aging job runs automatically at midnight (local time).

If you have the data retention period set for a longer period such as 365 days and you reduce the number of days for retention, you cannot later recollect the older data that was not retained. If you increase the number of days, the change only affects new message or session activity.

To set the data retention time period

- 1 Click **Data Collection**.
- 2 Enter the number of days that you want data to be retained in the Data Retention Length (days) field.

Setting the start date for data collection

On the Data Collection page, you set the date that you want data collection to start. By default, the initial data collection is set to start 30 days back for on-premise data sources and one day for Office 365 (Exchange Online) data sources.

For on-premise data sources, if you have stored data that goes more than 30 days back, you might want to change the start date for data collection. For example, typically Exchange tracking logs are retained for 30 days. If you normally archive your Exchange tracking logs and have logs that go back beyond 30 days, you could set the data collection start date to reflect the actual dates for the data that you have stored.

To set the data collection start date

- 1 Click **Data Collection**.
- 2 Enter the start date for which data is to be collected in the Data Collection Start Date field.

You can only change the start date if you have not added any data sources yet.

Adding and configuring data sources


You can configure the data sources to create the data collections that collect information from different environments. The workflow to add data sources and set up data collection is similar for each target environment.

Data source name	Types of information collected
Active Directory data sources for Active Directory Forest/Office 365 Hybrid target environment	
Domain Controller	User, group, and contact data from Active Directory. See Creating an Domain Controller data source on page 47.
Office 365 User Subscription Configuration	User subscription information including licenses and subscribed services such as Exchange Online, Skype for Business Online, SharePoint Online, using remote PowerShell. See Creating a data source for Office 365 user subscription configuration on page 48

Data source name	Types of information collected
Azure Active Directory data sources for native Office 365 target environment	
Office 365 User Subscription Configuration	User subscription information including licenses and subscribed services such as Exchange Online, Skype for Business Online, SharePoint Online, using remote PowerShell. See Creating a data source for Office 365 user subscription configuration on page 48
Exchange data sources for Active Directory Forest/Office 365 Hybrid target environment	
Exchange Configuration	Exchange configuration from the Exchange server including organization, server, DAG, mailbox, database status, database copy, mobile device configuration, personal archive and mailbox statistics. See Creating a data source for Exchange configuration on page 56.
Exchange Tracking Logs	Exchange message traffic and DLP rule matches from your Exchange message tracking logs. See Creating a data source for Exchange tracking logs on page 60.
Exchange Mailbox Contents	Exchange message data through Exchange Web Services (EWS) from your target mailboxes. See Creating a data source for Exchange mailbox contents (EWS) on page 61.
Exchange IIS Logs	Mobile device events and email statistics from the ActiveSync IIS log files on Exchange Client Access Services (CAS). See Creating a data source for Exchange IIS Logs on page 64.
Exchange Public Folders	Statistics for legacy public folders (Exchange 2007 and Exchange 2010) and for new public folders (Exchange 2013 and Exchange 2016). See Creating a data source for Exchange public folders on page 65.
Exchange Online Hybrid User Configuration	Detailed information for Exchange Online users and distribution groups in a hybrid environment. See Creating a data source for Exchange Online hybrid user configuration on page 67.
Exchange Online Hybrid Mailbox Configuration	Exchange Online mailbox statistics, permissions, and mobile device data from a hybrid (Office 365 and on-premise Exchange) environment using PowerShell and LDAP. See Creating a data source for Exchange Online hybrid mailbox configuration on page 68.
Exchange Online Mailbox Contents	Information about email traffic from Exchange Online user mailboxes using Exchange Web Services (EWS). You can collect from some or all mailboxes. See Creating a data source for Exchange Online mailbox contents (EWS) data source on page 71.
Exchange Online data sources for native Office 365 target environment	
Exchange Online Native User Configuration	Exchange Online user and group data from a native Office 365 environment using PowerShell. See Creating a data source for Exchange Online native user configuration on page 69.
Exchange Online Native Mailbox Configuration	Exchange Online mailbox configuration including statistics, permissions and mobile devices from native Office 365 using remote PowerShell. See Creating a data source for the Exchange Online native mailbox configuration data source on page 70.
Exchange Online Mailbox Contents	Information about email traffic from Exchange Online user mailboxes using Exchange Web Services (EWS). You can collect from some or all mailboxes. See Creating a data source for Exchange Online mailbox contents (EWS) data source on page 71.
Skype for Business/Lync data sources for on-premise Active Directory forest target environment	
Skype for Business/Lync Configuration	Server, pool, and user policy configuration data directly from the Skype for Business/Lync server. See Creating a data source for Skype for Business/Lync configuration on page 76.

Data source name	Types of information collected
Skype for Business/Lync CDR Database	Skype for Business/Lync peer-to-peer session and conference data from the CDR database. See Creating a data source for Skype for Business/Lync CDR on page 78.
Skype for Business/Lync QoE Database	Skype for Business/Lync Quality of Experience (QoE) data from the QoE database. Creating a data source for Skype for Business/Lync QoE on page 79.
Cisco data sources for on-premise Active Directory forest target environment	
Cisco Configuration	Cisco server and end-user data the Cisco Unified Communications Manager server and synchronized user data from Active Directory. See Creating a data source for Cisco configuration on page 87.
Cisco CDR Logs	Cisco peer-to-peer call, ad hoc conference call, and meet-me conference call information from Cisco CDR log files. See Creating a data source for Cisco CDR logs on page 88.

To set up data collection for your environments

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click + beside the name of the target environment.
- 4 Select the data sources for Active Directory (Domain Controller), Exchange, Skype for Business/Lync, Cisco, and Office 365 from which you want to collect data
Tiles for each data source are displayed.
- 5 Click the data source tile to enter the configuration information for the data collection.
- 6 For specific instructions about how to create data sources for each supported platform, see the following sections:
 - [Adding data sources for Active Directory or Azure Active Directory](#) on page 46
 - [Adding data sources, chargeback costs, and thresholds for Exchange and Exchange Online](#) on page 50
 - [Adding data sources, chargeback costs, and thresholds for Skype for Business/Lync](#) on page 75
 - [Adding data sources, chargeback, and thresholds for Cisco](#) on page 86

Adding more than one instance of the same data source

You must add data sources for each target environment that you have created. So if you have two Active Directory forests added as target environments, you would add the same data sources for each target environment.

For performance reasons, you can add multiple data sources of the same type to distribute the collection load or for different platform versions. For example, you would add an Exchange configuration data source for your Exchange 2007 servers and an Exchange configuration data source for your Exchange 2010 or Exchange 2013/2016 servers.

You can use the rename option to rename each data source so that it reflects the information that is being collected. For information about renaming data sources, see [Renaming a data source](#) on page 39.

Recommendations for collecting from Office 365

The data sources that collect from Office 365, including Exchange Online (hybrid or native), use remote PowerShell to collect the data from Office 365. You must specify the account is used to connect to Exchange Online when you configure the data source.

It is recommended that you specify a different account for the PowerShell credentials that are used for each Exchange Online or Office 365 data source that you configure.

When would I use the Delete Data option?

On the Data Collections page, there is a Delete Data button. Typically you only use this option if you want to delete all the data you have collected. To do this, you must first delete all the data sources that you have configured. By deleting all the data sources and then deleting all the collected data, you can essentially start over again.

If you only want to remove old data, adjust the time period for data retention.

How often do collections update the data?

By default, each data collection has a default schedule. When you configure a data collection, you have the option of changing the schedule for that collection.

After a data collection is created and is running on a regular basis, you can change the collection. For example, you could select additional types of data to be collected. For some types of data, the data is collected only once a day so the new data will not be available in insights until the following day.

Some of the data collections update each time that the collection job is run. For these data collections, only the data that changed from the previous job run is gathered. The data collections that update the data that changed since the last run are:

- Exchange Tracking Logs (message traffic, DLP matches)
- Exchange Mailbox Contents (message information using EWS)
- Exchange ActiveSync IIS Logs (mobile device event and email statistics)
- Exchange Online Mailbox Contents (message information using EWS)
- Skype for Business/Lync CDR Database (peer-to-peer session and conference information)
- Skype for Business/Lync QoE Database (Quality of Experience data)
- Cisco CDR Logs (peer-to-peer calls, ad hoc conference calls, and meet-me conference calls)

Some data collections take a once-a-day “snapshot” of the data even if they are scheduled to run more than once a day. Data collections that gather data once a day are:

- Active Directory Domain Controllers (user information through LDAP)
- Exchange Configuration (Exchange server, database, mailbox, mobile device configuration, and statistics using LDAP and PowerShell)
- Exchange Public Folders (statistics for legacy and new public folders)
- Exchange Online Hybrid Mailbox Configuration (mailbox information using LDAP and PowerShell)
- Exchange Online Hybrid User Configuration (users and distribution groups in hybrid environments using LDAP and remote PowerShell)

- Exchange Online Native Mailbox Configuration (mailbox information using remote PowerShell)
- Exchange Online Native User Configuration (user information using remote PowerShell)
- Office 365 User Subscription Configuration (subscription configuration using remote PowerShell)
- Skype for Business/Lync Configuration (server, pool, and user policy configuration using LDAP and PowerShell)
- Cisco Configuration (LDAP for Active Directory and Cisco Unified Communications Manager)

Additionally, there are data collections that run in the background and update data only as needed:

- Database Consistency (used to ensure messages received before and after UTC midnight are inserted in the database only once).
- Exchange Calculation (recipient response times, message delivery times)
- Cisco Calculation (duration times for ad hoc conferences)
- Data Query Availability (used to make the data available for queries. Data collected between runs of the Data Query Availability job are not available until the job runs)

For more information, see [How often do the data collections actually gather data and when do they run?](#) on page 122.

Viewing the collection job status

After the data collection start date, specified on the Data Collection page, the UC Analytics data collections begin to gather information. To view the progress of the collection jobs and to see if there are any errors during the collections, click **View Data Collection Status** at the top right of the Data Collections page.

Figure 2. Viewing data source collection job status.

The screenshot shows the 'Jobs' page in UC Analytics. At the top, there are filters for 'Jobs' (2/2/2017 to 2/16/2017) and 'Last 14 days'. Below the filters are dropdowns for 'No status filter' and 'No job type filter', and a search box 'Search by data source name'. A 'Clear all' button is also present. The main table has the following data:

Job Type	Data Source	Start Date	End Date	Details
✓ Data Query Availability		2/16/2017 12:30:00 PM	2/16/2017 12:30:00 PM	Data query availability job complete.
✓ Data Query Availability		2/16/2017 12:15:00 PM	2/16/2017 12:15:00 PM	Data query availability job complete.
✓ Data Query Availability		2/16/2017 12:00:00 PM	2/16/2017 12:00:00 PM	Data query availability job complete.
✓ Data Query Availability		2/16/2017 11:45:00...	2/16/2017 11:45:00...	Data query availability job complete.
✓ Data Query Availability		2/16/2017 11:30:00...	2/16/2017 11:30:00...	Data query availability job complete.
✓ Data Query Availability		2/16/2017 11:15:00...	2/16/2017 11:15:00...	Data query availability job complete.
✓ Data Query Availability		2/16/2017 11:00:00...	2/16/2017 11:00:00...	Data query availability job complete.
✓ Data Query Availability		2/16/2017 10:45:00...	2/16/2017 10:45:00...	Data query availability job complete.

At the bottom, the 'Details' pane shows 'No details' and has checkboxes for 'Show errors only' and 'Show warnings only', both of which are checked.

By default, UC Analytics only shows the status from the last job run for each data source (Most recent / running). If a collection job shows errors or warnings, you can view the details in the Details pane at the bottom. Use the

Show errors only check box and the **Show warnings only** check box to select whether you want to only errors, only warnings, or both errors and warnings.

Using the filter options, you can filter for a specific data source or job status to determine if a data collection has been running correctly. You can also narrow the returned results by specifying a custom date range.

Copying the job details information

If you want to copy the contents of the job details, double-click on the row that contains the details you want. The Details Copy option displays and you can copy the details and paste the text into a separate document.

Filtering job status results by state or type

You can filter the job status results to view only jobs that succeeded, jobs that failed, jobs that are currently running, or jobs that completed with warnings. You can also filter for a specific data source (job type).

- To filter by job status, click ▼ in status field and select a status of running, succeeded, failed, or warning.
- To filter by the data collection, click ▼ in the job type field and select the data source for which you want to see all data collection job status information.

To search for a specific data source, enter part or all of the data source name in the search field.

For information about how often data collections gather data, see [How often do collections update the data?](#) on page 36.

Filtering job status results by date range

The UC Analytics keeps 30 days of job status records by default but the job status page can only display 15,000 records. If there are more than 15,000 records in the database, you can filter the results by date range.

By default, the job status page shows the most recent result for each data source collection.

To filter by date range

- 1 Click the ▼ beside **Most recent only** in the date range field
- 2 Select the time period that you want:
 - Today
 - Yesterday
 - Last 7 days
 - Last 14 days
 - Custom
- 3 If you select **Custom**, specify the start date and the end date for the time period:
 - a Enter or select the start date in the left field.
 - b Enter or select the end date in the right field.

The job status page displays only the job status records that fall within the specified date range.

Forcing a data source collection to run now

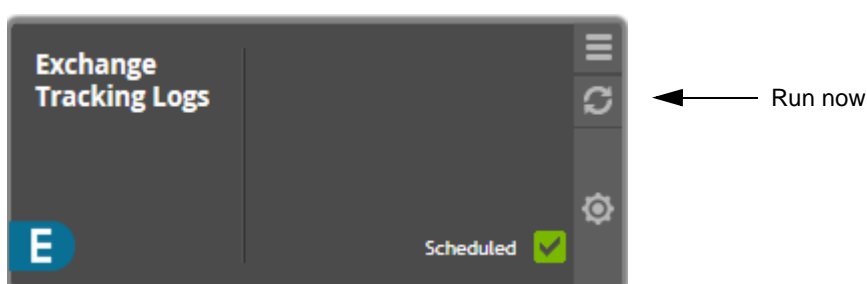
After you have created and scheduled your data source collections, you may want to have a collection run immediately. For example, you might have added new targets for a data source collection and want get updated data before the next scheduled collection.

To run a data source collection immediately

- 1 On the Admin Settings | Data Collections page, hover your cursor over the data source that you want to run.

Several icons display on the options menu on the right side of the data source tile.

Figure 3. Displaying the data source options menu.



- 2 Select the Run now icon.

A new data collection job is started for the data source.

If you use the Run now option for a data source that only takes a once-a-day “snapshot” of the data, no new data is collected unless you have added a new target. For information about how often data collections gather new data, see [How often do collections update the data?](#) on page 36.


Renaming a data source

You can create more than one data collection for the same data source. By adding multiple data source collections, you can distribute the load of collecting resource-intensive data.

For example, you can collect mailbox permissions and Active Directory extended permissions through the Exchange configuration data source. Collecting the permissions data can be very time-consuming. You can add and configure multiple Exchange configuration data sources to collect different subsets of the target mailboxes to collect this information.

Once you have added more than one Exchange configuration data source, you can rename each data source to more clearly identify what data is collected from each data source.


To rename a data source

- 1 On the Data Collection page, hover over the right corner of the data source tile and click .
- 2 Select **Rename**.
- 3 Enter a descriptive name for the data source and click **Save**.

Identifying your internal domains

Using the Classifications page, you identify the domains which are internal to your organization. You identify your internal domains to allow insights to show internal and external email traffic, Skype for Business/Lync sessions, and other related information correctly.

To identify internal domains

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Classifications** and click **Domain Classifications**.
- 3 Click **Add domain** for the Internal Domains.
- 4 Enter any domains that are considered internal to your business and click **Apply**.

You can modify or add domains at any time. After you have made changes, the existing data in the insights is reclassified (internal and external) immediately.

Guidelines when specifying domains

Some messages may not come directly from a domain so remember to add subdomains or to specify * in the domain name.

For example, you could add both DomainA.com and *.DomainA.com to the Domain Classifications list to include DomainA.com and all of its sub-domains. If you entered *.DomainA.com, you would get the domain and the subdomains but you would also get "AnotherDomainA.com" as well.

Entries in the Custom Domain Classifications list are not case-sensitive so it doesn't matter if you enter the domain names as capital letters or lower case letters.

You can also add your own custom classifications. Since there will not be a default insight for your new classification, you can create a new insight and use the filters to select the new classification.

Classifying domains for message traffic


You also can assign classifications to domains that determine how messages to and from the domains are classified in the insights. For example, using the Classification page you could specify that gmail.com should be classified as a "Personal" domain. For email where the sender or the receiver is in the specified domain, the messages are tagged with the associated classification.

You can assign domains to the following categories:

- Competitor
- Partner
- Customer
- Personal
- Social Network

i | **NOTE:** Domains which are not added to the internal domains list are treated as external, regardless of their classification.

To set classifications for messages

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Classifications** and click **Domain Classifications**.
- 3 Click **Add domain classifications**.
- 4 Enter the domain and the classification that is used for all messages sent from and to the domain, and click **Add**.
- 5 Repeat steps 3 and 4 for each domain you want to classify in your insights.

Granting full access to admin settings


In the Security settings, under the Access to Tenant Configuration heading, you can add a user to grant the account full access to the Admin Settings that are used to configure UC Analytics. This option only grants to the Admin Settings that are used to configure UC Analytics and does not grant access to the collected data.

For the same account to have access to the collected data in insights, both aggregate and unrestricted, you must also grant the account access for each type of data. For more information see [Granting users access to data](#) on page 41.

Adding a tenant administrator

If you have a multiple tenant implementation, you might want to create a tenant administrator who would only have access to the configuration settings for a specific environment (tenant).

To set up an account as a tenant administrator

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Security**.
- 3 Under the Access to Tenant Configuration section, click **Add Users**.
- 4 Leave the **Grant access to all users in all target environments** check box empty.
- 5 Select the target environment for the tenant.
- 6 Enter a specific user (email address or SAM account name) or enter a distribution group for all the users to be granted access.
- 7 Click **Add**.
- 8 Click **Save**.

Granting users access to data

The Security settings allow you to grant users access to data that is displayed in the insights. You can grant access (aggregate or unrestricted) to the different types of collected data such as:


- cross platform usage
- Exchange message data
- Exchange DLP data
- Skype for Business/Lync usage data

- Skype for Business/Lync QoE data
- Cisco usage data

You can grant access to all users in all target environments or grant access to specific users in a specific target environment.

i | **TIP:** When granting access to specific users, you can enter either an individual user or a distribution group. If you selected Office 365 as the target environment, it is not recommended that you enter a dynamic distribution group or a distribution group with a large number of members due to performance issues.

To grant data access to users

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Security**.
- 3 To grant access to a specific type of data, click **Add Users** in the section for that type of data.
- 4 Specify if the users have aggregate access or unrestricted access to the data.
- 5 If you want to grant access to all users, select the **Grant access to all users in all target environments** check box.

- OR -

Enter a specific user (email address or SAM account name) or enter a distribution group for all the users to be granted access.

i | **IMPORTANT:** For deployments with multiple target environments, if you are specifying a specific user or distribution group, ensure that the displayed target environment is the environment for which the user or group has rights. If necessary, select the correct environment from the dropdown list.

- 6 Click **Add**.
- 7 Click **Save**.

Unlike a product administrator who has access to the configuration settings for all environments (tenants), a tenant administrator can only configure settings for a specific environment.

Differences between aggregate and unrestricted access

If you have aggregate data access, you can view “public” information in insights. Public information is information that does not specifically identify both individuals in a messaging transaction or does not include “private” information such as message subject, file attachment name, or message ID. If you are collecting message body information, the message body is also considered private.

Other information is considered “sensitive” and may be available for aggregate access depending on the filters you have set in the insight. Sensitive information can include information such as: file attachment extensions, subject keywords, participants, send and received time of day, and importance.

Though you can view sensitive information with aggregate access, sensitive information is not available for specific individuals unless you have unrestricted access. Generally, with aggregate access, you can view insights that contain one sensitive item but not two or more sensitive items.

For example, if you have aggregate access to data, you can view insights that contain information such as

- How many emails were sent and received by your organization and who are the top senders?
- Which mail-enabled groups have not been active lately?
- How many emails were sent with file attachments of specific extensions and how big were they?

However, you cannot view an insight that shows private information such as

- what are the number of messages sent from one specific person to another specific person?

- what is the size of a specific mailbox and its last logon date?
- what are the number of messages from a department that contained a specific message subject? However, you can see the number of messages from a department that contain a specific subject keyword.
- which sent messages had a specific file attachment (such as “purchase.docx”)?

To see detailed and private information in insights, you must have a security access of Unrestricted for the type of data reported in the insight.

To grant access to specific collected data

For Exchange configuration and Exchange message data, there are different types of access that can be granted to users:

Table 15. Types of access that can be granted for Exchange message data.

If I have this type of access	I can see
Aggregate	summary (aggregate) information about the collected messages and public folder statistics. This access does not include details about individual messages or information that is considered “private”. By default, access to aggregate data is granted to everyone.
Unrestricted	Unrestricted access to the details of all the messages that everyone has sent or received in all the targeted mailboxes. It is recommended that this access be granted only to select personnel.

For cross-platform data, you can grant the following access to the aggregated data:

Table 16. Types of access that can be granted for cross-platform data.

If I have this type of access	I can see
Aggregate	summary (aggregate) information about the collected messages from Exchange and about the Skype for Business/Lync peer-to-peer sessions and conferences. This access does not include details about individual messages or about individual sessions and conferences. By default, access to aggregate data is granted to everyone.
Unrestricted	Unrestricted access to the details of all the messages that everyone has sent or received in all the targeted mailboxes, and details about all the sessions and conferences in which the targeted Skype for Business/Lync users participated. It is recommended that this access be granted only to select personnel.

For Skype for Business/Lync configuration and Skype for Business/Lync session, enterprise voice, and conference data, there are different types of access that can be granted to users:

Table 17. Types of access that can be granted for Skype for Business/Lync data.

If I have this type of access	I can see
Aggregate	summary (aggregate) information about Skype for Business/Lync sessions and conferences. This security access does not include details about individual sessions and conferences. By default, access to aggregate data is granted to everyone.
Unrestricted	unrestricted access to the details of all the Skype for Business/Lync sessions and conferences in the data collected from the CDR database. It is recommended that this access be granted only to select personnel.

For Skype for Business/Lync QoE data, you can provide separate access to users:

Table 18. Access that can be granted for Skype for Business/Lync QoE data.

If I have this type of access	I can see
Aggregate	summary (aggregate) information about Skype for Business/Lync QoE data. This access does not include details about individual calls, sessions, and conferences. By default, access to aggregate data is granted to everyone.
Unrestricted	unrestricted access to all the detailed QoE information for the individual calls, sessions, and conferences. It is recommended that this access be granted only to select personnel.

For DLP policy rule match data, you can provide separate access to users:

Table 19. Access that can be granted for DLP rule match data.

If I have this type of access	I can see
Unrestricted	unrestricted access to all the individual DLP policy rule matches. It is recommended that this access be granted only to select personnel.

For Cisco usage data, you can provide separate access to users:

Table 20. Access that can be granted for Cisco usage data.

If I have this type of access	I can see
Aggregate	summary (aggregate) information about Cisco usage data. This access does not include details about individual peer-to-peer sessions and conferences. By default, access to aggregate data is granted to everyone.
Unrestricted	unrestricted access to all the detailed Cisco information for the individual peer-to-peer sessions and conferences. It is recommended that this access be granted only to select personnel.

Setting display formats in UC Analytics

You access the insights through the UC Analytics home page at the following web site:

`http://<ServerName>/Analytics/`

i **NOTE:** If UC Analytics is configured for a native Office 365 environment, you will be prompted for your credentials. You can specify either your Windows credentials or your Office 365 credentials, depending on how UC Analytics was configured to for user authentication. For more information, see [Adding a target environment for native Office 365](#).

On the Welcome page, you have the option of either

- viewing a set of recommended insights
- accessing the insight library so you can select the insights you want

To view the list of all available insights, click the library icon  home page side bar.

Changing formats for date, time, and digit separators

By default, the UC Analytics date and time formats are set to month/day/year (M/d/yyyy) and hour:minute:seconds am or pm (h:mm:ss tt) for each user. When displaying numeric values, UC Analytics uses a period (.) for decimal values and a comma (,) as digit separators for thousands.

Each user can modify these settings to match their locale. These settings are configured per user. An administrator cannot set these values for all users.

To modify format settings

- 1 On the UC Analytics web site, click your user name that is displayed in top right corner.
- 2 Select **User Profile**.
- 3 Beside the Date Format or Time Format field, click ▼ and select the format you want.
- 4 Enter the values you want for the Decimal Separator and the Thousand Separator fields.
- 5 Click **Save**.

Overriding the time zone offset

UC Analytics determines the time zone that is used for scheduling and insights from the regional settings of the computer that is running the user browser. You can override the time zone that is used through your user profile.

To override the time zone setting

- 1 On the UC Analytics web site, click your user name that is displayed in top right corner.
- 2 Select **User Profile**.
- 3 Click the **Use Specific** button in the Time zone section.
- 4 In the field below, enter the value for the specific time zone offset that you want. For example, you might enter -5 or +2.
- 5 Click **Save**.

These settings are configured per user. An administrator cannot set these values for all users.

Adding data sources for Active Directory or Azure Active Directory

- [Adding data sources for different target environments](#)
- [Permissions needed to collect Active Directory data](#)
- [Adding data sources for Active Directory / Office 365 \(hybrid\)](#)
- [Adding data sources for native Office 365](#)

Adding data sources for different target environments

Depending on the target environment, you can add different data sources to collect information from either on-premise Active Directory or Azure Active Directory.

- If you are collecting from on-premise Active Directory or a hybrid Office 365 deployment in an Active Directory Forest / Office 365 (hybrid) target environment, you can add the following data sources:
 - Domain Controller
 - Office 365 User Subscription Configuration
- If you are collecting from a Azure Active Directory in an Office 365 (native) target environment, you can add the following data sources:
 - Office 365 User Subscription Configuration

For information about adding target environments, see [Adding multiple Active Directory forests](#) on page 30 and [Adding a target environment for native Office 365](#) on page 32.

Permissions needed to collect Active Directory data

To collect data from on-premise Active Directory or Azure Active Directory, you add different data sources to gather information. For each data source, you must specify the credential that is used to collect the data.

In most collections, you have the option to use the credential that is specified for the Data Engine service. If you want to use that credential for your data collections, ensure it has the permissions specified for that data source.

Permissions needed for the Domain Controller data source

By default, a target environment is configured for the Active Directory forest in which UC Analytics is installed. In the Active Directory Forest / Office 365 (hybrid) environment, you create a Domain Controller data source to collect on-premise Active Directory data.

The credentials that you must specify to collect the Active Directory data have the following permissions:

- Read permissions on all the Active Directory user, group and contact objects

Alternately, when you add a Domain Controller data source, you can choose to use the Data Engine service credentials if these credentials have read permissions on the Active Directory user objects. You specify the Data Engine credentials during installation.

Permissions needed for the Office 365 user subscription configuration data source

The Office 365 User Subscription data source data source collects the user subscription information for either a native or a hybrid Office 365 implementation including licenses and subscribed services such as Exchange Online, Skype for Business Online, SharePoint Online, etc.

The credentials that you specify to collect native Office 365 user subscription configuration are used to connect to Azure Active Directory using remote PowerShell. The credentials must:

- Have Azure Active Directory PowerShell access enabled.

Additional software prerequisites must be met if you want to collect Office 365 user subscription information. For more information, see the Software Requirements in the *UC Analytics Release Notes*.

Adding data sources for Active Directory / Office 365 (hybrid)

If you have an Active Directory on-premise or hybrid environment, there are two data sources that you can configure:

- Domain Controller which collects information about users, groups, and mail contacts from Active Directory using LDAP.
- Office 365 User Subscription Configuration which retrieves Office 365 user subscription information including licenses and services provisioning status.

Creating an Domain Controller data source

When you create a data collection for Active Directory, you configure the Domain Controller data source. The data collection collects information about users, groups, and mail contacts from Active Directory using LDAP.

If you have created more than one Active Directory forest target environment, remember to add a Domain Controller data source for each environment.

To set up an Active Directory collection

- 1 Click the gear icon  on the home page side bar.

- 2 Click **Data Collection**.
- 3 Click **+** beside the name of the target environment.
- 4 Select the **Domain Controller** check box.
- 5 Click the **Domain Controller** tile to open the configuration page.
- 6 Enable the data collection for the data source.
- 7 Select the types of data that you want to collect.
- 8 Specify the credentials to be used to collect from Active Directory.
- 9 If the Data Engine service credentials have read permissions on the Active Directory user objects, you can use Data Engine service account to collect from Active Directory.

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.
 - If you want to specify a specific domain controller instead of automatically discovering the domain controller, click **Show Advanced Settings** in the LDAP Collection Parameters section.
- 10 Click **Save**.

For information about creating Domain Controller data sources for a resource forest configuration, see [Configuring UC Analytics for resource forests](#) on page 31.

Creating a data source for Office 365 user subscription configuration


The Office 365 user subscription configuration data source is used to gather Office 365 user subscription information including licenses and service provisioning status using remote PowerShell. You can create the data source for both native and hybrid Office 365.

- For a hybrid Office 365 implementation, you add the Office 365 user subscription configuration data source in an Active Directory Forest / Office 365 (hybrid) target environment.
- For a native Office 365 implementation, you add the Office 365 user subscription configuration data source in an Office 365 (native) target environment.

i | **NOTE:** Office 365 data sources are not displayed in the Data Collection page until you have added Office 365 as a target environment. For more information, see [Adding a target environment for native Office 365](#) on page 32.

For more information about credential prerequisites, see [Permissions needed for the Office 365 user subscription configuration data source](#) on page 47.

To set up an Office 365 user subscription configuration collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click **+** beside the name of either an Active Directory / Office 365 (hybrid) target environment or an Office 365 (native) target environment.
- 4 Select the **Office 365 User Subscription Configuration** check box.
- 5 Click the **Office 365 User Subscription Configuration** tile to open the configuration page.
- 6 Enable the data collection for this data source.
- 7 Select the data that you want to collect.
- 8 Specify the Office 365 credential used to collect the user configuration data through PowerShell.

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.

9 Click **Save**.

Adding data sources for native Office 365

If you have a native Office 365 environment, there are one data source that you can configure to collect from Azure Active Directory using remote PowerShell:

- Office 365 User Subscription Configuration which retrieves Office 365 user subscription information including licenses and services provisioning status.

For information about how to configure the Office 365 User Subscription Configuration data source, see [Creating a data source for Office 365 user subscription configuration](#) on page 48.

Adding data sources, chargeback costs, and thresholds for Exchange and Exchange Online

- [Permissions needed to collect Exchange on-premise or hybrid data](#)
- [Permissions needed to collect from native Exchange Online](#)
- [Creating a data source for Exchange configuration](#)
- [Creating a data source for Exchange tracking logs](#)
- [Creating a data source for Exchange mailbox contents \(EWS\)](#)
- [Do I need both Exchange tracking log and mailbox contents \(EWS\) collections?](#)
- [Creating a data source for Exchange IIS Logs](#)
- [Creating a data source for Exchange public folders](#)
- [Adding Exchange Online hybrid data sources for hybrid Office 365](#)
 - [Creating a data source for Exchange Online hybrid user configuration](#)
 - [Creating a data source for Exchange Online hybrid mailbox configuration](#)
- [Adding Exchange Online data sources for native Office 365](#)
 - [Creating a data source for Exchange Online native user configuration](#)
 - [Creating a data source for the Exchange Online native mailbox configuration data source](#)
 - [Creating a data source for Exchange Online mailbox contents \(EWS\) data source](#)
- [Setting chargeback costs for Exchange](#)
- [Setting thresholds for Exchange metrics](#)
- [Omitting words when filtering by subject or body](#)

Permissions needed to collect Exchange on-premise or hybrid data

To collect data from Exchange, you add different data sources to gather information. For each data source, you must specify the credential that is used to collect the data. In most collections, you have the option to use the credential that is specified for the Data Engine service. If you want to use that credential for your data collections, ensure it has the permissions specified for that data source.

By default, a target environment is configured for the Active Directory forest in which UC Analytics is installed. In the Active Directory forest environment, you can create the following data sources to collect data from both on-premise Exchange and hybrid Exchange Online:

- **Exchange Configuration:** This data source collects Exchange configuration and mailbox properties using PowerShell and Active Directory LDAP queries. For information about the required permissions, see [Permissions needed for the Exchange configuration data source](#) on page 51.
- **Exchange Tracking Logs:** This data source collects message traffic and DLP rule match data from the Exchange tracking logs. For information about the required permissions, see [Permissions needed for the Exchange tracking logs data source](#) on page 52.
- **Exchange Mailbox Contents:** This data source collects mailbox contents from on-premise Exchange using Exchange Web Service (EWS). For information about the required permissions, see [Permissions needed for the Exchange mailbox contents data source](#) on page 52.
- **Exchange IIS Logs:** This data source collects ActiveSync information about ActiveSync mobile devices from IIS logs. For information about the required permissions, see [Permissions needed for the Exchange IIS logs \(ActiveSync\) data source](#) on page 53.
- **Exchange Public Folders:** This data source collects legacy public folder data from Exchange 2007 and Exchange 2010 and new public folder data from Exchange 2013 and Exchange 2016. For information about the required permissions, see [Permissions needed for the Exchange public folders data source](#) on page 53.
- **Exchange Online Hybrid User Configuration:** This data source collects detailed information about users and distribution groups in a hybrid environment using remote PowerShell. See [Permissions needed for the Exchange Online hybrid user configuration data source](#) on page 54.
- **Exchange Online Hybrid Mailbox Configuration:** This data source collects Exchange Online mailbox statistics such as mailbox size, permissions, and mobile devices from hybrid (on-premise and Office 365) environments. For information about the required permissions, see [Permissions needed for the Exchange Online hybrid mailbox configuration data source](#) on page 54. For a native Office 365 implementation, use the Exchange Online Mailbox Configuration data source.
- **Exchange Online Mailbox Contents:** This data source collects mailbox content information from hybrid or native Office 365 using Exchange Web Services (EWS). For information about the required permissions, see [Permissions needed for Exchange Online mailbox contents data](#) on page 54.

About adding an account to an Exchange role group

When the credential used in a data source must be a member of an Exchange role group, it is important to add the account to the role group correctly. If you simply add a user to a group using Active Directory Users & Computers, (ADUC) utility, the user is not granted the full rights of the Exchange group role. In this case, the data source permissions are not met.

To provide the full grant of an Exchange admin role group to an account, you must use the PowerShell Exchange cmdlet `Add-ExchangeAdministrator` or the Exchange Admin Center (EAC).

Permissions needed for the Exchange configuration data source

The credentials that are used to collect Exchange configuration and mailbox properties through PowerShell and LDAP must be a member of specific security groups.

For Exchange 2007-only environments and Exchange mixed environments that include Exchange 2007, the credential must:

- Be a member of Exchange View-Only Administrator role group.

For Exchange 2010-only, Exchange 2013-only, or Exchange 2016-only environments, or Exchange mixed environments without Exchange 2007, the credential must:

- Be a member of the Public Folder Management role group or Organization Management role group. (The View-Only Organization Management role group is not sufficient.)



NOTE: If you require information about database copies, ensure that the LDAP connection credential has read permissions to ms-Exch-MDB-Copy objects in Active Directory:

- To grant permissions directly, modify permissions in Active Directory.
- To grant permissions indirectly, add the credential to either the Exchange Organization Management or the Exchange Public Folder Management role group.

Permissions needed for the Exchange tracking logs data source

For Exchange tracking logs, the specified credentials are used to collect message traffic and DLP incident data from the Exchange message tracking logs. (DLP data is collected only for Exchange 2013 and later.) The credentials must:

- Have access rights to the share that contains the Exchange message tracking logs
- Be a member of the View-Only Organization Management security group
- Be a member of the Public Folder Management security group

You can gather Exchange message tracking logs from any of the following locations:

- Directly from tracking log folders on each Exchange server
- From user-created file shares on each Exchange server, such as `\\ServerName\MessageTracking`
- From any file share that contains up-to-date Exchange tracking log copies. Original log files and compressed zip files (using either WinZip or Windows native compressed format) are supported in the share.

For information about gathering historical tracking logs from a centralized location, see [About collecting historical tracking logs](#) on page 60.

Default locations for tracking logs

By default, the Exchange tracking log folder is located as follows:

- Exchange 2007
`\\ServerName\c$\Program Files\Microsoft\Exchange Server\TransportRoles\Logs\MessageTracking`
- Exchange 2010
`\\ServerName\c$\Program Files\Microsoft\Exchange Server\V14\TransportRoles\Logs\MessageTracking`
- Exchange 2013
`\\ServerName\c$\Program Files\Microsoft\Exchange Server\V15\TransportRoles\Logs\MessageTracking`
- Exchange 2016
`\\ServerName\c$\Program Files\Microsoft\Exchange Server\V15\TransportRoles\Logs\MessageTracking`

Permissions needed for the Exchange mailbox contents data source

Any credentials that are used to collect Exchange mailbox contents using Exchange Web Services (EWS) must have “Exchange Impersonation” permissions to all the target mailboxes:

- For information about how to set Exchange impersonation for Exchange 2013 or Exchange 2016, see [Setting impersonation for Exchange 2013 or Exchange 2016](#) on page 97.

- For information about how to set Exchange impersonation for Exchange 2010, see the Microsoft article: [Configuring Exchange Impersonation](#).
- For information about how to set Exchange impersonation for Exchange 2007, see the Microsoft article: [Configuring Exchange Impersonation \(Exchange Web Services\)](#).
- For information about how to set impersonation for Exchange Online (Office 365), see [Setting impersonation for Office 365](#) on page 98.

Permissions needed for the Exchange IIS logs (ActiveSync) data source

When you configure an Exchange IIS logs data source, you specify two sets of credentials:

- credentials for the LDAP connection to Active Directory
- credentials for collecting the IIS log files from the Exchange CAS (Client Access Services)

Credentials to collect data through LDAP

The credentials that are used to collect device and user information through LDAP from Active Directory must have the following Active Directory permissions:

- Read permission on User objects.
- Read permission on the msExchMailboxGuid property of all User objects

Usually, these read permissions are available to members of the Authenticated Users group. Consequently, you only need to be an authenticated user of the domain or of another domain that is trusted by the domain.

Credentials to collect the IIS log files

The credentials that are used to collect information from the IIS log files on the Exchange Client Access Server (CAS) must have the following permissions:

- Local Administrators rights on all Exchange CAS servers

The Local Administrators rights are required to access to the IIS logs through an administrative volume share, such as C\$.

As an alternative to providing Local Administrators rights, you could create a non-administrative share for the IIS log folder. You could then grant read access to the credentials to the IIS log files through the share.

Also, IIS logging must be configured on the Exchange CAS servers. For more information, see [Appendix C: Configuring IIS Log Files to capture ActiveSync events](#) on page 103.

Permissions needed for the Exchange public folders data source

When you configure an Exchange public folders data source, you can select whether you want to collect data from legacy public folders (Exchange 2007, Exchange 2010) or new (Exchange 2013, Exchange 2016) public folders.

For Exchange 2007, ensure that the Microsoft Exchange Server 2007 Management Tools are installed.

For Exchange 2010, Exchange 2013, or Exchange 2016 the credentials that you specify to collect public folder data using remote PowerShell must:

- Be a member of the Exchange View-Only Administrator role.
- Have remote PowerShell access enabled.

Permissions needed for the Exchange Online hybrid user configuration data source

You would create an Exchange Online hybrid user configuration data source only if you have an Exchange/Office 365 hybrid environment. The credentials that you specify to collect Active Directory user data using LDAP must have the following permissions:

- Read permissions on all the Active Directory user objects.

The credentials that you specify for Exchange Online PowerShell to collect user data using remote PowerShell and must have the following permissions:

- Be a member of the Exchange View-Only Organization Management role in the Exchange Online tenant.
- Have PowerShell access enabled.

Permissions needed for the Exchange Online hybrid mailbox configuration data source

The credentials that you specify to collect Active Directory user data using LDAP must have the following permissions:

- Read permissions on all the Active Directory user objects.

The credentials that you specify to collect Exchange Online hybrid mailbox configuration data (including statistics, permissions, and mobile devices) using remote PowerShell must:

- Be a member of the Exchange View-Only Organization Management role in the Exchange Online tenant.
- Have PowerShell access enabled.

To collect mailbox recipient “send as” permissions, the credentials must also:

- Be a member of the Recipient Management role group in the Exchange Online tenant.

The credentials that are used to collect Active Directory user data must have the following permissions:

- Read permissions on all the Active Directory user objects.

Permissions needed for Exchange Online mailbox contents data

The Exchange Online mailbox contents data source can be used for either hybrid or native Office 365 deployments.

- You create an Exchange Online mailbox contents data source for a hybrid deployment by adding the data source in an Active Directory Forest/Office 365 Hybrid target environment.
- You create an Exchange Online mailbox contents data source for a native Office 365 deployment by adding the data source to an Office 365 (native) target environment.

The credentials that you specify to collect Exchange Online mailbox contents data using EWS must:

- Have Exchange impersonation permissions for all of the target mailboxes.
- Have Exchange Web Services (EWS) access enabled.
- Be a member of the Exchange View-Only Organization Management role in the Exchange Online tenant.
- Have PowerShell access enabled.

Permissions needed to collect from native Exchange Online

If you have a native Office 365 target environment, you can create the following data sources to collect information from your native Exchange Online deployment. You must set up an Office 365 (native) target environment and specify the required credentials when you configure the Exchange Online data sources.

- **Exchange Online Native User Configuration:** This data source collects user and distribution group data using remote PowerShell. For information about the required permissions, see [Permissions needed for Exchange Online native user configuration data source](#) on page 55.
- **Exchange Online Native Mailbox Configuration:** This data source collects Exchange Online native mailbox statistics such as mailbox size, permissions, and mobile devices. For information about the required permissions, see [Permissions needed for Exchange Online native mailbox configuration](#) on page 55. For a hybrid Office 365 implementation, use the Exchange Online Hybrid Mailbox Configuration data source.
- **Exchange Online Mailbox Contents:** This data source can collect mailbox content information from either hybrid or native Office 365 using Exchange Web Services (EWS). For information about the required permissions, see [Permissions needed for Exchange Online mailbox contents data](#) on page 54.

Permissions needed for Exchange Online native user configuration data source

The Exchange Online Native User Configuration data source collects user and distribution group data using remote PowerShell. The credentials that you specify are used to collect user configuration data using remote PowerShell from native Office 365. The credentials must:

- Be a member of the Exchange View-Only Organization Management role in the Exchange Online tenant.
- Have PowerShell access enabled.

Permissions needed for Exchange Online native mailbox configuration

The credentials that you specify to collect Exchange Online native mailbox configuration data (including statistics, permissions, and mobile devices) using remote PowerShell must:

- Be a member of the Exchange View-Only Organization Management role in the Exchange Online tenant.
- Have PowerShell access enabled.

To collect mailbox recipient “send as” permissions, the credentials must also:

- Be a member of the Recipient Management role group in the Exchange Online tenant.

Creating a data source for Exchange configuration

The Exchange configuration data collection gathers information about the Exchange hierarchy such as organizations, servers, DAGs, database status, database copies, and mailbox properties through PowerShell and LDAP queries to Active Directory.

When you configure an Exchange configuration collection, you specify the LDAP connection credentials, target mailboxes (optional), the Exchange server from which you want to collect, and remote PowerShell connection details.


Why should I specify target mailboxes?

If you specify target mailboxes, mailbox statistical and configuration information is also collected for those mailboxes including:

- Mailbox configuration (quotas, home database)
- Mailbox statistics (size, Item counts, deleted Item counts, storage status, and last logon time)
- Mobile devices associated with the mailbox (device model, device type, first sync time, last sync time, policy update time)
- Mailbox permissions and Active Directory extended permissions
- Archive mailbox (personal archive) statistics

This information is not collected for mailboxes that are not set as target mailboxes. (You can create a dynamic distribution group for all your mailboxes and specify the group as the target in the data source.)

To set up an Exchange configuration collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click **+** beside the name of the target environment.
- 4 Select the **Exchange Configuration** check box.
- 5 Enable the data collection for this data source.
- 6 Select the types of data that you want to collect.
- 7 Specify credentials to be used for the LDAP connection or select the **Use Data Engine service credential** option. See [Permissions needed for the Exchange configuration data source](#) on page 51.

For an Exchange resource forest deployment

If you have an Exchange resource forest deployment, you must configure the options in the **Account Forest LDAP Connection Parameters** section.

- a Select the **Search additional forests for user accounts** check box.
The options under **Show advanced settings** are expanded.
 - b Select the appropriate option for the Account Forest Domain Controller:
 - Automatically discover domain controller in specific domain
 - Use specific domain controller
 - c Specify the credentials that are used to access the domain controller using LDAP.
- 8 Specify any target mailboxes. You can enter the common name or email address of a mailbox, a group, or a dynamic distribution group. For more information about using a dynamic distribution group, see [Using](#)

[dynamic distribution groups to group target users](#) on page 58.

- 9 Specify the Exchange version from which you are collecting data. For a mixed environment, select the latest version of Exchange that is installed.

i | **NOTE:** For Exchange 2007 mailboxes, you must create a separate data source even for mixed Exchange organizations.

For a mixed organization containing Exchange 2010 and Exchange 2013/2016 servers, select Exchange 2013/2016 to get the most comprehensive configuration data.

To collect from Exchange 2010, Exchange 2013, or Exchange 2016, you must

- Click **Add server** and enter the Client Access Services (CAS) server name (and port number if applicable).

- 10 Specify the explicit credentials to be used to make the PowerShell connection.

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.
- To specify a specific domain controller instead of automatically discovering the domain controller, click **Show Advanced Settings** in the LDAP Collection Parameters section. The domain controller must be a global catalog server.

i | **NOTE:** If the Data Engine is in a different forest from the data to be collected, specify the domain controller you want to use for data collection. If you leave the default setting to automatically discover domain controller, the data collector selects a domain controller from the forest that hosts the Data Engine, not a domain controller from the forest from which you are collecting the data.

- To use HTTPS instead of HTTP to connect to the Exchange 2010/2013/2016 CAS server, click **Show Advanced Settings** in the PowerShell Connection Parameters section. This option requires that TLS/SSL is enabled for remote PowerShell on the Exchange CAS server.

The Advanced Settings also provide options if you have a customized PowerShell virtual folder, or if you want specify either an implicit credential or the Data Engine credential be used to make the PowerShell connection.

- 11 Click **Save**.

For more information about credential prerequisites, see [Permissions needed for the Exchange configuration data source](#) on page 51.

For information about creating Exchange Configuration data sources for a resource forest configuration, see [Configuring UC Analytics for resource forests](#) on page 31.

Best practices for gathering performance

To obtain the best performance from the Exchange configuration data collection, you can add several Exchange configuration data sources to collect the data. Each data source would specify a different group of target users.

It is recommended that you group the target users according to the geographic location of the mailboxes. You can add a separate data source to collect from each geographical area. Select the server hosting Client Access Services (CAS) that is nearest to the mailbox server for each data source you configure.

In another example of distributing the data collection, you might want configure a separate Exchange configuration data source to collect mailbox permissions data. You could configure one Exchange configuration data source to collect mailbox permissions from a select number of target mailboxes. You could add another data source to collect the rest of the data for the same mailboxes.

Once you have added multiple data sources for Exchange configuration data, you can name each data source to reflect the information that is being collected. For information about renaming a data source, see [Renaming a data source](#) on page 39.

Limitations for the Exchange configuration collection

For Exchange 2007, the Exchange configuration data collection does not gather the following information:

- Messages in the Exchange dumpster (deleted messages)
- Attachment size and attachment last modified date
- Time of day (send/receive) and “after hours” information
- Store (database) size for mailbox databases and public folder databases

Using dynamic distribution groups to group target users

A dynamic distribution group includes any recipient in Active Directory with attributes that match the filter and conditions that you have defined.

To group target users by geographic location, you can create multiple dynamic distribution groups. All the mailboxes in each group would belong to same geographic location.

Suppose you have several data center locations within a global domain such as LON01, NYC01, and TKY01. You would create multiple data sources, one for each location.

For example, the data source for LON01 site would use the dynamic distribution group that contains only mailboxes from LON01 and would use a CAS server from LON01.

If you wanted to use an organization unit (OU) as the collection target, you could create a dynamic distribution group that has an OU as its scope.

i **TIP:** To create dynamic distribution group for mailboxes from a single location, you could use a filter based on a mailbox custom attribute.

- 1 Run a PowerShell script to get all the mailbox databases.
- 2 Get the CAS server for each mailbox database.
- 3 Set the corresponding CAS server name to be one of the Custom Attributes for each mailbox.
- 4 After the attribute is set, use that Custom Attribute as a filter when you create dynamic distribution groups.

Specifying the types of mailboxes to be included

When you create your dynamic distribution group, you set the filters you determine which mailboxes and users will be included. If you want the data collection to gather all user mailboxes, shared mailboxes, resource mailboxes, and system mailboxes you would select **All recipient types**.

If you do not want to include system mailboxes, you could select specific filters to include only certain types of users and mailboxes such as:

- Users with Exchange mailboxes
- Users with external email addresses
- Resource mailboxes
- Contacts with external email addresses
- Mail-enabled groups

Can I enter the Domain Users group as the target for the data collection?

When you specify target users for data collections, you can use distribution or security groups. However, UC Analytics does not support the group Domain Users as a target group. Domain Users is a special Microsoft group that does not contain a membership setting so UC Analytics cannot resolve it. UC Analytics supports all the groups with membership information. You can check a group's membership setting using the ADSI Edit tool or ADAC (Active Directory Administrative Center).

For Exchange configuration, the data collection gathers the configuration for mailboxes, not for users. For this reason, it is not recommended that you use a group such as Domain Users that contains users without mailboxes, mail contacts, and other types of recipients. The data collection will complete but it will be full of warnings in the job details for every “not applicable” user.

When creating a group to use as a target, it is recommended you create a static group with a specified membership or create a dynamic group. For a dynamic group, select the check boxes for users with Exchange mailboxes. That will cover all the Exchange mailboxes in the organization. For static groups, use only groups that have members with Exchange mailboxes and do not include any other kind of users.

Troubleshooting the Exchange configuration collection

In some situations, if you are collecting mailbox properties from an Exchange server without using TLS/SSL, you might see error messages such as the following:

“Connecting to remote server failed with the following error message: The WinRM client cannot process the request. Unencrypted traffic is currently disabled in the client configuration. Change the client configuration and try the request again. For more information, see the about_Remote_Troubleshooting Help topic.”

You must enable unencrypted traffic for the WinRM client. You can use the Local Group Policy Editor to modify the WinRM client settings.

To enable unencrypted traffic

- 1 Click **Start**, type **gpedit.msc** in the Start Search box and press **ENTER** to open the Local Group Policy Editor.
- 2 In the tree view, navigate to Local Computer Policy | Computer Configuration | Administrative Templates | Windows Components | Windows Remote Management (WinRM) | WinRM Client.
- 3 Enable the following settings:
 - Allow unencrypted traffic
 - Trusted Hosts
- 4 Select the Trusted Hosts setting.
- 5 Enter the servers to which you want to connect in the Trusted Hosts list. Wild cards are accepted.


Creating a data source for Exchange tracking logs

The Exchange tracking log collection gathers email message information for all mailboxes. You must gather the Exchange message tracking logs from all Exchange hub transport servers and all Exchange mailbox servers. If you have historical tracking logs stored in a central location, you can gather from those tracking logs as well. To configure the Exchange tracking log data collection, you specify the following:

- the locations of the Exchange message tracking log shares from which you want to gather data
- the credentials used to access the tracking logs

For information about the permissions needed, see [Permissions needed for the Exchange tracking logs data source](#) on page 52.

To set up an Exchange tracking log collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click **+** beside the name of the target environment.
- 4 Select the **Exchange Tracking Logs** check box.
- 5 Enable the data collection for this data source.
- 6 Select the types of data that you want to collect. (DLP data is collected only for Exchange 2013 and later.)
- 7 Specify credentials to be used for the LDAP connection or select the **Use Data Engine service credential** option.
- 8 Specify each folder location from which the tracking logs are to be gathered.
- 9 Specify the credentials that are used to gather the tracking logs.

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.
- To specify a specific domain controller instead of automatically discovering the domain controller, click **Show Advanced Settings** in the LDAP Collection Parameters section.

If the Data Engine is in a different forest from the data to be collected, you must specify the domain controller you want to use. If you automatically discover domain controller, the data collector selects a domain controller from the forest that hosts the Data Engine, not a domain controller from the forest from which you are collecting data.

- 10 Click **Save**.

About collecting historical tracking logs

Tracking logs from each Exchange server have the same file name format, so you must have a separate folder for each server. Do not change the tracking log file names. The naming convention for the message tracking log files is as follows:

MSGTRKyyyyymmdd-nnnn.log
MSGTRKMAyyyyymmdd-nnnn.log
MSGTRKMDyyyyymmdd-nnnn.log
MSGTRKMSyyyyymmdd-nnnn.log

When you configure the tracking log data collection, you specify each folder from which you are gathering tracking logs.

The data collection job goes through the files each time it is scheduled to run (every hour by default). However, the job reads only new changes (new files or files that have changed since the last time the job ran). This method allows the data collection job to collect and update the data in a timely manner without reading the log files multiple times.

Creating a data source for Exchange mailbox contents (EWS)

The Exchange mailbox contents data collection gathers email message information from user mailboxes using Exchange Web Services (EWS). When you configure a mailbox contents data collection, you must specify the target mailboxes and the Client Access Server (CAS) URL that is used to connect to the mailbox server. For Exchange 2016, you specify the URL for the Client Access Services on the mailbox server.

TIP: It is recommended that you add the user mailboxes to a distribution group. You can then add the distribution group for the target mailboxes.


If you have an Exchange hybrid environment, you can configure an Exchange Online Mailbox Contents data source to collect Office 365 mailboxes. For more information, see [Creating a data source for Exchange Online mailbox contents \(EWS\) data source](#) on page 71.

In addition to message traffic, the mailbox contents collection gathers the following message information from the target mailboxes:

- Subject
- Body
- File attachments
- Localized send / receipt time of day (used to show “after hours” and “response time” data in insights)

For information about how times are calculated on insights when you use the “response time” or “after hours” filters, see the section titled *How the Filters Work* in the *UC Analytics User Guide*.

To set up a mailbox contents (EWS) data collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click **+** beside the name of the target environment.
- 4 Select the **Exchange Mailbox Contents** check box.
- 5 Enable the data collection for this data source.
- 6 Select the data to collect.

If you specify Body, it is recommended that you limit the number of mailboxes to a key group of user mailboxes for this specific data collection.

- 7 Specify the Windows credentials for the LDAP connection or select the **Use Data Engine service credential** option.
- 8 Specify the target mailboxes. You can enter a mailbox, distribution group, dynamic distribution group, or a security group in one of the following formats:
 - Account name
 - Domain\Account name
 - Email address

- 9 Select the versions of Exchange from which you want to collect data:
 - Collect from mailboxes on all supported Exchange versions
 - Specify a specific Exchange version from which you want to collect data.
- 10 Specify the credential used to collect the mailbox data. For more information see [Permissions needed for the Exchange mailbox contents data source](#) on page 52.
- 11 Specify the URLs for the CAS (Client Access Servers or Client Access Services) to be used by Exchange Web Services (EWS).
 - Use Exchange auto-discovery for each mailbox
 - Specify the specific Exchange Web Services URLs
 - i** | **TIP:** In many Exchange environments, Exchange auto-discovery does not always work so it is recommended that you specify individual Exchange Web Services URLs.

When you specify the CAS for the EWS gathering, you must enter the URL for the server. For example, the URL might be:

```
https://MyCASServer/ews/Exchange.asmx
```

You can specify multiple CAS URLs to be used for the EWS gathering. If you specify multiple CAS URLs, data collection is faster. If you do not specify one or more CAS URLs, the Exchange Autodiscover service is used to find the CAS servers for the target mailboxes.

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.
- If you want to specify a specific domain controller instead of automatically discovering the domain controller, click **Show Advanced Settings** in the LDAP Collection Parameters section.

If the Data Engine is in a different forest from the data to be collected, you must specify the domain controller you want to use for data collection. If you leave the default setting to automatically discover domain controller, the data collector selects a domain controller from the forest that hosts the Data Engine, not a domain controller from the forest from which you are collecting the data.

- 12 Click **Save**.

Tips for better performance for mailbox contents collection

You can specify more than one set of Exchange gathering credentials and you can specify more than one CAS server for EWS data collection.

If you specify multiple EWS credentials and multiple CAS servers, it allows the EWS gathering to collect data from more mailboxes in parallel. This will reduce the time that is needed to collect the data.

For information about the types of data collections and the rate at which message data is collected, see [Determining where to install services](#) on page 17.

You could also increase the Exchange throttling quota for the accounts used to collect the message data.

i | **IMPORTANT:**

If you configure a mailbox content data gathering to collect the message body, storage requirements can be doubled. It is strongly recommended that if you collect the message body, you specify only a limited number of target mailboxes for this data collection.

How many CAS URLs and credentials are needed for mailbox content data collection?

By increasing the number of CAS URLs and the number of credentials that are used for collection of the Exchange data, you can shorten the time that it takes EWS to collect the message data.

Assuming that Exchange throttling is set to the default value for the gathering credentials, the estimated collection times are as follows:

- The initial data collection is a collection that collects 30 days data from all the target mailboxes. Messages that are older than 30 days are not collected even if they are still in the mailbox. You can expect a collection rate of 25 mailboxes an hour per CAS server/credential pair set.

You can change the number of days for the initial data collection period when you configure the connection on the Data Collection page under Admin Settings.

- After the initial collection, each subsequent collection is an ongoing collection in which only new or changed messages are collected. The ongoing collection runs every 30 minutes. Collection times decrease and you can expect a collection rate of 100 mailboxes an hour per CAS server/credential pair set.

Do I need both Exchange tracking log and mailbox contents (EWS) collections?

Though the mailbox contents (EWS) data collection and the Exchange tracking log data collection both gather message information, there are some differences.

Table 21. Comparison of Exchange tracking log and mailbox contents collections

Exchange tracking log data collection	Mailbox contents (EWS) data collection
The Exchange tracking log collection gathers message information for all mailboxes.	The mailbox contents collection gathers message information for only the specified target mailboxes.
The Exchange tracking log collection gathers both message data and DLP rule matches.	The mailbox contents collection (EWS) gathers only message data but includes specific data not found in tracking logs such as <ul style="list-style-type: none"> • file attachments • “in reply to” • time-of-day data. The mailbox contents data collection can also include message body text if configured. <p>NOTE: If a message is collected through EWS and the message was sent from a mailbox user to a distribution group to which the user belongs, the user is counted only as a sender, not a recipient. If the tracking log collection was also run, that user is also counted as a recipient.</p>
The Exchange tracking log collection is much faster than the mailbox contents collection.	The mailbox contents collection is much slower than the Exchange tracking log collection.
You run the Exchange tracking log collection against both Exchange mailbox servers and hub transport servers.	The mailbox contents collection accesses the Exchange CAS (Client Access Servers) to gather data.

For more detailed information, see [What are the differences between the Exchange Mailbox Contents and Exchange Tracking Logs data sources?](#) on page 125.


Creating a data source for Exchange IIS Logs

You configure an Exchange IIS logs data source to collect mobile device activities such as messages sent, messages received, and device information from the IIS log files.

You must collect the logs of the IIS sites for both on your front-end Exchange Client Access Server (CAS) and on your back-end Exchange Mailbox servers. The back-end IIS logs are required for information about the number of messages that are downloaded and uploaded using ActiveSync.

- For information about what permissions are required by the credentials for the data collection, see [Permissions needed for the Exchange IIS logs \(ActiveSync\) data source](#) on page 53.
- For information about how IIS logging must be configured on your Exchange CAS servers, see [Appendix C: Configuring IIS Log Files to capture ActiveSync events](#) on page 103.
- For information about the ActiveSync events that are collected, see [What ActiveSync events are collected and displayed in the insights?](#) on page 105.

To set up an Exchange IIS logs collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click **+** beside the name of the target environment.
- 4 Select the **Exchange IIS Logs** check box.
- 5 Enable the data collection for this data source.
- 6 Specify the Windows credentials to be used for the LDAP connection to Active Directory or select the **Use Data Engine service credential** option.
- 7 Enter the path location for the IIS log files that you want to collect.
- 8 Enter the credentials needed to access the Exchange server that hosts the log files.

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.
- If you want to specify a specific domain controller instead of automatically discovering the domain controller, click **Show Advanced Settings** in the LDAP Collection Parameters section.

i **NOTE:** If the Data Engine is in a different forest from the data to be collected, you must specify the domain controller you want to use for data collection. If you leave the default setting to automatically discover domain controller, the data collector selects a domain controller from the forest that hosts the Data Engine, not a domain controller from the forest from which you are collecting the data.

- 9 Click **Save**.

About the IIS log file locations

The Exchange ActiveSync IIS log files are located either in user-created file shares or in the IIS log directories of your Exchange Client Access Servers (CAS) and Exchange Mailbox servers.

- If IIS is configured for one log file per server, the default path for the IIS log files is: \\servername\c\$\inetpub\logs\LogFiles\W3SVC.
- If IIS is configured for one log file per site, the default path for the IIS log files is: \\servername\c\$\inetpub\logs\LogFiles\W3SVC n where n is the site ID of the site.

For example, if the front-end site is site #1 and the back-end site is site #2 and both are on the same Exchange server, the default paths will be as follows:


- \\servername\ c\$\inetpub\logs\LogFiles\W3SVC1
- \\servername\ c\$\inetpub\logs\LogFiles\W3SVC2

Creating a data source for Exchange public folders

By configuring an Exchange public folders data source, you can collect statistics for your legacy public folders (Exchange 2007 and Exchange 2010) and for your new public folders (Exchange 2013 and Exchange 2016). You can use the public folder insights to track your migration from legacy to new public folders.

For information about what permissions are required by the credentials for the data collection, see [Permissions needed for the Exchange public folders data source](#) on page 53.

To set up an Exchange public folders collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click **+** beside the name of the target environment.
- 4 Select the **Exchange Public Folders** check box.
- 5 Enable the data collection for this data source.
- 6 Select the types of data that you want to collect.
- 7 Under Targets, select **All Public Folders** or enter the folder paths for specific public folders.
- 8 Specify the type of public folders that you are collecting.
 - If you are collecting legacy public folders (Exchange 2007 and Exchange 2010), click **Legacy**.
 - a If Exchange 2007 is the highest Exchange legacy version, and you have a mixed Exchange 2003 and Exchange 2007 environment, enter the Exchange 2007 server name for a server that hosts public folders. (Not required for Exchange 2007 native organization.)
 - b If Exchange 2010 is the highest Exchange legacy version, enter the Exchange CAS server and specify the credential to be used to create the remote PowerShell connection.
 - If you are collecting public folders from Exchange 2013 or Exchange 2016, click **New**.
 - a Enter the Exchange CAS server and specify the credential to be used to create the remote PowerShell connection.

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.
- To use HTTPS instead of HTTP to connect to the Exchange 2010/2013/2016 CAS server, click **Show Advanced Settings** in the Credential section. This option requires that TLS/SSL is enabled for remote PowerShell on the Exchange CAS server.

The Advanced Settings also provide options if you have a customized PowerShell virtual folder, or if you want specify either an implicit credential or the Data Engine credential to be used to make the PowerShell connection.

- 9 Click **Save**.

Adding Exchange Online hybrid data sources for hybrid Office 365

If you have a hybrid Exchange Online (on-premise and Exchange Online) environment, there are three data sources that you configure to collect from your hybrid Office 365 environment.

- Exchange Online Hybrid User Configurations which retrieves user and distribution group data using LDAP and remote PowerShell.
- Exchange Online Hybrid Mailbox Configuration which retrieves Exchange Online mailbox configuration including statistics, permissions and mobile devices.
- Exchange Online Mailbox Contents which retrieves information about email traffic from Exchange Online user mailboxes using Exchange Web Services (EWS). You can collect from some or all mailboxes. See [Creating a data source for Exchange Online mailbox contents \(EWS\) data source](#) on page 71.

i | **NOTE:** The Exchange Online Mailbox Contents data source can be added to an Active Directory Forest / Office 365 (hybrid) target environment to collect from hybrid Exchange Online or to a native Office 365 target environment to collect from native Exchange Online.

About AD synchronization methods for hybrid Exchange Online

For hybrid Exchange Online environments, there are several different tools that are used to synchronize users and mailboxes between on-premise Active Directory and Azure Active Directory.

When you configure the Exchange Online Hybrid User Configuration or the Exchange Online Hybrid Mailbox Configuration data source, you must select the synchronization method that is deployed in your environment.

Essentially, all the different synchronization options use one of two methods to synchronize on-premise Active Directory to Azure Active Directory.


Table 22. Hybrid synchronization methods for on-premise Active Directory and Azure Active Directory.

Synch method	Synchronization tool and environment	On-premise AD attribute	Azure AD attribute
Method 1	Azure AD Connect (with Exchange 2016)	msds-ExternalDirectoryObjectId	ExternalDirectoryObjectId
Method 2	Azure AD Connect (without Exchange 2016) Azure AD Connect (upgraded from DirSync) Azure AD Sync (all environments) DirSync (all environments)	ObjectGUID	ImmutableId

Creating a data source for Exchange Online hybrid user configuration

The Exchange Online hybrid user configuration gathers information about users and distribution groups in a hybrid environment only.

To set up an Exchange Online hybrid user configuration collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click **+** beside the name of the target environment.
- 4 Select the **Exchange Online Hybrid User Configuration** check box.
- 5 Enable the data collection for this data source.
- 6 Select the types of data that you want to collect.
- 7 Specify Windows credentials to be used for the LDAP connection or select the **Use Data Engine service credential** option. See [Permissions needed for the Exchange Online hybrid user configuration data source](#) on page 54.
- 8 Specify the Office 365 credential to be used to collect the user configuration data through remote PowerShell.

To improve data collection performance you can provide multiple credentials.

i | **TIP:** Office 365 has a per-user throttling mechanism to protect the Exchange systems. It is recommended that you use dedicated user accounts on this page to avoid throttling caused by concurrent usage of the same account in other applications.

- 9 Select the tool that is used to synchronize your Azure Active Directory with on-premise Active Directory.
 - Azure AD Connect with Exchange 2016
 - Azure AD Connect
 - Azure AD Sync
 - DirSync

i | **NOTE:** If you use Azure AD Connect and have Exchange 2016 (mixed or native environment), select **Azure AD Connect with Exchange 2016** for optimal performance.

For information about synchronization methods, see [About AD synchronization methods for hybrid Exchange Online](#) on page 66.

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.
- To specify a specific domain controller instead of automatically discovering the domain controller, click **Show Advanced Settings** in the LDAP Collection Parameters section. The domain controller must be a global catalog server.

i | **NOTE:** If the Data Engine is in a different forest from the data to be collected, specify the domain controller you want to use for data collection. If you leave the default setting to automatically discover domain controller, the data collector selects a domain controller from the forest that hosts the Data Engine, not a domain controller from the forest from which you are collecting the data.


- 10 Click **Save**.

Creating a data source for Exchange Online hybrid mailbox configuration

For a hybrid Office 365 environment, you can create a mailbox configuration collection to gather mailbox statistics such as mailbox size, mailbox permissions, and remote devices from Exchange Online mailboxes. For a native Office 365 environment, see [Creating a data source for the Exchange Online native mailbox configuration data source](#) on page 70.

For more information about credential prerequisites, see [Permissions needed for the Exchange Online hybrid mailbox configuration data source](#) on page 54.

To set up an Exchange Online hybrid mailbox configuration collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click **+** beside the name of the target environment.
- 4 Select the **Exchange Online Hybrid Mailbox Configuration** check box.
- 5 Enable the data collection for this data source.
- 6 Select the types of data to collect:
 - Mailbox permissions
 - Mailbox (permissions)
 - Recipient permissions (Send As permissions)
 - Delegates (Send On Behalf Of permissions)
 - Mobile devices
 - Mailbox statistics
- 7 Specify the target mailboxes you want to collect:
 - All Exchange Online mailboxes
 - Specific Exchange Online mailboxes
For specific mailboxes, you can enter the common name or email address of a mailbox, a group, or a dynamic distribution group.
- 8 Specify the Windows credentials to be used for the LDAP connection or select the **Use Data Engine service credential** option.
- 9 Specify the Office 365 credential to be used to collect the Exchange Online mailbox configuration data through remote PowerShell.
To improve data collection performance you can provide multiple credentials.

i | **TIP:** Office 365 has a per-user throttling mechanism to protect the Exchange systems. It is recommended that you use dedicated user accounts on this page to avoid throttling caused by concurrent usage of the same account in other applications.

- 10 Select the tool that is used to synchronize your Azure Active Directory with on-premise Active Directory.
 - Azure AD Connect with Exchange 2016
 - Azure AD Connect
 - Azure AD Sync
 - DirSync

i | **NOTE:** If you use Azure AD Connect and have Exchange 2016 (mixed or native environment), select **Azure AD Connect with Exchange 2016** for optimal performance.

For information about synchronization methods, see [About AD synchronization methods for hybrid Exchange Online](#) on page 66.

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.
- To specify a specific domain controller instead of automatically discovering the domain controller, click **Show Advanced Settings** in the LDAP Collection Parameters section. The domain controller must be a global catalog server.

i **NOTE:** If the Data Engine is in a different forest from the data to be collected, you must specify the domain controller you want to use for data collection. If you leave the default setting to automatically discover domain controller, the data collector selects a domain controller from the forest that hosts the Data Engine, not a domain controller from the forest from which you are collecting the data.

11 Click **Save**.

Adding Exchange Online data sources for native Office 365

If you have a native Office 365 environment, there are three data sources that you configure to collect from Exchange Online using remote PowerShell:

- Exchange Online Native User Configurations which retrieves user and distribution group data using remote PowerShell.
- Exchange Online Native Mailbox Configuration which retrieves Exchange Online mailbox configuration including statistics, permissions and mobile devices.
- Exchange Online Mailbox Contents which retrieves information about email traffic from Exchange Online user mailboxes using Exchange Web Services (EWS). You can collect from some or all mailboxes. See [Creating a data source for Exchange Online mailbox contents \(EWS\) data source](#) on page 71.

i **NOTE:** The Exchange Online Mailbox Contents data source can be added to an Active Directory Forest / Office 365 (hybrid) target environment to collect from hybrid Exchange Online or to a native Office 365 target environment to collect from native Exchange Online.


Creating a data source for Exchange Online native user configuration

For a native Office 365 target environment, you can create an Exchange Online native user configuration collection to gather user and distribution group details from Exchange Online using remote PowerShell.

Exchange Online data sources do not display in the Data Collection page until you have added Office 365 as a target environment. For more information, see [Adding a target environment for native Office 365](#) on page 32.

For more information about credential prerequisites, see [Permissions needed for Exchange Online native user configuration data source](#) on page 55.

To set up an Exchange Online native user configuration collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.

- 3 Click **+** beside the name of the Office 365 target environment.
- 4 Select the **Exchange Online Native User Configuration** check box.
- 5 Click the **Exchange Online Native User Configuration** tile to open the configuration page.
- 6 Enable the data collection for this data source.
- 7 Select the data that you want to collect.
- 8 Specify the Office 365 credential used to collect the user configuration data through PowerShell.

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.
- 9 Click **Save**.


Creating a data source for the Exchange Online native mailbox configuration data source

For a native Office 365 target environment, you can create a mailbox configuration collection to gather mailbox statistics such as mailbox size, mailbox permissions, and remote devices from Exchange Online mailboxes using remote PowerShell. For hybrid Office 365 (Exchange Online and on-premise Exchange), see [Creating a data source for Exchange IIS Logs](#) on page 64.

Exchange Online data sources do not display in the Data Collection page unless you have added Office 365 as a target environment. For more information, see [Adding a target environment for native Office 365](#) on page 32.

For more information about credential prerequisites, see [Permissions needed for Exchange Online native mailbox configuration](#) on page 55.

To set up an Exchange Online native mailbox configuration collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click **+** beside the name of the Office 365 (native) target environment.
- 4 Select the **Exchange Online Native Mailbox Configuration** check box.
- 5 Click the **Exchange Online Native Mailbox Configuration** tile to open the configuration page.
- 6 Enable the data collection for this data source.
- 7 Select the types of data to collect:
 - Mailbox permissions
 - Mailbox (permissions)
 - Recipient permissions (Send As permissions)
 - Delegates (Send On Behalf Of permissions)
 - Mobile devices
 - Mailbox statistics
- 8 Specify the target mailboxes you want to collect:
 - All Exchange Online mailboxes
 - Specific Exchange Online mailboxes

For specific mailboxes, you can enter the common name or email address of a mailbox, a group, or a dynamic distribution group.

- 9 Specify the Office 365 credential to be used to collect the mailbox configuration data through remote PowerShell.

To improve data collection performance you can provide multiple credentials.

i | **TIP:** Office 365 has a per-user throttling mechanism to protect the Exchange systems. It is recommended that you use dedicated user accounts on this page to avoid throttling caused by concurrent usage of the same account in other applications.

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.

- 10 Click **Save**.


Creating a data source for Exchange Online mailbox contents (EWS) data source

For a native Office 365 or for an Active Directory/Office 365 hybrid target environment, you can create a mailbox contents collection to gather email traffic from Exchange Online user mailboxes using Exchange Web Services (EWS).

To improve data collection performance, you can provide multiple credentials when you configure the data source. You can also create different mailbox contents data collections that contain different groups of users.

For more information about credential prerequisites, see [Permissions needed for Exchange Online native mailbox configuration](#) on page 55.

To set up an Exchange Online mailbox contents data collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click **+** beside the name of an the Active Directory/Office 365 Hybrid or Office 365 (native) target environment.
- 4 Select the **Exchange Online Mailbox Contents** check box.
- 5 Click the **Exchange Online Mailbox Contents** tile to open the configuration page.
- 6 Enable the data collection for this data source.
- 7 Select the data that you want to collect.

If you specify Body, it is recommended that you limit the number of mailboxes for this specific data collection.

- 8 Specify the target mailboxes. You can enter a mailbox, group, or a dynamic distribution group in SMTP address format.

It is recommended that you add the Exchange Online users to a distribution group. You can then add the distribution group for the target mailboxes.

- 9 Specify the Office 365 credential used to collect the mailbox data

To improve data collection performance you can provide multiple credentials.

i | **TIP:** Office 365 has a per-user throttling mechanism to protect the Exchange systems. It is recommended that you use dedicated user accounts on this page to avoid throttling caused by concurrent usage of the same account in other applications.

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.

10 Click **Save**.

Setting chargeback costs for Exchange

You can specify the costs at which your Exchange email is charged. The calculated values appear in specific insights so you can determine the chargeback amounts for users or departments. You can use the chargeback insights to compare with external and internal billing.

You can specify the unit on which costs are calculated such per instance or per MB.

To set chargeback costs for Exchange


- 1 On the Admin Settings page, click the **Chargeback** tile.
- 2 Specify the currency that should be used in the Currency Symbol field.
- 3 Under Exchange - Base Charges, you can set chargeback costs for email messages and for mailboxes.
 - a Click the Exchange row for which you want to assign a charge type and a cost.
 - b Select the charge type (instance or MB).
 - c Enter the cost to be changed for the charge type.
- 4 When you have completed entering charge types and costs to be calculated, click **Save**.

Setting thresholds for Exchange metrics

You can set thresholds for Exchange metrics, using different colors to identify minimum and maximum values that you want to track. When you view the metrics in an insight table, if a threshold is reached or exceeded, the value shows the color that you specified. When you set a threshold, you set the color that displays when the number is over or under a specified value.

By default, the Thresholds Classification page lists key Lync QoE thresholds that can be set. You can add threshold classifications for Exchange ActiveSync, Exchange database, Exchange mailbox, Exchange message, Exchange DLP, and Exchange public folder metrics.

To add a new threshold classification

- 1 Click **Classifications** and click **Thresholds**.
- 2 Click **Add classification**. For a list of Exchange fields for which you can set thresholds, see [Table 23](#).
- 3 Beside the displayed field, click the down arrow  and select the field that you want from the list.
- 4 Specify the operator for the threshold:
 - greater than or equal to
 - greater than
 - less than or equal to
 - less than
- 5 Set the value for the threshold and the units of measure.
- 6 Select the color (red, yellow, or green) that should display when the threshold is met.

- 7 To add another threshold for the metric, click **Add threshold** and specify the value, units of measure, and color.

Table 23. Exchange metric thresholds that can be added.

Type	Threshold metric
Exchange ActiveSync Event	<ul style="list-style-type: none"> Elapsed Time Number of Email Attachments Downloaded Number of Email Attachments Uploaded Number of Email Attachments Transferred Number of Emails Downloaded Number of Emails Uploaded Number of Emails Transferred Response Code Sequence Number Size of Email Attachments Uploaded Size of Email Attachments Downloaded Size of Email Attachments Transferred Size of Emails Uploaded Size of Emails Downloaded Size of Emails Transferred Size of Items Uploaded Size of Items Downloaded Size of Items Transferred
Exchange Database	<ul style="list-style-type: none"> Available New Mailbox Space Deleted Item Retention Period Issue Warning Quota Log File Size Mailbox Retention Period Prohibit Send and Receive Quota Prohibit Send Quota Size
Exchange Database Copy	<ul style="list-style-type: none"> Copy Queue Length Replay Queue Length
Exchange DLP Match	<ul style="list-style-type: none"> Data Classification Confidence Data Classification Count
Exchange Email File Attachment	<ul style="list-style-type: none"> Size
Exchange Email Message	<ul style="list-style-type: none"> Size
Exchange Email Message Participant	<ul style="list-style-type: none"> Delivery Time First Response Time
Exchange Mailbox	<ul style="list-style-type: none"> Issue Warning Quota At Item Count Prohibit Send and Receive Quota At Prohibit Send Quota At Size
Exchange Legacy Public Folder	<ul style="list-style-type: none"> Age Limit Issue Warning Quota At Maximum Item Size Prohibit Post Quota At

Table 23. Exchange metric thresholds that can be added.

Type	Threshold metric
Exchange Legacy Public Folder Replica	<ul style="list-style-type: none"> • Associated Items Count • Associated Items Size • Deleted Items Count • Deleted Items Size • Item Count • Number of Contacts • Number of Owners • Size
Exchange Public Folders	<ul style="list-style-type: none"> • Age Limit • Associated Items Count • Associated Items Size • Deleted Item Retention Period • Deleted Items Count • Deleted Items Size • Issue Warning Quota At • Item Count • Maximum Item Size • Number of Contacts • Number of Owners • Prohibit Post Quota At • Size

Omitting words when filtering by subject or body

For insights that contain information about Exchange email messages, filters are available for Subject Keyword and Body Keyword. The keyword filters list the words that occur most often in the subject or in the body of messages.

Typically you do not want to include common words such as “the”, “of”, or “they”. To omit words from the subject or body keywords, you add them to the Stop Words list.

To add words to be omitted from keyword filters

- 1 Click **Queries** and click **Add Stop Words**.
The entity is Email Message.
- 2 Click the down arrow and select the appropriate field:
 - Subject Keywords
 - Body Keywords
- 3 Enter any additional words that should be omitted when ranking keywords in messages.

Adding data sources, chargeback costs, and thresholds for Skype for Business/Lync

- [Permissions needed to collect Skype for Business/Lync data](#)
- [Creating a data source for Skype for Business/Lync configuration](#)
- [Creating a data source for Skype for Business/Lync CDR](#)
- [Creating a data source for Skype for Business/Lync QoE](#)
- [Setting call classifications for Skype for Business/Lync](#)
- [Setting chargeback costs for Skype for Business/Lync](#)
- [Setting thresholds for Skype for Business/Lync metrics](#)

Permissions needed to collect Skype for Business/Lync data

To collect data from Skype for Business/Lync, you add different data sources to gather information. You can create the following data sources:

- **Skype for Business/Lync Configuration:** collect Skype for Business/Lync server configuration and user using PowerShell
- **Skype for Business/Lync CDR Database:** peer-to-peer session and conference details from a Skype for Business/Lync CDR (Call Detail Recording) SQL database using SQL queries
- **Skype for Business/Lync QoE Database:** Quality of Experience (QoE) information from a Skype for Business/Lync QoE SQL database using SQL queries

For each data source, you must specify the credential that is used to collect the data. In most collections, you have the option to use the credential that is specified for the Data Engine service. If you want to use that credential for your data collections, ensure it has the permissions specified for that data source.

Permissions needed for Skype for Business/Lync configuration data source

The account you specify for a Skype for Business/Lync configuration data source is used to collect the user, server, and pool configuration data from the Skype for Business/Lync servers. The account must have the CsViewOnlyAdministrator RBAC (role-based access control) role in the Skype for Business/Lync organization.

The easiest method to assign this role is to add the user to the CS View-Only Administrators build-in security group.

Permissions needed for Skype for Business/Lync CDR data source

You create a Skype for Business/Lync CDR database data source to collect usage information about peer-to-peer activities including instant messaging, Voice-over-Internet-Protocol (VoIP) calls, application sharing, file transfers, and conferences from the CDR database.

The account that is used to collect the Skype for Business/Lync CDR data must have the following database role membership in the CDR SQL databases:

- db_datareader

For UC Analytics to collect statistical conference and session data from Skype for Business/Lync server, the Monitoring role must be installed on the Skype for Business/Lync server and Call Detail Recording (CDR) must be enabled and running. For the steps to configure your Skype for Business/Lync server, see [Appendix B: Configuring the Skype for Business or Lync Server](#) on page 99.

Permissions needed for Skype for Business/Lync QoE data source

You create a Skype for Business/Lync QoE database data source to collect Quality of Experience (QoE) numeric data that tracks the quality of audio and video calls in your organization. QoE must be enabled and running on the Skype for Business/Lync server. For more information see [Enable Quality of Experience on Lync Server 2013](#).

The account that is used to collect Skype for Business/Lync QoE data must have the following database role membership on the QoE SQL database:

- db_datareader

Creating a data source for Skype for Business/Lync configuration

The general workflow in adding data sources and setting up data collection is similar for each target environment that you have create

When you configure a Skype for Business/Lync configuration data source, you specify the credential to be used to connect to the Skype for Business/Lync server using remote PowerShell. You collect the server, service, user, and pool configuration information directly from the Skype for Business/Lync server using remote PowerShell.


If you specify target users, the data collection will gather additional user configuration data such as:

- Effective user policies
- Audio/video enabled status
- Skype for Business/Lync user enabled status
- Public network enabled status

i **NOTE:** If you want to enter specific target users (rather than select all users or an organizational unit), you can enter a user, a group, or a dynamic distribution group. However, you cannot use the Domain Users group. For more information see [Can I enter the Domain Users group as the target for the data collection?](#) on page 59.

For more information about dynamic distribution groups, see [Using dynamic distribution groups to group target users](#) on page 58.

To set up a Skype for Business/Lync configuration collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click **+** beside the name of the target environment.
- 4 Select the **Skype for Business/Lync Configuration** check box.
- 5 Enable the data collection for this data source.
- 6 Select the types of data that you want to collect.
- 7 Specify the Windows credentials to be used for the LDAP connection or select the **Use Data Engine service credential** option.
- 8 Specify the target users for the data collection. You can select either:
 - All Lync or Skype for Business enabled users in all domains.
 - All Lync or Skype for Business enabled users in specified organizational units (OUs) or containers.
 - Specific Lync or Skype for Business enabled users.
- 9 For the PowerShell connection, specify the Skype for Business/Lync front-end server from which you are collecting data.
 - Click **Add server** and enter the Skype for Business/Lync front-end server name (and port number if applicable).
- 10 Specify the credentials to be used to connect through PowerShell.

The credential that is specified must have the appropriate permissions. For more information, see [Permissions needed for Skype for Business/Lync configuration data source](#) on page 75.

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.
 - If you want to specify a specific domain controller instead of automatically discovering the domain controller, click **Show Advanced Settings** in the LDAP Collection Parameters section. The domain controller must be a global catalog server.
 - i** | **NOTE:** If the Data Engine is in a different forest from the data to be collected, you must specify the domain controller you want to use for data collection. If you leave the default setting to automatically discover domain controller, the data collector selects a domain controller from the forest that hosts the Data Engine, not a domain controller from the forest from which you are collecting the data.
 - To use HTTPS instead of HTTP to connect to the Skype for Business/Lync front-end server, click **Show Advanced Settings** in the PowerShell Connection Parameters section. This option requires that TLS/SSL be enabled for remote PowerShell on the Skype for Business/Lync front-end server.

The PowerShell Advanced Settings also provide options if you have a customized PowerShell virtual folder, or if you want specify either an implicit credential or the Data Engine credential be used to make the PowerShell connection.
- 11 Click **Save**.


Creating a data source for Skype for Business/Lync CDR

You collect peer-to-peer session, enterprise voice, and conference data from the Skype for Business/Lync Call Details Recording (CDR) database using a SQL query. You can collect session and conference information such as:

- Start date and duration
- Participants
- Media types used
- Software clients

When you configure the Skype for Business/Lync CDR data source, you specify the CDR database and provide the credential needed to access the database. The credential must have the appropriate permissions. For more information, see [Permissions needed for Skype for Business/Lync CDR data source](#) on page 76.

To set up a Skype for Business/Lync CDR database collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click **+** beside the name of the target environment.
- 4 Select the **Skype for Business/Lync CDR Database** check box.
- 5 Enable the data collection for this data source.
- 6 Select the types of data that you want to collect.
- 7 Specify the CDR SQL databases from which you want to gather data.

If you are not using the default port for SQL and must specify a custom port number, enter the SQL database name (FQDN or host name) followed by a comma and the port number.

i | **NOTE:** You do not need to specify an instance when you enter an explicit port number. For example, you would specify UCAServer1,2014 to connect to the instance on port 2014.

- 8 Specify how chargeback costs are calculated if enterprise voice calls go through both a PSTN gateway and a Skype for Business/Lync mediation server by selecting either:
 - PSTN gateway
 - Skype for Business/Lync Mediation Server
- 9 Enter a credential to be used to access the CDR database. You can select one of the following:
 - Use the Data Engine service credential
 - Specify a Windows credential
 - Specify a SQL credential

Advanced settings (not required for most deployments)


- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.
- 10 Click **Save**.

Creating a data source for Skype for Business/Lync QoE

You collect Skype for Business/Lync Quality of Experience (QoE) data from the QoE database using a SQL query. Quality data is available for peer-to-peer sessions, conferences, and enterprise voice calls.

When you configure the Skype for Business/Lync QoE data source, you specify the QoE database and provide the credential needed to access the database. The credential must have the appropriate permissions. For more information, see [Permissions needed for Skype for Business/Lync QoE data source](#) on page 76.

To set up an Skype for Business/Lync QoE database collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click **+** beside the name of the target environment.
- 4 Select the **Skype for Business/Lync QoE Database** check box.
- 5 Enable the data collection for this data source.
- 6 Select the type of data that you want to collect.
- 7 Specify the QoE SQL databases from which you want to gather data.

If you are not using the default port for SQL and must specify a custom port number, enter the SQL database name (FQDN or host name) followed by a comma and the port number.

NOTE: You do not need to specify an instance when you enter an explicit port number. For example, you would specify UCAServer1,2014 to connect to the instance on port 2014.

- 8 Enter a credential to be used to access the QoE SQL database. You can select one of the following:
 - Use the Data Engine service credential
 - Specify a Windows credential
 - Specify a SQL credential

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.
- 9 Click **Save**.

Setting call classifications for Skype for Business/Lync

You can configure call classifications for your Skype for Business/Lync (such as local, long-distance, international, toll-free, and so on). Later, when you enter chargeback costs, you select the call classifications to which costs are assigned. For information about setting up chargeback costs, see [Setting chargeback costs for Skype for Business/Lync](#) on page 80.

For Skype for Business/Lync calls, you can configure call classifications for enterprise voice calls. You can define call classifications for outgoing calls, incoming calls, and internal peer-to-peer sessions.

About specifying Enterprise Voice call classifications

When you define call classifications for Skype for Business/Lync Enterprise Voice outgoing and incoming calls, you must specify the gateway or the mediation server that is being used for the call type. You can enter the FQDN or IP address or server name. Typically you might enter the IP address for a gateway or the server name for a mediation server.

- i** | **TIP:** If you encounter a problem trying to resolve the value that you entered for the mediation server or gateway server, check an insight that contains the gateway and mediation server information, such as the Skype for Business/Lync Peer-to-Peer Session Details insight. Enter the name or the IP address exactly as it appears in the insight.

You also enter phone number (usually a destination mask) and number of digits for different call classifications. For example, you can specify a phone number (or destination mask) for internal, local, toll-free and other types of calls. You can use one wild card (“*”) in the destination mask to match the pattern for the normalized destination phone numbers in Skype for Business/Lync.

To set classifications for Skype for Business/Lync incoming and outgoing calls

- 1 Click **Classifications** and click **Call Classifications**.
- 2 Click **Add Classification** for type of call that you want to define.
- 3 Enter a description for the call type.
- 4 Enter the mediation or gateway server name for the server used.
- 5 Enter the phone number associated with the call type. Since you can enter * wild cards, you usually would enter a destination mask for a phone number. For example, you might specify 1800* for toll free calls.
- 6 Enter the number of digits for the call type. This field is optional. If the call type can be uniquely identified through the phone number mask, you could leave this field empty.

Setting chargeback costs for Skype for Business/Lync

You can specify chargeback costs for Skype for Business/Lync peer-to-peer sessions, conferences, and voice calls. You can view the calculated values in the chargeback insights to determine the chargeback amounts for users or departments. You can use the chargeback insights to compare with external and internal billing.

Essentially you can set two different types of charges:

- base charges for infrastructure costs
- additional charges which include call minutes for each call classification

You can specify the unit on which peer-to-peer sessions and conferences costs are calculated such as per minute, per MB, per person, or per instance.

For Skype for Business/Lync enterprise voice calls and peer-to-peer session calls, you can also create the call classifications that are used to identify the different types of calls for chargeback. For more information about specifying call classifications, see [Setting call classifications for Skype for Business/Lync](#) on page 79.

To set chargeback costs for Skype for Business/Lync

- 1 On the Admin Settings page, click the **Chargeback** tile.
- 2 Specify the currency that should be used in the Currency Symbol field.
- 3 Click a Skype for Business/Lync item or media type for which you want to assign a charge type and a cost.

- 4 Under Skype for Business/Lync - Base Charges, you specify the chargeback for each media type for both peer-to-peer sessions and conferences.
 - a Click on the row for the media type you want.
 - b Select the charge type (per instance, per MB, per minute, per person, and so on).
 - c Enter the cost to be charged.
 - d Specify if rounding should be used to round the calculated cost up to the next minute (60 seconds) or set the rounding value to whatever you want.
- 5 Under Skype for Business/Lync - Additional Charges, you can set specific chargeback costs for the call classifications that you have defined (such as local, toll free, long distance, and so on).
 - a Specify the charge per minute for each call type.
 - b Specify if rounding should be used to round the calculated cost up to the next minute (60 seconds) or set the rounding value to whatever you want.
- 6 When you have completed entering charge types and costs to be calculated, click **Save**.

Setting thresholds for Skype for Business/Lync metrics

You can set thresholds for Skype for Business/Lync Quality of Experience (QoE) metrics, using different colors to show good calls or the severity of poor calls. When you set a threshold, you set the color that displays when the number is over or under a specified value. When you view a QoE insight, if a value shown in a table meets the set threshold, a colored underline appears below the value.

By default, the Thresholds Classification page lists key Lync QoE thresholds that can be set. You can also add thresholds for the conference and session metrics that are displayed in the Skype for Business/Lync insights. You can use color to identify minimum and maximum values that you want to track.

To set thresholds for the default Skype for Business/Lync QoE metrics

- 1 Click **Classifications** and click **Thresholds**.
- 2 Click the metric for which you want to set a threshold.

By default, you can set or modify threshold levels for several key QoE stream quality metrics. For more information, see [About the default Skype for Business/Lync quality metrics](#) on page 82.
- 3 Specify the operator for the threshold:
 - greater than or equal to
 - greater than
 - less than or equal to
 - less than
- 4 Set the value for the threshold and the units of measure.
- 5 Select the color (red, yellow, or green) that should display when the threshold is met.
- 6 To add another threshold, click **Add threshold** and specify the value, units of measure, and color.

About the default Skype for Business/Lync quality metrics

The default metrics that display on the Thresholds Classifications page are preset with the recommended thresholds for Skype for Business 2015 and for Lync 2013. Many of them also apply to Lync 2010. To add recommended thresholds that apply only to Lync 2010, you can add a threshold classification for the specific quality metric. For more information, see [Adding new threshold classifications](#) on page 83

The following table provides basic information about the quality metrics thresholds that are displayed by default:

Table 24. Default Skype for Business/Lync quality metrics and recommended thresholds.

Quality metric	Description	Possible thresholds
% of application sharing content lost	The percentage of the content from the sharer that did not reach the viewer. Content can be discarded (or spoiled) when the sharer discards tiles from the graphics source or when the ASMCU tiles discards tiles from sharer respectively. Available for Microsoft Lync Server 2013 and later only.	<ul style="list-style-type: none"> > 36% is bad.
% of call with high video CPU load	Percentage of the call where the client experienced high CPU load when processing video (dynamic capability flag was active). Available for Microsoft Lync Server 2013 and later only.	<ul style="list-style-type: none"> > 10% is bad
% of call with low video frame rate	The percentage of the call that is below the low frame rate threshold. Available for Microsoft Lync Server 2013 and later only.	<ul style="list-style-type: none"> > 10% is bad
% of local video frames lost	The percentage of the total video frames that are lost.	<ul style="list-style-type: none"> > 10% is bad
Average % of packets lost	Packet loss (%) represents the percentage of packets that did not make it to their destination. Packet loss will cause the audio to be distorted or missing (on the receiver end).	<ul style="list-style-type: none"> < 3% packet loss is considered good > 5% packet loss will affect audio > 7% packet loss is poor (+7% packet loss can be considered a major degradation of quality) > 10% is extremely bad > 50% packet loss - essentially no service
Average % of samples concealed	Average ratio of concealed samples generated by audio healing to typical samples.	<ul style="list-style-type: none"> < 2% is good > 3% is poor > 7% is bad
Average % of video packets lost with error correction	The packet loss rate after forward error correction (FEC) has been applied. Available for Microsoft Lync Server 2013 and later only.	<ul style="list-style-type: none"> > 10% is bad

Table 24. Default Skype for Business/Lync quality metrics and recommended thresholds.

Quality metric	Description	Possible thresholds
Average application sharing RDP latency	Average processing time for remote desktop protocol (RDP) tiles over the duration of the viewing session. A higher total equates to a longer delay in the viewing experience. Available for Microsoft Lync Server 2013 and later only.	<ul style="list-style-type: none"> > 400 msec is bad.
Average jitter	Jitter (ms) measures the variability of packet delay and results in a distorted or choppy audio experience. Jitter can increase latency on networks.	<p>Generally, jitter metrics can be qualified as follows:</p> <ul style="list-style-type: none"> < 20 ms is good > 30 ms is not good (but may be acceptable) > 45 ms is poor
Average network MOS degradation	Network Average Mean Opinion Score (MOS) is the key measurement used to gauge the perceived audio quality (based on an algorithm that calculates how a typical user would rate the voice quality). This metric shows the amount the Network MOS was reduced because of jitter and packet loss. It is an integer rating from 0 to 5.	<ul style="list-style-type: none"> 5 – excellent 4 – good 3 – fair 2 – poor 1 – bad
Average relative endpoint latency	Average amount of one-way latency between the two media end points involved in the application sharing. This is a single-hop latency measure. Relative one-way latency measures the delay between the client and the server. Available for Microsoft Lync Server 2013 and later only.S	<ul style="list-style-type: none"> > 1.75 ms is bad.
Average round trip latency	Network Round Trip Time (RTT) is the most common measure of latency and is measured in ms. This measure is the average round trip time for RTP packets between endpoints. When latency is high, users will likely hear the words, but there will be delays.	<p>For RTP packets as reported in the monitoring reports:</p> <ul style="list-style-type: none"> < 200 ms is good > 200 ms is poor > 500 ms is bad
Average video frame rate	The average video frame rate sent (outbound) during the call and the average video frame rate received (inbound) during the call. (frames/s)	<ul style="list-style-type: none"> A value of < 7 frames per second is considered poor video quality.
Average video frame rate used	Average frames per second received for all video streams and computed over the duration of the session. This metric is reported for video streams when available. (frames/s)	<ul style="list-style-type: none"> A value of < 7 frames per second is considered poor video quality.

Adding new threshold classifications

In addition to the default threshold metrics, you can set thresholds for additional metrics. You can add threshold classifications for more QoE metrics or for conference and peer-to-peer session metrics.

To add a new threshold classification

- 1 Click **Classifications** and click **Thresholds**.
- 2 Click **Add classification**.
- 3 Beside the displayed field, click the down arrow ▼ and select the field that you want from the list.
- 4 Specify the operator for the threshold:
 - greater than or equal to
 - greater than
 - less than or equal to
 - less than
- 5 Set the value for the threshold and the units of measure.
- 6 Select the color (red, yellow, or green) that should display when the threshold is met.
- 7 To add another threshold for the metric, click **Add threshold** and specify the value, units of measure, and color.

You can add threshold classifications for the quality metrics that are specific only to Lync 2010.

Table 25. Lync 2010-only quality metrics and recommended thresholds.

Quality metric	Description	Possible thresholds
Average estimated bandwidth	The available bandwidth estimated on the client-side. Absolute thresholds are not useful, but when the client detects bandwidth is low (< 100 kbps) audio quality can easily be affected by other applications or network congestion. In Lync 2010 only.	<ul style="list-style-type: none"> • < 100 kbps is bad
Average Listen MOS	The average predicted wideband listening MOS score for audio received from and sent to the network including speech level, noise level, codec, network conditions and capture device characteristics. Mean Opinion Score (MOS) is the gold standard measurement to gauge the perceived audio quality (an algorithm calculates how a typical user would rate the voice quality). In Lync 2010 only.	<ul style="list-style-type: none"> • 5 excellent • 4 good • 3 fair • 2 poor • 1 bad

In addition to adding the specific Lync 2010 recommended metric thresholds, you can add threshold classifications for a wide range of numeric values that are reported in the Skype for Business/Lync insights.

The following table lists the additional thresholds that you can add and set.

Table 26. Skype for Business/Lync threshold classifications that can be added.

Type	Threshold metric
QoE media session	<ul style="list-style-type: none"> • Average Conversation MOS
QoE session	<ul style="list-style-type: none"> • Device CPU Number of Cores • Device CPU Processor Speed

Table 26. Skype for Business/Lync threshold classifications that can be added.

Type	Threshold metric
QoE stream	<ul style="list-style-type: none"> • % of Call at CIF Resolution • % of Call at VGA Resolution • % of Call Competing For Network Resources • % of Call in Loss Congestion State • % of Call in Loss Congestion State From Delayed Packets • % of Call at HD720 Resolution • % of Error Correction Packets Used • % of Packets Dropped By Healer • % of Video Frames Lost • Average % of Samples Compressed • Average % of Samples Stretched • Average Application Sharing RDP Latency • Average Echo • Average Estimated Bandwidth • Average Listen MOS • Average Network MOS • Average Network MOS Degradation From Jitter • Average Network MOS Degradation From Packet Loss • Average Video Bandwidth • Average Video Bit Rate • Maximum % of Packets Lost • Maximum % of Samples Concealed • Maximum Estimated Bandwidth • Maximum Jitter • Maximum Network MOS Degradation • Maximum Round Trip Latency • Maximum Video Bit Rate • Minimum Estimated Bandwidth • Minimum Listen MOS • Minimum Network MOS
Conference participant	<ul style="list-style-type: none"> • (Media session) Duration • Number of IM Messages
Conference	<ul style="list-style-type: none"> • Duration
Peer-to-peer session	<ul style="list-style-type: none"> • Duration • Response Code
Peer-to-peer session participant	<ul style="list-style-type: none"> • Number of IM messages
Conferencing policy	<ul style="list-style-type: none"> • Maximum Application Sharing Bit Rate • Maximum Audio Bit Rate • Maximum File Transfer Bit Rate • Maximum Meeting Size • Maximum Video Bit Rate

Adding data sources, chargeback, and thresholds for Cisco

- [Permissions needed to collect Cisco data](#)
- [Creating a data source for Cisco configuration](#)
- [Creating a data source for Cisco CDR logs](#)
- [Setting call classifications for Cisco](#)
- [Setting chargeback costs for Cisco](#)
- [Setting thresholds for Cisco metrics](#)

Permissions needed to collect Cisco data

Cisco is supported only if Active Directory is present and Cisco end-users are synchronized to Active Directory users using the SAM account name.

To collect data from Cisco, you add different data sources to gather information. You can create the following data sources:

- **Cisco Configuration:** end-user data from the Cisco Unified Communications Manager (CUCM) and user data from Active Directory using LDAP
- **Cisco CDR Logs:** Peer-to-peer session and conference details from the Cisco CDR (call detail records) log files.

UC Analytics does not directly access your Cisco CDR database. You must copy the CDR logs to a share location from which you want to gather the data. For example, you can configure the Cisco Unified Communications Manager server to automatically upload the CDR logs to the directory you specify.

Permissions needed for the Cisco configuration data source

You specify the credentials that are used to collect user data from Active Directory and the credentials used to collect end-user data from the Cisco Unified Communications Manager (CUCM) server.

- The credential that is used to collect Active Directory (LDAP) user data must have read permissions on all the Active Directory user, group and contact objects.
- The credential that is used to collect Cisco end-user data (including phone number and SAM account) must have read permissions on the Cisco Unified Communications Manager server. The credential must be a member of user group Standard AXL API Access.

There are different methods that Active Directory (LDAP) can be integrated with Cisco. UC Analytics requires that the SAM account name be used to synchronize the Cisco end-users and Active Directory users.

Permissions needed for the Cisco CDR logs data source

To configure the Cisco CDR logs data source, you require the following credentials and permissions:


- You must specify a credential to access the Cisco server to collect Cisco configuration for the specified Cisco Unified Communications Manager (CUCM). The credential must be a member of user group Standard AXL API Access.
- To access the directories that contain the CDR logs, you can enter specific Windows credentials or you can use the credential that is specified for the Data Engine service. The credential used to access the Cisco CDR logs must have read rights on all the file shares on which you have stored the log files.

Creating a data source for Cisco configuration

By creating a Cisco configuration data collection, you can collect Cisco end-user information including Cisco phone numbers, device IDs, and the user information from Active Directory (LDAP).

For information about prerequisites, see [Permissions needed for the Cisco configuration data source](#) on page 86.

To set up a Cisco configuration collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click **+** beside the name of the target environment.
- 4 Select the **Cisco Configuration** check box.
- 5 Enable the data collection for this data source.
- 6 Select the types of data that you want to collect.
- 7 Specify the Windows credentials to be used for the LDAP connection to Active Directory or select the **Use Data Engine service credential** option.
- 8 Specify the Cisco Unified Communications Manager (Call Manager) server and credentials to collect Cisco data.

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.
- If you want to specify a specific domain controller instead of automatically discovering the domain controller, click **Show Advanced Settings** in the LDAP Collection Parameters section.

i **NOTE:** If the Data Engine is in a different forest from the data to be collected, you must specify the domain controller you want to use for data collection. If you leave the default setting to automatically discover domain controller, the data collector selects a domain controller from the forest that hosts the Data Engine, not a domain controller from the forest from which you are collecting the data.

- 9 Click **Save**.

Creating a data source for Cisco CDR logs

By creating a Cisco CDR log collection, you can gather data about Cisco voice-over-IP and other peer-to-peer sessions, ad hoc conferences, and MeetMe Conferences.


You gather the Cisco CDR (call detail recording) logs from one or more specified file shares. Typically, you set up a process to copy the log files from your Cisco server to a specified directory. For example, you could configure the Cisco Unified Communications Manager server to automatically upload the CDR logs to a directory you specify.

To configure a Cisco CDR log data collection, you must specify the following:

- the Cisco Unified Communications Manager (CUCM) server used to collect the Cisco configuration and the associated credential
- the locations of the CDR log file shares from which you want to gather data
- the credentials used to access the CDR logs

For information about prerequisites, see [Permissions needed for the Cisco CDR logs data source](#) on page 87.

To set up a Cisco CDR log collection

- 1 Click the gear icon  on the home page side bar.
- 2 Click **Data Collection**.
- 3 Click **+** beside the name of the target environment.
- 4 Select the **Cisco CDR Logs** check box.
- 5 Enable the data collection for this data source.
- 6 Select the types of data that you want to collect.
- 7 Specify the Cisco Unified Communications Manager (CUCM) server to which you want to connect and the port to be used.
- 8 Specify the Cisco credential to access the CUCM.
- 9 Specify each folder location of the directories from which the Cisco CDR logs are to be gathered.
- 10 Specify the credentials that are used to gather the CDR logs.

Advanced settings (not required for most deployments)

- If you want to change the minimum interval between the data collection runs, click **Show Advanced Settings** in the Scheduling section.
- 11 Click **Save**.

Setting call classifications for Cisco

You can configure call classifications for your Cisco calls (such as local, long-distance, international, toll-free, and so on). When you set up chargeback costs, you can select the call classifications to which costs are assigned. For information about setting up chargeback costs, see [Setting chargeback costs for Cisco](#) on page 89.

For Cisco calls, you can define call classifications for outgoing peer-to-peer sessions, incoming peer-to-peer sessions, and internal peer-to-peer sessions.

About specifying Cisco call classifications

When you define call classifications for Cisco outgoing and incoming calls, you must specify the endpoint that is being used for the call type. The endpoint is the port that is used by the call on the Cisco gateway. Gateways can have multiple endpoints dependent on the number of ports that each gateway contains.

You also enter phone number (usually a destination mask) and number of digits for different call classifications. For example, you can specify a phone number (or destination mask) for internal, local, toll-free and other types of calls. You can use one wild card (“*”) in the destination mask to match the pattern for the phone numbers in Cisco.

To set classifications for incoming and outgoing Cisco calls

- 1 Click **Classifications** and click **Call Classifications**.
- 2 Click **Add Classification** for type of call that you want to define.
- 3 Enter a description for the call type.
- 4 Enter the endpoint name for the Cisco endpoint used.
- 5 Enter the phone number associated with the call type. Since you can enter * wild cards, you usually would enter a destination mask for a phone number. For example, you might specify 1800* for toll free calls.
- 6 Enter the number of digits that are required for the targeted call. This field is optional. If the call type can be uniquely identified through the phone number mask, you could leave this field empty.

For example, a Toll Free 800 number is 11 digits so you could enter 11 digits for the Toll Free entry.

For both Lync and Cisco internal peer-to-peer sessions, you specify phone number information and the number of digits for both the caller and the callee.

Setting chargeback costs for Cisco

You can specify the costs at which your Cisco peer-to-peer calls and conference calls are charged. You can view the calculated values in the chargeback insights to determine the chargeback amounts for users or departments. You can use the chargeback insights to compare with external and internal billing.

For Cisco voice calls and peer-to-peer sessions, you can also create the call classifications that are used to identify the different types of calls for chargeback. For more information about specifying call classifications, see [Setting call classifications for Cisco](#) on page 88.

To set chargeback costs for Cisco

- 1 On the Admin Settings page, click the **Chargeback** tile.
- 2 Specify the currency that should be used in the Currency Symbol field.
- 3 Click a Cisco item or media type for which you want to assign a charge type and a cost.
- 4 Under Cisco - Base Charges, you specify the chargeback for each media type for both peer-to-peer sessions and conferences.
 - a Click on the row for the media type you want.
 - b Select the charge type (per instance, per MB, per minute, per person, and so on).
 - c Enter the cost to be charged.
 - d Specify if rounding should be used to round the calculated cost up to the next minute (60 seconds) or set the rounding value to whatever you want.
- 5 Under Cisco - Additional Charges, you can set specific chargeback costs for the call classifications that you have defined (such as local, toll free, long distance, and so on).
 - a Specify the charge per minute for each call type.

- b Specify if rounding should be used to round the calculated cost up to the next minute (60 seconds) or set the rounding value to whatever you want.
- 6 When you have completed entering charge types and costs to be calculated, click **Save**.

Setting thresholds for Cisco metrics

You can set thresholds for Cisco metrics, using different colors to identify minimum and maximum values that you want to track. When you set a threshold, you set the color that displays when the number is over or under a specified value. When you are viewing the metrics in an insight table, if the threshold value is reached, the metric displays in the color you specified.

By default, the Thresholds Classification page lists key Lync QoE thresholds that can be set. You can set thresholds for the following Cisco session and conference metrics:

- Cisco Peer-to-Peer Session - Duration
- Cisco Conference - Duration
- Cisco Conference Session - Duration

To add a new threshold classification

- 1 Click **Classifications** and click **Thresholds**.
- 2 Click **Add classification**.
- 3 Beside the displayed field, click the down arrow ▼ and select the field that you want from the list.
- 4 Specify the operator for the threshold:
 - greater than or equal to
 - greater than
 - less than or equal to
 - less than
- 5 Set the value for the threshold and the units of measure.
- 6 Select the color (red, yellow, or green) that should display when the threshold is met.
- 7 To add another threshold for the metric, click **Add threshold** and specify the value, units of measure, and color.

Configuring subscription global settings

- [What are insight subscriptions?](#)
- [Configuring settings for subscriptions](#)

What are insight subscriptions?

In addition to viewing insights on the Analytics web site, you can create subscriptions for insights. Subscriptions allow you to automatically export selected insights to a format such as PDF, Microsoft Word (.docx), HTML, MHTML, or to csv (raw or localized), or tsv (raw or localized) formats.

Each subscription is scheduled to automatically send the exported insight to specific email addresses or to a file share, such as a SharePoint folder.

i | **NOTE:** For information about the difference between raw and localized .csv file format, see the section titled “Can I export an insight?” in the *Unified Communications Analytics User Guide*.

Users can create different schedules for each subscription so that the subscription insights are exported and sent on a regular basis.

Configuring settings for subscriptions

Before users can create and send subscriptions, you must specify the SMTP relay server that is used to email the subscriptions. (If you are configuring subscriptions to post to file shares only, you do not have to specify a SMTP relay server.) If you want, you can also specify a custom From Address and display name that appears on the subscription emails as the sender. You must be a product administrator to access Admin Settings.


Prerequisites

Optionally, you can require authentication be used to access the SMTP server for email subscriptions. If you enabled authentication by checking the Authentication check box, the authentication credentials must have Sent As permissions for the Active Directory user associated with the email address that is in the Sender Address field.

The Sender Email Address is used to send SMTP email subscriptions and the Sender Display Name appears in the From field in the email.

i | **NOTE:** The Sender Display Name that you specify may not always be shown in the From field when a recipient views a subscription email in Outlook or through OWA. In some case, the recipient will see the user name from the matching user object in Active Directory. This is due to changes made in Exchange 2010 and later that affect how the From header is populated.

To specify configuration settings for all subscriptions

- 1 Click the gear icon  on the home page side bar to access Admin Settings.
- 2 Click the **Subscriptions** tile.
- 3 Ensure that **Enable subscriptions** is selected.
 - a For subscriptions where exported insights are sent by email, ensure that **Enable email as a target output** is selected.
 - b For subscriptions where exported insights are posted to a file share, ensure that **Enable file share as a target output** is selected.
- 4 Enter the SMTP mail server to be used to send subscription emails.
- 5 If required, specify a port for the SMTP server in the Port: field.

i **NOTE:** Since not all SMTP communication is done over port 25, you can specify the port number for the SMTP server. When the server name is displayed, the port number appears after a colon (:) at the end of the server name.

Advanced settings (not required for most deployments)

If you require authentication to access the SMTP server or if you want to change the From Address used as the sender in subscription emails, click **Show advanced settings**.

- 6 If authentication is required, select **Use authentication for SMTP server** and click **Set credential**.
 - a Enter the credentials that are to be used to access the SMTP server.
- 7 If TLS/SSL is required to connect to the SMTP server, select the **Use TLS/SSL** check box.
- 8 If you want to change the sender address and the display name that appears on subscription emails, click the default value of UCAnalytics@Quest.com.
 - a Enter the display name that should appear as the From name on email subscriptions.
 - b Enter the email address from which you want to send the email subscription.
- 9 Click **Save**.

Once you have configured the Admin Settings for subscriptions, users can create subscriptions for insights. For information about how to create and schedule subscriptions, see the *UC Analytics User Guide*.

What if I configure authenticated SMTP with an Exchange server that uses Integrated Windows authentication?

Suppose you are configuring subscriptions to use authenticated SMTP and the following conditions apply:

- The SMTP server is an Exchange server.
- The Exchange server is configured to authenticate SMTP connections with Integrated Windows authentication.

In this case, the specified credential must be from a domain that is trusted by the domain which contains the server hosting UC Analytics. Integrated Windows authentication will likely use Kerberos authentication. Windows will not allow UC Analytics to construct the required Kerberos token if the token is from a domain that the server does not trust.

As in all other cases of authenticated SMTP through an Exchange server, the credential must have SendAs permission on the Active Directory user to which the Sender Address belongs (if the address is not one of the credential's own email addresses).

Making changes to your deployment

- [Options available in the Deployment Manager](#)
- [Deploying additional Storage Engines](#)
- [Using the Tools menu for support activities](#)

Options available in the Deployment Manager

After you have installed UC Analytics, you can make changes to your deployment. Typically you use the UC Analytics.exe file to access the following options that can be used to manage your installation:

Table 27. Options available for managing the UC Analytics installation.

Option Name	Function
Uninstall All	Allows you remove all the installed services.
Manage Deployment	Allows you to update the properties for the installed services, add additional Storage Engines, add or remove product administrators, and modify customer feedback settings.
Manage Licenses	Shows a list of the installed licenses if you are running the DeploymentManager.exe on the server on which Data Engine service is installed. Though you can see only the licenses installed on this Data Engine server, if you add a new license, the new license is automatically installed on all your Data Engine servers if you have a multiple Data Engine deployment.
Tools	Allows you to set detailed logging and to create a support bundle that can be used by Quest Support to troubleshoot any issues. In some circumstances the TLS/SSL certificates for the Storage Engine and the Query Engine might become corrupted such as when the Storage Engine and Query Engine folders are not excluded from your antivirus scans. In this situation, you can use the Tools option to regenerate the certificates.

To make changes in your deployment

- 1 Start the installer by double-clicking the **UC Analytics.exe** file.
- OR -
From the Start menu, select **Programs | Quest UC Analytics | UC Analytics**.
- 2 Click **Manage Deployment**.
- 3 Select the option that you want:

Table 28. Options available for managing the UC Analytics deployment.

Option Name	Function
Change Product Administrators	Allows you to add users who have access to the Admin Settings. Product administrators can classify domains for message reporting, set user access for insights, and perform other administrative functions.
Modify Deployment	Allows you to move or update the properties for the installed services. You can also add more Storage Engines.
Change Customer Feedback Settings	Allows you to specify anonymous participation in the customer feedback program.

About the customer feedback program

If you participate in the customer feedback program, you help improve UC Analytics by allowing Quest to collect product usage telemetry. When you choose to participate in the program, data is automatically transmitted to Quest.

We only collect:

- Information about the hardware that hosts the Data Engine service.
- Anonymous information about which insights are used to help identify and prioritize future improvements.

We do NOT collect any personally identifying information.

Participation in the program is voluntary. Product functionality is not affected by participation. To opt out, clear the **Allow Quest to collect anonymous usage data** check box.

Deploying additional Storage Engines

When you install UC Analytics, you specify the locations on which the different services are to be installed. You can install multiple instances of the Storage Engine service.

At a later date, you might want to add more Storage Engines to distribute the collected data.

i | **NOTE:** If you must reboot the computer on which the Storage Engine service is installed and you have a larger database, there can be a period of 5 to 10 minutes before UC Analytics will be available.

To add more Storage Engines

- 1 Start the installer by double-clicking the **UC Analytics.exe** file.
- OR -
From the Start menu, select **Programs | Quest UC Analytics | UC Analytics**.
- 2 Click **Modify Deployment**.
- 3 Click **Add Server** and select the Storage Engine role.
- 4 Specify the server on which the Storage Engine service will be installed.
- 5 To change the default values for a server role, click **Edit Properties**.
- 6 Click **Advanced**, enter any changes to the properties for the selected server role and click **OK**.

Table 29. Storage Engine service properties.

Physical Memory (MB) Calculate automatically	Ensure the check box is selected to calculate whether sufficient memory is available. If you clear the check box, the value you enter in text box is assigned. NOTE: Physical memory is automatically assigned during installation. After you successfully install, if you display the properties dialog, the currently installed physical memory is displayed in the text box.
Query Port Number	If the port is already used by another application, change the port number. This port is used by the Query Engine service to access the Storage Engine service.

- 7 In the Storage Directory Path field, specify the directory path in which all the collected data is to be stored.

When you install additional Storage Engines, it can take several minutes to start and synchronize all the Storage Engine services before they are ready.

Using the Tools menu for support activities

If you have opened a case with Quest Support, you may be requested to enable detailed logging and then send a package of the log files for analysis. You can use the tools available through the UC Analytics to do this.

To enable the button that allows you to create the support package, you must run the UC Analytics.exe installer file from the installation directory (C:\Program Files\Quest\UC Analytics).

The Tools menu also provides an option that lets you regenerate TLS/SSL certificates for the Storage Engine and the Query Engine when the certificates have been corrupted.

To create a support package or regenerate certificates

- 1 Start the installer by double-clicking the **UC Analytics.exe** file.

- OR -

From the Start menu, select **Programs | Quest UC Analytics | UC Analytics | Deployment Manager.exe**.

- 2 Click **Tools**.
- 3 Select the option that you want:

Table 30. Deployment manager tool options.

Option Name	Function
Set Detailed Logging	Allows you to set detailed logging for the selected service. Detailed logging includes informational messages in addition to the default logging level which includes errors and warnings.

Table 30. Deployment manager tool options.

Option Name	Function
Create Support Package	<p>After you have run the product and reproduced the issue with detailed logging, you can use this option to create a zipped file that contains the log files.</p> <p>The zipped file, which is named UC Analytics Support.zip, is copied to the desktop.</p> <p>NOTE: In a distributed installation, you must run the Create Storage Package option on the server that has the Data Engine service installed.</p>
Regenerate Certificates	<p>Use this option to regenerate the TLS/SSL certificates used by the Storage Engine and Query Engine if they become corrupted. Keystore and Truststore certificates can become corrupted in certain circumstances, such as when the Storage Engine or Query Engine folders are not excluded from your antivirus software.</p>

Appendix A: Configuring Impersonation

Any credentials that are used to collect Exchange data must have application impersonation rights or “Exchange Impersonation” permissions to all the target mailboxes:

- For information about how to set Exchange impersonation for Exchange 2013 or Exchange 2016, see [Setting impersonation for Exchange 2013 or Exchange 2016](#) on page 97.
- For information about how to set Exchange impersonation for Exchange 2010, see the following Microsoft article: [Configuring Exchange Impersonation](#).
- For information about how to set Exchange impersonation for Exchange 2007, see the following Microsoft article: [Configuring Exchange Impersonation \(Exchange Web Services\)](#).
- For information about how to set impersonation for Office 365, see [Setting impersonation for Office 365](#) on page 98.

Setting impersonation for Exchange 2013 or Exchange 2016

To collect data from Exchange 2013 or Exchange 2016, the credentials used to collect Exchange data must have Application Impersonation rights for the mailboxes from which you are collecting.

To set impersonation for the collection credentials

- 1 Open the web-based Exchange Admin Center (EAC) for your Exchange server.
- 2 Click **Permissions**.
- 3 Under Admin Roles, click **+**.
- 4 Enter a name for the new role group, such as Impersonation.
- 5 Under Roles, click **+**.
- 6 Select **ApplicationImpersonation** and click **add->**.
- 7 Click **OK**.
- 8 Under Users, click **+**.
- 9 Select the account that will be used to collect the mailbox data from Exchange 2013 or Exchange 2016 and click **OK**.
- 10 Click **Save**.

Setting impersonation for Office 365

To collect data from Exchange Online (Office 365), the credentials used to collect Exchange data must have Application Impersonation rights, which means the account must be assigned to a Role-Based Access Control group that has Application Impersonation rights. By default, no groups have Application Impersonation rights in Office 365.

You must sign in as an administrator to the Office 365 portal and add this right either to an existing role group or to a new role group that you create.

For example, you could create a new role group named insight Impersonation and add the Application Impersonation right to the group. Role groups are created in the Role Groups page of the Office 365 portal.

To create a group with impersonation rights

- 1 Log into the Office 365 portal and go to the Outlook settings and protection page.
- 2 Select **Options | See All Options**.
- 3 Select **Manage Myself | My Organization**.
- 4 Select **Roles & Auditing** and click **New**.
- 5 For the Name, enter Analytics Impersonation (or another name if you want).
- 6 Under Roles click **Add**.
- 7 Select the **ApplicationImpersonation** role, click **Add** and **OK**.
- 8 Under Members click **Add**.
- 9 Select the names of the users that you want to add to the new role group, click **Add** and **OK**.
Ensure that the account you are using for data collection is added to this role.
- 10 Click **Save**.

Appendix B: Configuring the Skype for Business or Lync Server

For UC Analytics to collect statistical data from a Skype for Business or a Lync server, the Monitoring role must be installed on the Skype for Business or Lync server and Call Detail Recording (CDR) must be enabled and running.

See the following sections for the steps to configure your Lync or Skype for Business server.

[Configuring Lync Server 2010](#)

[Configuring Lync Server 2013 or Skype for Business 2015](#)

Configuring Lync Server 2010

Microsoft Lync Server gathers statistical data through the Monitoring role and from Active Directory. To collect data from Lync 2010 server, you must install and enable the Monitoring server role on your Lync server.

The procedures are divided into the following parts:

1. [Adding the SQL store for monitoring](#)
2. [Installing the Monitoring role](#)
3. [Enabling Call Detail Recording \(CDR\)](#)
4. [Starting the monitoring services](#)

1. Adding the SQL store for monitoring

As a prerequisite, you must have created a SQL instance which will be used to store the monitoring records. Then you can add your instance to the SQL store in Topology Builder.

To add the SQL store used for monitoring

- 1 In the Lync Topology Builder, select **SQL Store | New SQL Store**.
- 2 Enter the FQDN for the SQL Server and the SQL instance name (if you are not using the default instance).
- 3 Verify that the new store appears under the SQL stores folder.

2. Installing the Monitoring role

After you extract the .iso file on the Lync server on which you want to install the monitoring server role, you can run the appropriate .msi files to install the roles.

To install the monitoring server role on the Lync server

- 1 Navigate to the following file path:
 \Setup\amd64\Setup\monitoringserver.msi
- 2 Double-click the **monitoringserver.msi** file.
- 3 Follow the steps in the wizard to complete the installation.

3. Enabling Call Detail Recording (CDR)

For your Lync server, you enable call detail recording (CDR) by setting the global properties for the associated forest. CDRs are logs of usage statistics from conferences, instant messaging, and phone sessions that take place across your Lync servers.

To enable CDR on your Lync server

- 1 In the Start menu, select **Lync Server Control Panel**.
- 2 Select **Monitoring and Archiving** in the list on the left side of the panel.
- 3 Click the **Call Detail Recording** tab.
- 4 Double-click **Global**.
- 5 Select the following check box:
 - **Enable monitoring of call detail recordings (CDRs)**

4. Starting the monitoring services

After you have enabled and configured monitoring, you must start the monitoring services.

To start services on your Lync server

- 1 In the Start menu, select **Lync Server Control Panel**.
- 2 Select **Topology** in the list on the left side of the panel.
- 3 Click the **Status** tab and select the appropriate Lync server.
- 4 Select **Action | Start all services**.

You can verify that the monitoring services have started using the Services MMC tool or using Services Manager.

Configuring Lync Server 2013 or Skype for Business 2015

As of Lync Server 2013, the Monitoring role and the Archiving role no longer exist as separate roles. Both the monitoring and archiving services are collocated on each Front-End server. In Skype for Business Server 2015, monitoring is enabled or disabled on a pool-by-pool basis.

The procedures are divided into the following parts:

1. [Associating the store with the Front-End pool](#)
2. [Updating the Lync or Skype for Business Server](#)
3. [Enabling and configuring monitoring](#)

Prerequisite

As a prerequisite, you must have created a SQL instance which will be used to store the monitoring records. During configuration you will associate the SQL instance with the Front-End Lync server on which the monitoring services will run.

1. Associating the store with the Front-End pool

You must associate a monitoring store (database) with the Front-End pool. A single monitoring store can be associated with multiple pools.

The monitoring store is used to collect call detail recording (CDR). Call detail recording tracks the usage of Lync server activities such as Voice over IP (VoIP) phone calls; instant messaging (IM); file transfers; audio/video (A/V) conferencing; and application sharing sessions.

In this procedure, the Standard Edition Server is referenced. The same procedure is also used for the Enterprise version.

To associate the store with the Front-End pool and publish the topology

- 1 Open the Topology Builder.
- 2 Select **Standard Edition Front-End** and select the Lync or Skype for Business server.
- 3 Right-click and select **Edit Properties**.
- 4 Select the **Monitoring (CDR and QoE metrics)** check box and click **New**.
- 5 In the Define New SQL Server Store dialog, enter the FQDN for the SQL Server and the SQL instance name (if you are not using the default instance).
- 6 Click **OK**.
- 7 Review SQL server information for the Monitoring database and click **OK**.
- 8 Select the Lync server, right-click and select **Topology**.
- 9 Select **Publish**.

2. Updating the Lync or Skype for Business Server

Now you must update the server to include the monitoring store information. From the installation path or from your application DVD, run the **Setup.exe** file. (For a 64-bit server, the file path would be Setup | AMD64 | Setup.exe.)

To update the Lync Server

- 1 In the Deployment Wizard, select **Install or Update Lync Server System**.
The wizard now updates the Front-End server to include the changes that you made in the Topology Builder.

To update the Skype for Business Server

- 1 In the Deployment Wizard, click **Install or Update Skype for Business Server System**.

- 2 On the Deploy page, under Step 2: Setup or Remove Skype for Business Server Components, click **Run Again**.
- 3 In the Setup Skype for Business Server components wizard, on the Setup Skype for Business Server components page, click **Next**.
- 4 On the Specify path to MSIs page, type the path to the file Ocscore.msi (a file included with your Skype for Business Server installation media) and click **Next**.
- 5 On the Executing Commands page, click **Finish**.

3. Enabling and configuring monitoring

Now you can enable and configure monitoring (CDR).

To enable monitoring on the Lync or Skype for Business server

- 1 Open the Lync or Skype for Business Server Control Panel and select **Monitoring and Archiving** in the left panel.
- 2 In the Call Detail Recording tab, double-click the **Global** default policy and ensure that **Enable Monitoring of CDRs** is selected.
- 3 Now you must enable CDR:
 - For Lync, click **Commit**.
 - For Skype for Business, click the appropriate site from the table, click **Action**, and then click **Enable CDR**.

4. Starting the monitoring services

After you have enabled and configured monitoring, you must start the monitoring services.

To start services on your server

- 1 In the Start menu, select **Lync Server Control Panel** or select the **Skype for Business Server Control Panel**.
- 1 Select **Topology** in the list on the left side of the panel.
- 2 Click the **Status** tab and select the appropriate server.
- 3 Select **Action | Start all services**.

If the services are already running, stop and then start the services. You can verify that the monitoring services have started using the Services MMC tool or using Services Manager.

Appendix C: Configuring IIS Log Files to capture ActiveSync events

- [Configuring IIS Logging on the Exchange CAS and Mailbox servers](#)
- [Configuring IIS Logging if Exchange is hosted on Windows Server 2008 or later](#)
- [What ActiveSync events are collected and displayed in the insights?](#)

Configuring IIS Logging on the Exchange CAS and Mailbox servers

For UC Analytics to collect data from the ActiveSync IIS log files, IIS logging must be configured on the front-end Exchange Client Access Server (CAS) and on the back-end Exchange Mailbox servers to include the required data. You configure IIS logging on the servers that have the ActiveSync role installed.

UC Analytics only supports the following log file format:

- W3C Extended Log File Format (set by default)

UC Analytics does not support a weekly, monthly, or yearly frequency for IIS log file rollover. Only hourly or daily log file rollover is supported.

When using W3C Extended Log File Format, you must configure the extended properties. The default configuration of IIS logging for the W3C Extended Log File Format is insufficient for UC Analytics to gather the information required for insights.

Use IIS Manager to configure the W3C Extended Log File Format on the Exchange CAS server from which you want to gather data. If the Exchange CAS server is hosted on Window Server 2008, Windows Server 2008 R2, or Windows Server 2012, see [Configuring IIS Logging if Exchange is hosted on Windows Server 2008 or later](#) on page 104.

To configure IIS logging on Exchange hosted on Windows 2003 Server

If you do not have the Internet Information Services (IIS) Manager MMC snap-in added, start at step 1.

If the snap-in is already installed, start at step 7.

- 1 From the Start menu, select **Run**.
- 2 Type **mmc**, and click **OK**.
- 3 In the Console dialog box, select **File | Add/Remove Snap-in**.
- 4 Click **Add**.
- 5 Select **Internet Information Services** and click **Add**.
- 6 Click **Close** and click **OK**.

- 7 In the treeview, browse to the Web Sites node.
- 8 Right-click the web site that you want and select **Properties**.
- 9 On the Web Site property sheet, click **Properties** in the Enable Logging section.
- 10 Select the **Advanced Properties** tab.
- 11 Ensure that the following fields are selected:
 - Date (date)
 - Time (Time)
 - User Name (cs-username)
 - Server name (s-computername)
 - Server IP Address (s-ip)
 - URI Stem (cs-uri-stem)
 - URI Query (cs-uri-query)
 - Protocol Status (sc-status)
 - Bytes Sent (sc-bytes)
 - Bytes Received (cs-bytes)
 - Time Taken (time-taken)
 - User Agent (cs(User-Agent))
 - Protocol Substatus (sc-substatus)
- 12 Click **OK**.
- 13 Under the General tab, ensure that the New Log Schedule is set to hourly or daily. UC Analytics does not support weekly, monthly, or unlimited log file size for the log file rollover schedule.
- 14 Click **OK** and close the console.

Configuring IIS Logging if Exchange is hosted on Windows Server 2008 or later

If you have Exchange hosted on Windows Server 2008 or Windows Server 2008 R2 (IIS 7), or Windows Server 2012 (IIS 8), you must install the IIS Management Scripts and Tools on the Exchange CAS server. The IIS Management Scripts and Tools are required to allow UC Analytics to gather the IIS log files.

You must also configure the W3C Extended Log File Format for IIS logging on the Exchange server.

To configure IIS logging on Exchange hosted on Windows 2008 Server or later

- 1 Click **Start** and select **Administrative Tools | Internet Information Services (IIS) Manager**.
- OR -
In the Server Manager, select the **Web Server (IIS)** and open the **Internet Information Services (IIS) Manager**.
- 2 On the left, open the dropdown menu under the Start Page option.
- 3 Select the default Web Site.
- 4 Double-click **Logging**.

- 5 On the Logging page, in the Log file section set the Format to **W3C** and click **Select Fields**.
- 6 In the W3C Logging Fields dialog box, ensure the following fields are selected:
 - Date (date)
 - Time (time)
 - User Name (cs-username)
 - Server name (s-computername)
 - Server IP Address (s-ip)
 - URI Stem (cs-uri-stem)
 - URI Query (cs-uri-query)
 - Protocol Status (sc-status)
 - Bytes Sent (sc-bytes)
 - Bytes Received (cs-bytes)
 - Time Taken (time-taken)
 - User Agent (cs(User-Agent))
 - Protocol Substatus (sc-substatus)
- 7 Click **OK**.
- 8 If you are running IIS 8.5 on Windows Server 2012 R2 or later, ensure that the Log Event Destination is set one of the following:
 - Log file only
 - OR -
 - Both log file and ETW event
- 9 Under the Log File Rollover section, ensure that a schedule of hourly or daily is selected. UC Analytics does not support weekly, monthly, or yearly log file rollover time periods.

What ActiveSync events are collected and displayed in the insights?

If you have configured your Exchange CAS servers as described in this appendix, you can view data about different types of ActiveSync events in the insights.


To see the event type in an insight, click the details browser icon  in the top right corner of the main insight. The following events are collected and displayed in the ActiveSync insights.

Table 31. ActiveSync event types that are collected and reported.

Event Type	Description
Sync	Synchronizes the changes in a folder between the client and the server.
SendMail	Sends mail to the server. This command is issued in the HTTP POST command's URI and does not contain an XML body. The body contains the MIME message instead.
SmartForward	Forwards a message object without retrieving the full message object from the server.
SmartReply	Replies to a message object without retrieving the full message object from the server.

Table 31. ActiveSync event types that are collected and reported.

Event Type	Description
GetAttachment	Retrieves an email attachment from the server.
FolderSync	Synchronizes the folder hierarchy but does not synchronize the items in the folders.
FolderCreate	Creates an email, calendar, or contacts folder on the server.
FolderDelete	Deletes a folder from the server.
FolderUpdate	Moves a folder from one location to another location on the server and is used to rename folders.
MoveItems	Moves items from one folder to another.
GetItemEstimate	Gets an estimate of the number of items in a folder that is synchronized.
MeetingResponse	Used to accept, tentatively accept, or decline a meeting request in the user's Inbox folder.
Search	Finds and retrieves information about contacts and recipients in the Global Address List (GAL).
Settings	Supports getting and setting global properties, such as Out-of-Office (OOO) and device information.
Ping	Requests that the server monitor all specified folders for changes that require the client to resynchronize.
ItemOperations	Identifies the body of a request or response as containing a set of commands operating on items.
Provision	Gets the security policy settings set by the server administrator, such as the user's minimum password length requirement.
ResolveRecipients	Resolves a list of supplied recipients and, optionally, fetches their S/MIME certificates so that clients can send encrypted messages.
ValidateCert	Validates a certificate that has been received through an S/MIME mail.
HealthCheck	Internal access from Microsoft Exchange to check the health of the ActiveSync service.

The ActiveSync data collector skips any IIS log entries for ActiveSync events that are missing critical fields. The collector writes partial information for ActiveSync events that are missing non-critical fields.

Table 32. Critical and non-critical fields in ActiveSync event entries.

Critical Fields	Non-Critical Fields
date	sc-bytes
time	cs-bytes
cs-uri-stem	cs-username
cs-uri-query	s-computername
sc-status	s-ip
	time-taken
	cs(User-Agent)
	sc-substatus

If you see error or warnings in the data collection status for Exchange IIS Logs jobs that indicate there are missing fields, ensure that all of the recommended W3C logging fields were selected when you configured IIS logging on your Exchange CAS server.

For more information, see [Configuring IIS Logging on the Exchange CAS and Mailbox servers](#) on page 103 and [Configuring IIS Logging if Exchange is hosted on Windows Server 2008 or later](#) on page 104.

Appendix D: PowerShell cmdlets used by data sources

This section contains information about the PowerShell cmdlets that are used by the different data sources to collect data that is used in the insights.

[Exchange configuration data source](#)

[Exchange public folder data source](#)

[Skype for Business/Lync configuration data source](#)

[Exchange Online hybrid and native user configuration data sources](#)

[Exchange Online hybrid and native mailbox configuration data sources](#)

[Exchange Online mailbox contents data source](#)

[Office 365 user subscription configuration data source](#)

Exchange configuration data source

The Exchange configuration data source uses the following PowerShell cmdlets when collecting Exchange configuration data:

- `Get-ExchangeServer`
- `Get-MailboxDatabase -status`
- `Get-PublicFolderDatabase -status`
- `Get-MailboxDatabaseCopyStatus -Server` (Exchange 2010/Exchange 2013/Exchange 2016)
- `Get-Mailbox`
- `Get-MailboxStatistics`
- `Get-MobileDeviceStatistics` (Exchange 2013/Exchange 2016)
- `Get-MobileDevice` (Exchange 2013/Exchange 2016)
- `Get-ActiveSyncDeviceStatistics` (Exchange 2010/Exchange 2007)
- `Get-ActiveSyncDevice` (Exchange 2010)
- `Get-MailboxPermission`

Exchange public folder data source

The Exchange public folders data source uses the following PowerShell cmdlets when collecting Exchange public folder data:

- Get-PublicFolderDatabase
- Get-PublicFolder
- Get-PublicFolderStatistics

Skype for Business/Lync configuration data source

The Skype for Business/Lync configuration data source uses the following PowerShell cmdlets when collecting Lync configuration data:

- Get-CsArchivingPolicy
- Get-CsConferencingPolicy
- Get-CsExternalAccessPolicy
- Get-CsComputer
- Get-CsPool
- Get-CsService
- Get-CsUser

Exchange Online hybrid and native user configuration data sources

Both the Exchange Online hybrid and native user configuration data sources use the following PowerShell cmdlets when collecting Exchange user configuration data from Office 365:

- Get-Mailbox
- Get-Group
- Get-DynamicDistributionGroup
- Get-Contact
- Get-User
- Connect-MsolService
- Get-MsolUser

Exchange Online hybrid and native mailbox configuration data sources

Both the Exchange Online hybrid and native mailbox configuration data sources use the following PowerShell cmdlets when collecting Exchange mailbox configuration data from Office 365:

- Get-Mailbox
- Get-Group
- Get-DynamicDistributionGroup
- Get-User
- Get-Recipient
- Get-MailboxStatistics
- Get-MailboxPermission
- Get-RecipientPermission
- Get-MobileDeviceStatistics
- Connect-MsolService
- Get-MsolUser

Exchange Online mailbox contents data source

The Exchange Online mailbox contents data source uses the following PowerShell cmdlets when collecting Exchange mailbox contents data from Office 365:

- Get-Group
- Get-DynamicDistributionGroup
- Get-User
- Get-Recipient

Office 365 user subscription configuration data source

- Connect-MsolService
- Get-MsolUser

Appendix E: Custom configurations and backup and recovery options

This section contains information about how to handle specific situations. See the following topics for instructions about how to configure UC Analytics to meet certain requirements.

- [Setting up a multi-forest environment with a one-way trust](#)
- [Firewall configuration: ports for data collection](#)
- [Backing up and restoring your data using scripts](#)
- [Scheduling the backup batch file to run automatically](#)
- [Performing a manual backup of the storage folder before upgrade](#)
- [Moving your storage location](#)
- [Recommendations for disaster recovery](#)
- [Setting a custom title page for exported or subscription insights](#)
- [Changing the interval time before job status is purged](#)

Setting up a multi-forest environment with a one-way trust

You can configure UC Analytics to support multiple forests that have a one-way trust. In the following example of a one-way trust scenario, the forests are configured as follows:

- Forest A - Admin forest contains Active Directory users
- Forest R - Resource forest (all mailboxes are here).
- Forest R trusts Forest A.

UC Analytics is installed in the resource forest (Forest R). The UC Analytics Data Engine service is running with Forest R credentials. For this scenario, you would perform the following steps to allow users from the trusted domain (Forest A) to access UC Analytics:

- 1 Add a second target environment that points to the trusted domain.
- 2 Specify an account from the trusted domain as the Authentication Credential that is used to authenticate a user (User X) in Forest A and allows access to UC Analytics.
 - The credential must have read rights to the Active Directory forest specified in the Forest A environment.
 - The credential must have sufficient rights to browse users and groups and to resolve group memberships for all users and groups from this environment that are specified in the Security settings.

Once the second environment (for Forest A) is added and an authentication account from Forest A is specified for authentication, User X can log on to UC Analytics. For information about adding target environment for an additional Active Directory forest, see [Adding multiple Active Directory forests](#) on page 30.

User X can view the data in the insights based on the security permissions that were granted to the ForestA\Group in the Security settings. When you add users in the Security settings, ensure that the Forest A target environment is selected when you add the Forest A users. For information about the Security settings, see [Granting users access to data](#) on page 41.

Firewall configuration: ports for data collection

When you configure your firewall, UC Analytics requires access to the following TCP and UDP ports for Exchange 2010 and later:

Table 33. Ports required for data collection

Network access	Required ports
Microsoft file sharing SMB (Server Message Block)	TCP/UDP 135-139
Direct-hosted SMB traffic	TCP/UDP 445
LDAP (Lightweight Directory Access Protocol) Active Directory Lookups	TCP/389
Kerberos	TCP/UDP 88
LDAP Global Catalog	TCP/3268
DNS (Domain Name System)	TCP/UDP 53
WINS (Windows Internet Name Service)	TCP/UDP 137
Remote PowerShell for Exchange	TCP Port 80 and access to the IIS PowerShell virtual directory on your Exchange servers
Remote PowerShell for Skype for Business/Lync	TCP Port 443 and access to the IIS OcsPowerShell virtual directory on your Skype for Business/ Lync Servers

If you have implemented an intelligent firewall, you can configure the firewall to allow the specific types of network access (SMB, LDAP, and so on) that UC Analytics requires instead of access to specific ports.

Ports used for service-to-service communication in a distributed installation

In a distributed installation where you have a multiple UC Analytics servers deployed and there are firewalls between the UC Analytics servers, certain ports are required for service-to-service communication.

- For the Storage Engine and the Query Engine, open the following TCP ports: 1223, 7100, 7101, 7299, 9042, 9260, and 10099.
- For the Data Engine, open the following TCP ports: 1336 and 1337.

Backing up and restoring your data using scripts

UC Analytics provides two scripts (batch files) that allow you to back up data from your UC Analytics Storage Engine. The backup batch file uses the “snapshot” command that is included in the Cassandra nodetool utility.

When you run the backup script, you take a snapshot of all your stored data while the Storage Engine service is running. When you take a snapshot, you are backing up your data at a specific point in time.

NOTE: This will temporarily double the amount of data in your Storage directory while the batch file runs.

Supported scenarios

About the backup and restore batch files

Step 1: [Edit the backup.bat file parameters for your installation](#)

Step 2: [Run the backup batch file](#)

Step 3: [Edit the restore batch file](#)

Step 4a: [Restore the storage data in an existing installation](#)

Step 4b: [Restore the storage data to a new installation](#)

For information about creating a task to back up your data automatically, see [Scheduling the backup batch file to run automatically](#) on page 116.

Supported scenarios

The following scenarios are supported by the backup and restore scripts:

- You can back up an existing installation and restore to the same installation
- You can back up an existing installation and restore to a new installation. The Data Engine and Query Engine roles installed in a new installation can be different than the original installation. However, the new installation must meet the following criteria:
 - have the same number of installed Storage Engines as the old installation
 - use the same time zone settings on all Storage Engines as the old installation. UC Analytics performs the backup using a date-related label.
 - have the same server name and IP address as the old installation in which the backup was created.

For example, any of the following installations could be backed up and interchangeably restored:

Scenario 1

Server A: All components - Web Site, Data Engine (insights role), Query Engine (insights role), Data Engine (collector role), Query Engine (collector role), Storage Engine

Server B: Storage Engine

Scenario 2

Server A: Web Site, Data Engine (insights role), Query Engine (insights role), Storage Engine

Server B: Data Engine (collector role), Query Engine (collector role), Storage Engine

Scenario 3

Server A: Web Site, Data Engine (insights role), Data Engine (collector role), Storage Engine

Server B: Query Engine (insights role), Query Engine (collector role), Storage Engine

What is backed up by the backup script?

All your configuration and collected data is backed up when you run the backup script. When you run the restore script, the configuration and collected data are restored to the state when backup was created.

The backup process saves data from all the tables in the `\Storage\data\Doradus` folder and creates a separate subdirectory for each saved table. Each snapshot directory contains numerous files of data captured at the time of the snapshot. The backup process also saves configuration data from the `\Storage\data\System` folder.

The batch file then copies the snapshot data created to the backup folder (named based on the current date and time), then removes the snapshot directories from the original Storage folder.

About the backup and restore batch files

In the UC Analytics installation path, there is a folder named Storage Engine Backup Scripts. By default, the directory is located in the following path:

```
C:\Program Files\Quest\UC Analytics\Storage Engine Backup Scripts
```

The folder contains two batch files:

- backup.bat
- restore.bat

The folder also contains several PowerShell script files:

- BackupScript.ps1
- PurgeScript.ps1
- RestoreScript.ps1

To use the scripts, copy all the files to the folder on each server that hosts the UC Analytics Storage Engine. If you have multiple storage nodes, you must run the script on each node. All nodes must be in the same time zone.

Prerequisites

To run the backup and restore batch files, you must have PowerShell 3.0 or later installed on the server that hosts the UC Analytics Storage Engine.

To run the backup.bat file to create the backup, all the UC Analytic services must be running.

For deployments with multiple storage nodes, the scripts must be run on each node. All nodes must be within the same time zone.

The prerequisites for running the restore.bat differ depending whether you are restoring to the existing installation or to a new installation.

Step 1: Edit the backup.bat file parameters for your installation

After you have copied all the backup and restore files to the server that hosts the UC Analytics Storage Engine, you must edit the backup.bat file to specify the parameters for your environment.

The backup.bat file looks as follows:

```
-storageEngineDirectoryPath "C:\Program Files\Quest\UC Analytics\Storage Engine\bin" -storageDirectoryPath "C:\Program Files\Quest\UC Analytics\Storage" -backupDirectoryPath "c:\UCA\backup" -backupLabel "label" -portNumber 7299 -numberOfBackupCopies 2 -force
```

The parameters that you can edit are as follows:

Table 34. Backup batch file parameters

Parameter	Default value	Required information
storageEngineDirectoryPath	<i>C:\Program Files\Quest\UC Analytics\Storage Engine\bin</i>	Shows the location of the storage engine. If you installed using the default values, you can leave this parameter as shown.
storageDirectoryPath	<i>C:\Program Files\Quest\UC Analytics\Storage</i>	Shows the location of the storage folder. If you installed using the default values, you can leave this parameter as shown.
backupDirectoryPath	<i>c:\UCA\backup</i>	Enter the path and folder name for the folder to which you want to copy the backup files. You can enter a UNC path if you are copying the backup files to a different server or share.
backupLabel	<i>label</i>	Enter a unique name to identify the storage folder and server that contains the Storage Engine data that you want to back up. For example, you could enter Store1Server1.
portNumber	7299	Port used by the Cassandra nodetool.
numberOfBackupCopies	2	Enter the number of backup copies that will be retained at any one time. The backup folders are date-stamped so if you are retaining two copies, the two newest folders are kept and any older copies are removed.

NOTE: Do not change the port number from 7299. This is the port that is used by the Cassandra nodetool.

Step 2: Run the backup batch file

Once you have edited the backup.bat file to contain your local values, you can run the file to create a backup copy of your data. The backup script automatically detects the path to the latest version of Java that is installed on the computer.

You run the backup batch file locally on the server on which the Storage Engine resides.

To create a snapshot backup file

- 1 Ensure that all the UC Analytics services are running.
- 2 Navigate to the folder that contains the backup and restore script files.
- 3 Open the command prompt and enter:

```
backup.bat
```

The batch file creates a folder with the current date in YYYY-MM-DD format that contains subfolders with files that contain the data structure information and files that contain the actual data.

Step 3: Edit the restore batch file

Before you restore a backup file, you must edit the restore.bat file to specify the parameters for your environment. The restore.bat file looks as follows:

```
-backupDirectoryPath "c:\UCA\backup" -backupLabel "label" -backupDate "YYYY-MM-DD" -storageDirectoryPath "C:\Program Files\Quest\UC Analytics\Storage" -storageEngineDirectoryPath "C:\Program Files\Quest\UC Analytics\Storage Engine\bin" -portNumber 7299
```

The parameters that you can edit are as follows:

Table 35. Restore batch file parameters

Parameter	Default value	Required information
backupDirectoryPath	C:\UCA\backup	Enter the path and folder name for the folder to which the backup batch file previously copied the backup files.
backupLabel	label	Enter the unique name that was assigned as the backup label when the backup was created. The label name must exactly match the name that was entered in the backup.bat file when the backup was created.
backupDate	YYYY-MM-DD	Enter the date for the backup file from which you want to restore the data. The backup folder is named with the date of the backup in YYYY-MM-DD format.
storageDirectoryPath	C:\Program Files\Quest\UC Analytics\Storage	Shows the location of the storage folder. If you installed using the default values, you can leave this parameter as shown.
storageEngineDirectoryPath	C:\Program Files\Quest\UC Analytics\Storage Engine\bin	Shows the location of the storage engine. If you installed using the default values, you can leave this parameter as shown.
portNumber	7299	Port used by the Cassandra nodetool.

NOTE: Do not change the port number from 7299. This is the port that is used by the Cassandra nodetool.

Step 4a: Restore the storage data in an existing installation

Use this procedure to restore your existing database if the data has become corrupted.

When you restore data from a snapshot, there is intensive processor and I/O activity on the node that is being restored. You run the backup batch file locally on the server on which the Storage Engine resides.

To restore a snapshot backup file in the same installation

- 1 Stop the following two UC Analytics services on all the servers in the deployment in the following order:
 - Data Engine
 - Query Engine
- 2 Navigate to the folder that contains the backup and restore script files.
- 3 Open the command prompt and enter:


```
restore.bat
```
- 4 Restart all the UC Analytics services in the following order:
 - a Storage Engine.
 - b Query Engine
 - c Data Engine

Step 4b: Restore the storage data to a new installation

Use this procedure to restore a database to a new installation.

When you restore data from a snapshot, there is intensive processor and I/O activity on the node that is being restored. You run the backup batch file locally on the server on which the Storage Engine resides.

To restore a snapshot backup file in a new installation

- 1 Install UC Analytics. For information about the criteria that the new installation must meet, see [Supported scenarios](#) on page 112.

If you have a distributed installation, you must specify the same number of Storage Engines as were installed when the backup was performed.

- 2 Stop the UC Analytics services on all the servers in the deployment in the following order:

- Data Engine
- Query Engine
- Storage Engine

- 3 Navigate to the folder that contains the backup and restore script files.

- 4 Open the command prompt and enter:

```
restore.bat
```

- 5 Restart all the UC Analytics services in the following order:

- a Storage Engine.
- b Query Engine
- c Data Engine

Scheduling the backup batch file to run automatically

You can use the Windows task scheduler to run the backup batch file automatically on a regular basis. For information about configuring the backup and restore batch files, see [Backing up and restoring your data using scripts](#) on page 112.

i | IMPORTANT: The restore.bat file should only be run manually.

To set up a weekly schedule for backup

- 1 Start the Windows Task Scheduler.
- 2 Find and click the task folder in the console tree under which you want to create a folder.
 - a If you want to create a sub-folder in the root task folder, click the **Task Scheduler Library** folder.
 - b In the Actions pane, click **New Folder**.
 - c In the *Enter name of the new folder* dialog box, type the name of folder.
 - d Click **OK**.
- 3 In the Actions Pane, click **Create Task**.
- 4 Step through the details to establish the new schedule with the following settings:

- On the Triggers tab, set the schedule a trigger for once a week.
- Point the scheduled task to the backup.bat file provided by UC Analytics.
- Ensure that you configure the task to run at the highest level.

Using the -force parameter

In the backup.bat file, the parameter -force is useful if you have scheduled the batch file to run automatically as part of a scheduled task. If a backup file of the same name is found in the backup directory path, -force causes the script to overwrite the file without issuing a prompt saying that a duplicate file name has been detected.

If you are running the script manually, you might want to remove the -force parameter.

Performing a manual backup of the storage folder before upgrade

Before you upgrade UC Analytics, it is recommended that you perform a manual backup of your storage folder. If you have multiple Storage Engines deployed, you must take a backup of each storage folder on each server.

To perform a manual backup of the storage folder

- 1 Stop the following UC Analytics services:
 - UC Analytics Query Engine
 - UC Analytics Storage Engine
- 2 Navigate to the UC Analytics storage folder. By default, the directory is located in the following path:
C:\Program Files*<VendorName>*\UC Analytics\Storage
- 3 Use a compression utility to compress the Storage folder, such as to a zipped file format.
- 4 Copy the compressed file to a separate location.

NOTE: It may take several hours to complete the manual copy.

This type of backup copy is different from the backup script. The backup script copies only the collected data, data source configuration, and Admin settings. The backup script does not copy the Cassandra system files. To back up all of the database, you must copy the entire storage folder.

Moving your storage location

You can move the location of your UC Analytics storage database within the same computer. When you move the database, you must also change a number of places that point to the database.

To move the storage directory and configure the cassandra.yaml file for the new location

- 1 Stop the following UC Analytics services:
 - UC Analytics Query Engine
 - UC Analytics Storage Engine
- 2 Navigate to the UC Analytics installation folder. By default, the installation directory is located in the following path:
C:\Program Files\Quest\UC Analytics\

- 3 Move the directory named Storage to the new location.
- 4 Navigate to the following directory:
C:\Program Files\Quest\UC Analytics\Storage Engine\conf
- 5 Open the following text file for editing:
cassandra.yaml
- 6 Locate the following text in the file:
C:\Program Files\Quest\UC Analytics\Storage
You will find this text in three locations related to the data, the commitlog, and the saved_caches. Note that this path name uses forward slashes, not back slashes.
- 7 In the three locations, change the path name to the new location you selected in [Step 3](#).
- 8 Locate the following file in the Configuration directory:
C:\Program Files\Quest\UC Analytics\Configuration\DeploymentConfiguration.txt
- 9 In the Configuration.txt file, locate the following parameter:
"Settings":{"Key":"StorageEnginePhysicalMemory","Value":"8192"},{"Key":"StorageDirectoryPath","Value":
"C:\\Program Files\\Quest\\UC Analytics\\Storage"},
- 10 Change {"Key":"StorageDirectoryPath","Value":"**C:\\Program Files\\Quest\\UC Analytics\\Storage**"} to reflect the new path for the storage location.
- 11 Start the Query Engine and Storage Engine services.

i **TIP:** If User Account Control (UAC) is on when you perform the procedure, UAC file virtualization might be active when you try to update DeploymentConfiguration.txt and cassandra.yaml files. As a result, the actual files in Program Files are not updated. In this case, you might have to temporarily disable UAC to update the files.

Recommendations for disaster recovery

For disaster recovery planning, there are two options that you can consider.

- Use redundant storage such as a RAID array when you install UC Analytics.
- Back up your configuration and data storage.

To back up the UC Analytics configuration and data

- 1 Stop all the UC Analytics services:
 - UC Analytics Data Engine
 - UC Analytics Query Engine
 - UC Analytics Storage Engine
- 2 On each computer on which the Storage Engine is installed, back up your storage folder. By default, the storage directory is located in the following path:
C:\Program Files\Quest\UC Analytics\Storage
- 3 Start all the UC Analytics services.

To restore the UC Analytics configuration and data

- 1 Stop all the UC Analytics services:

- UC Analytics Data Engine
 - UC Analytics Query Engine
 - UC Analytics Storage Engine
- 2 On each computer on which the Storage Engine is installed, restore your storage folder. By default, the storage directory is located in the following path:
C:\Program Files\Quest\UC Analytics\Storage
 - 3 Start all the UC Analytics services.

Setting a custom title page for exported or subscription insights

When you export an insight (or create an insight subscription) that uses PDF or .docx format, the exported insight displays UC Analytics and the Quest logo on the title page. You can customize the title page to include your own logo and company name by editing a file called UC.Analytics.Insights.DataEngine.InsightLogic.dll.config.

The UC.Analytics.Insights.DataEngine.InsightLogic.dll.config file contains the following text:

```
<?xml version="1.0" encoding="utf-8" ?>
<configuration>
<configSections>
    <section name="insightRenderingConfiguration"
type="UC.Analytics.Common.General.Core.Configuration.SettingsNameValueCollection,
UC.Analytics.Common.General.Core"/>
</configSections>
<insightRenderingConfiguration>
    <add name="exportLogoImagePath" value=""/>
    <add name="exportTitle" value=""/>
</insightRenderingConfiguration>
</configuration>
```

When you edit the file and enter your own values, the characters that you enter in the value fields must use valid XML characters. For certain characters, such as double quotation marks (“), ampersand (&), and apostrophe (’), you must escape the characters to have them render correctly on the title page.

Table 36. Using XML special characters

Name	Character	How entered in file
Double quotation marks	“	"
Apostrophe	’	'
Ampersand	&	&
Less than	<	<
Greater than	>	>

To customize the title page for insight exports or subscriptions

- 1 On the server that hosts the UC Analytics Data Engine, navigate to the folder in which the Data Engine is installed. By default, the Data Engine folder is located at:
C:\Program Files\Quest\UC Analytics\Data Engine

- 2 Create a backup copy of the UC.Analytics.Insights.DataEngine.InsightLogic.dll.config file and save it to a different location.
- 3 Using a text editor, open the UC.Analytics.Insights.DataEngine.InsightLogic.dll.config file.
Suppose that you want to show the company name **Sitraka & Co.** as the title and also include a graphic file that is located at C:\Users\Admin\My Pictures\CompanyLogo.png.
- 4 You would edit the <add key> values as shown in the following example:

```
<?xml version="1.0" encoding="utf-8" ?>
<configuration>
  <configSections>
    <section name="insightRenderingConfiguration"
type="UC.Analytics.Common.General.Core.Configuration.SettingsNameValueCollection,
UC.Analytics.Common.General.Core"/>
  </configSections>
<insightRenderingConfiguration>
  <add name="exportLogoImagePath" value="C:\Users\Admin\My
Pictures\CompanyLogo.png"/>
  <add name="exportTitle" value="Sitraka & Co."/>
</insightRenderingConfiguration>
</configuration>
```

- 5 Save your changes.
- 6 For the changes to take effect, restart the Data Engine service.

Changing the interval time before job status is purged

By default, UC Analytics purges the job status information for a data collection job 30 days after the job has run. Depending on the size and number of your data collections, you might want to modify the time period that UC Analytics will keep job status details.

To change the time period for which UC Analytics will retain job run status details, you must edit the UC.Analytics.Insights.DataEngine.DataCollector.dll.config file. The configuration file is located in the data engine directory.

- 1 On the server that hosts the UC Analytics Data Engine, navigate to the folder in which the Data Engine is installed. By default, the Data Engine folder is located at:
C:\Program Files\Quest\UC Analytics\Data Engine
- 2 Create a backup copy of the UC.Analytics.Insights.DataEngine.DataCollector.dll.config file and save it to a different location.
- 3 Using a text editor, open the UC.Analytics.Insights.DataEngine.DataCollector.dll.config file.
- 4 Locate the following section and change the 30 (beside retentionPeriodInDays=) to the number of days that you want retain the job status data:

```
<purgeJobStatusJobSettings
  enabled="true"
  executionAlignmentTimeInMinutes="60"
  executionFrequencyInMinutes="1440"
```

```
retentionPeriodInDays="30"
```

```
/>
```

- 5 Save your changes.
- 6 For the changes to take effect, restart the Data Engine service.

Appendix F:

Questions and answers about data source configuration

- [How often do the data collections actually gather data and when do they run?](#)
- [Why did an insight show no data for a 30-day range though I initially set the data source to collect 30 days back?](#)
- [If I collect both Exchange Tracking Logs and Exchange Mailbox Contents, are there duplicate items?](#)
- [If I collect message data only from the Exchange tracking logs, is the message Send Date and delivery time available?](#)
- [What are the differences between the Exchange Mailbox Contents and Exchange Tracking Logs data sources?](#)
- [Why do I get an error when collecting Exchange configuration from multiple Exchange versions?](#)
- [Why do I have to specify domain name when doing a multi-forest collection?](#)

Introduction

This section includes common questions and provides detailed answers to help you understand and troubleshoot your UC Analytics deployment.

How often do the data collections actually gather data and when do they run?

Does a data collection only collect data once per UTC day? When I set the frequency for how often a collection will run, does the counter start at midnight (UTC/local) or does it start when the Data Engine service starts?

Answer

All user-created data sources (except for Mailbox Contents) run aligned with UTC midnight and, when you add the data source and start the Data Engine, at the interval that you specified in the data source. You can set the execution interval for all data sources except Mailbox Contents. The default interval is different for each data source, but all are more frequent than once a day.

The Configuration data sources (Exchange/Lync/Cisco/Domain Controllers) and the Exchange Public Folder data source are unique in that they only collect data once a day (first run in the UTC day). Even if you scheduled these collections to run more than once a day, they take a “snapshot” of the data only once per UTC day. You might schedule the data sources to run multiple times per day to handle any situation when data collection fails.

For example, suppose you add an Exchange Configuration data source, schedule it to run every 6 hours, and you created the data source at 8:00 am UTC. The data source job will run (in UTC) as follows:

- 8:00 am (it runs when scheduled and collects data)
- 12:00 pm (this the next multiple of 6 hours after midnight – no data is collected because it is the second run of the day)
- 6:00 pm (no data collected)
- 12:00 am (collects data since first run of the day)
- 6:00 am (no data collected)

If an Exchange Configuration data source job is scheduled every 24 hours, and was created at 8:00 a.m. (UTC), then the data source will run (in UTC) as follows:

- 8:00 a.m.
- 12:00 a.m.
- 12:00 a.m.

If a data source job is still running during its next scheduled execution time, the job execution is skipped.

The Mailbox Contents does not have a scheduled interval nor is it aligned on UTC midnight. The Mailbox Contents collection job runs continuously, with a minimum interval between the job execution start times. The start time is unpredictable. If the job execution takes less time than the minimum interval, the job will start X minutes (minimum interval) after the previous job execution started. Otherwise, job execution starts immediately after the previous execution completes.

Why did an insight show no data for a 30-day range though I initially set the data source to collect 30 days back?

I have just installed UC Analytics. Why does the Lync Organizational Summary User Activity insight show no data when I use a custom date filter for 4 Jan 2015 - 10 Jan 2015? I set my initial data source collection to go back 30 days and I started it on January 30.

Answer

There are two data source collections that provide information for the Lync Organizational Summary User Activity insight:

- Domain Controller (AD)
- Lync Configuration

When you started your Domain Controller data collection on Jan 30, the initial data collection period was set to 30 days (default). The Domain Controller data source collects information from AD and retroactively creates AD user “snapshots” in the database for each user. This ensures that there is a user object to link to historical data from data sources that collect historical information such as the Exchange Tracking Logs, Lync Users, and so on.

The Lync Configuration data source collection, however, collects data using PowerShell and only creates objects for the current day. The data only covers the date on which the Lync Configuration data collection accessed PowerShell to gather the data.

Therefore, on Jan 30, after the data source collections had finished, you had 30 days of AD user data and but only one day of Lync configuration data (Jan 30). The AD user snapshots from Jan 30 are linked to the Lync information, but user snapshots from Jan 29, 28, and earlier are not linked since Lync configuration data does not exist in the database for those dates.

The Lync Organizational Summary User Activity insight queries the AD users that are Lync-enabled (users that have a link to collected Lync data). When you specify “today”, the time range that the query looks for is “from midnight today until midnight tomorrow”.

When you specify an explicit date range such as Jan 1 - Jan 10, the query looks for the AD user snapshots that have links to Lync data from Jan 1 until end of Jan 10 (midnight of Jan 11) and takes the latest snapshot from this period, which is the Jan 10 snapshot. If you apply the rule about how UC Analytics collects historical data to the specified date range - the only linked AD-Lync data exists for Jan 30, not for Jan 10. For this reason, the insight shows no data.

After you have had UC Analytics running for a few months, you will have a few months of Lync configuration data accumulated and can use filters to view historical data in the Lync Organizational Summary User Activity insight.

This same scenario applies to other insights that use either the Lync Configuration or the Exchange Configuration data source collection to provide information such as:

- Lync User Adoption
- Lync Archiving Policies - Inventory
- Mailboxes - Inventory
- Mailboxes - Permissions inventory
- Mailboxes – Inactive

Since the Lync Configuration or the Exchange Configuration data source collections only gather data for the current day, even though the Domain Controller (AD) data source collection initially contained 30 days of historical user data, there would be only one day of data for the linked AD user - Lync/Exchange configuration data.

Over time, as you run the Lync Configuration and the Exchange Configuration data source collections on an ongoing basis, you would gradually have more days available of linked user Lync/Exchange configuration data. You can then use custom date range filters on the insights.

If I collect both Exchange Tracking Logs and Exchange Mailbox Contents, are there duplicate items?

I have configured both of the following data source collections:

- Exchange Tracking Logs
- Exchange Mailbox Contents (all mailboxes)

You can see information for the same mailboxes in the Exchange Tracking Logs data collection and in the Exchange Mailbox Contents data collection. If you run both data source collections, are there duplicate items in the UC Analytics Exchange insights?

Answer

No. The Tracking Log and Mailbox Contents data collections both collect email messages from Exchange mailboxes. However, if you collect the same messages using both types of data collections the message objects are not duplicated.

Determining which data collections you want to run is dependent on the information that you want to get from the messages. If you want a faster data collection, or if you need DLP / journaling messages, set up an Exchange Tracking Logs data collection. If you want information about response time or email attachments, and are willing to wait longer for the data collection, set up the Exchange Mailbox Contents data collection.

If I collect message data only from the Exchange tracking logs, is the message Send Date and delivery time available?

If you are running only the Exchange Tracking Logs data collection but not the Exchange Mailbox Contents data collection, the insights will contain the Send Date and delivery time for both MAPI and SMTP submitted messages. However, in some special circumstances, there are a couple of limitations:

- If an SMTP message is both undeliverable and has no internal Exchange server hops, the Send Date is unavailable and no delivery time is calculated.
- For MAPI and SMTP messages with at least one recipient in a moderated distribution group, the displayed Send Date can be the date when the message was approved by the moderator (and delivery was resumed to the group recipients) instead of the date when the message was originally sent. The delivery time is calculated from the time the message resumed its transit. When this occurs, for any recipients who had the message delivered directly to their mailbox before the moderator approved it, a received time can display that is before the Send Date.

If you are also running Exchange Mailbox Contents data collection, the Send Date is always obtained from the Date field of the message.

What are the differences between the Exchange Mailbox Contents and Exchange Tracking Logs data sources?

Both the Exchange Mailbox Contents (EWS) and the Exchange Tracking Logs data sources can be configured to collect Exchange messages. What are the differences in the information that the two data collections gather?

Answer

The following table shows the differences in email message statistics between the two data sources that collect Exchange email message information:

- Exchange Mailbox Content (EWS) data source
- Exchange Tracking Log data source

Table 37. How Exchange message data is collected by different UC Analytics data sources.

What gets collected	Exchange Mailbox Content (EWS) data source	Exchange Tracking Log data source
Exchange user messages from the specified data source targets	You select Exchange target mailboxes by specifying groups (distribution groups or security groups) or individual users. Only the messages from selected mailboxes are collected.	You specify the target tracking log folders on your Exchange servers. Only the messages that transferred through the specified servers are collected.

Table 37. How Exchange message data is collected by different UC Analytics data sources.

What gets collected	Exchange Mailbox Content (EWS) data source	Exchange Tracking Log data source
<p>Message history</p> <p>Collection of message history is limited by</p> <ul style="list-style-type: none"> • data retention time • start date specified for the data collection. 	<p>Only the messages within the “deleted mailbox retention period” are collected. For Exchange 2007, messages in the Exchange dumpster (deleted messages) are not collected.</p> <p>By default, a 30 day maximum message history is set in the data collection configure file. To collect more history, you must manually update the configuration file.</p>	<p>Only the messages that fall within the tracking log maximum age (MessageTrackingLogMaxAge) as set in the Exchange server configuration are collected.</p>
<p>System messages include:</p> <ul style="list-style-type: none"> • legacy public folder synchronization messages • health mailbox testing messages • system attendant messages • system monitoring messages 	<p>Are not collected.</p>	<p>Are collected and are marked as “system messages”.</p>
<p>Journaling messages</p>	<p>Are not collected unless the journaling mailbox is configured as a data source target. In this case, the messages are not marked as journaling messages.</p>	<p>Are collected and are marked as Journaling messages.</p>
<p>SMTP messages</p>	<p>Are collected if they are sent to one of the target mailboxes.</p> <p>Are not collected if they are sent to non-target mailboxes or to external recipients.</p>	<p>Are collected, but the “message sent time” for inbound messages is approximate. Delivery time is set only after the Exchange Calculation job has run.</p>
<p>Messages from outside organizations</p>	<p>Will be collected if they are sent to one of the configured target mailboxes.</p>	<p>Will be collected, but the “message sent time” inbound messages is approximate. Delivery time is set only after the Exchange Calculation job has run.</p>
<p>Messages sent to outside organizations</p>	<p>Are collected if they are sent from one of the target mailboxes. The “message receive time” and “message delivery time” are not set.</p>	<p>Are collected. The “message receive time” and “message delivery time” are not set.</p>
<p>Missing message properties</p>	<p>None</p>	<p>Importance</p> <p>Sensitivity</p> <p>Encryption</p> <p>Conversation</p> <p>Body</p> <p>Response Time</p>

Table 37. How Exchange message data is collected by different UC Analytics data sources.

What gets collected	Exchange Mailbox Content (EWS) data source	Exchange Tracking Log data source
Missing message participant properties	None NOTE: If a message is collected through EWS and the message was sent from a mailbox user to a distribution group to which the user belongs, the user is counted only as a sender, not a recipient. If the tracking log collection was also run, that user is also counted as a recipient.	Sender On Behave Of To, CC,BCC (shows message recipients but not whether the To, CC, or BCC box was used) Sent Time Of Day Was Sent After Hours Was Received After Hour In Reply To
Message size	Message size in mailbox. The size in sender's mailbox takes precedence.	Message size in transport. The earliest transport size takes precedence.

Why do I get an error when collecting Exchange configuration from multiple Exchange versions?

I am collecting from Exchange 2007, Exchange 2010, and Office 365 hybrid. For the Exchange configuration data source collection, I get the following message with job type Exchange Configuration for the mailboxes that are stored in Office 365:

“Unable to detect the Exchange Mailbox version or Exchange configuration data source does not support the target mailbox version.”

I've verified that Impersonation is granted in Office 365, and I checked the PowerShell connection parameters. Why I am getting this message?

Answer

You must create separate Exchange configuration data collections for Exchange 2007 and for Exchange 2010/Office 365 hybrid.

If you only added an Exchange configuration data source for Exchange 2007, you will see that the Exchange 2010/2013 option is greyed out so it is not being used. Create another Exchange configuration data collection for Exchange 2010 and that will collect the mailboxes from Exchange 2010 and Office 365 hybrid as well.

Also in the Exchange 2007 Exchange configuration data source, remove the server name that you added under Exchange 2010/2013.

After you have configured multiple Exchange configuration data sources, you can use the Rename option to uniquely identify each data collection. For information about how rename a data source, see [Renaming a data source](#) on page 39.

Why do I have to specify domain name when doing a multi-forest collection?

During a multi-forest collection for the Exchange configuration data source, the Common Name or Email Address would not work when specifying the Target Mailboxes for the collection. We had to put in DomainName\User or DomainName\Group

The domain name is the netbios name of the other forest from which I was collecting.

Answer

The key to making the multi-forest configurations work is you cannot use the “Automatically discover domain controller” for data sources in the secondary forests. This setting is selected by default and is hidden. To access the setting, you must click the Show Advanced Settings link in the LDAP Connection Parameters section. Specify a specific domain controller.

We are more than just a name

We are on a quest to make your information technology work harder for you. That is why we build community-driven software solutions that help you spend less time on IT administration and more time on business innovation. We help you modernize your data center, get you to the cloud quicker and provide the expertise, security and accessibility you need to grow your data-driven business. Combined with Quest's invitation to the global community to be a part of its innovation, and our firm commitment to ensuring customer satisfaction, we continue to deliver solutions that have a real impact on our customers today and leave a legacy we are proud of. We are challenging the status quo by transforming into a new software company. And as your partner, we work tirelessly to make sure your information technology is designed for you and by you. This is our mission, and we are in this together. Welcome to a new Quest. You are invited to Join the Innovation.

Our brand, our vision. Together.

Our logo reflects our story: innovation, community and support. An important part of this story begins with the letter Q. It is a perfect circle, representing our commitment to technological precision and strength. The space in the Q itself symbolizes our need to add the missing piece—you—to the community, to the new Quest.

Contacting Quest

For sales or other inquiries, visit <http://quest.com/company/contact-us.aspx> or call +1-949-754-8000.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request.
- View Knowledge Base articles.
- Sign up for product notifications.
- Download software and technical documentation.
- View how-to-videos.
- Engage in community discussions.
- Chat with support engineers online.
- View services to assist you with your product.

