



Release notes

BlueXP classification

NetApp
June 20, 2024

Table of Contents

- Release notes 1
- What's new with BlueXP classification 1
- Known limitations 7

Release notes

What's new with BlueXP classification

Learn what's new in BlueXP classification (Cloud Data Sense).

5 June 2024 (Version 1.32)

This release includes the following update.

New Mapping status column in the Configuration page

This release now shows a new Mapping status column in the Configuration page. The new column helps you identify if the mapping is running, queued, paused or more.

For explanations of the statuses, see [Change scan settings](#).

15 May 2024 (Version 1.31)

Classification is available as a core service within BlueXP

BlueXP classification is now available as a core capability within BlueXP at no additional charge for up to 500 TiB of scanned data. No Classification license or paid subscription is required. As we focus BlueXP classification functionality on scanning NetApp storage systems with this new version, some legacy functionality will only be available to customers who had previously paid for a license. The use of those legacy features will expire when the paid contract reaches its end date.

[Learn more about the deprecated features](#).

1 April 2024 (Version 1.30)

Support added for RHEL v8.8 and v9.3 BlueXP classification

This release provides support for Red Hat Enterprise Linux v8.8 and v9.3 in addition to previously supported 9.x, which requires Podman, rather than the Docker engine. This is applicable to any manual on-premises installation of BlueXP classification.

The following operating systems require using the Podman container engine, and they require BlueXP classification version 1.30 or greater: Red Hat Enterprise Linux version 8.8, 9.0, 9.1, 9.2, and 9.3.

Learn more about [BlueXP classification deployments overview](#).

BlueXP classification is supported if you install the Connector on a RHEL 8 or 9 host that resides on-premises. It's not supported if the RHEL 8 or 9 host resides in AWS, Azure, or Google Cloud.

Option to activate audit log collection removed

The option to activate audit log collection has been disabled.

Scan speed improved

Scan performance on secondary scanner nodes has been improved. You can add more scanner nodes if you

need additional processing power for your scans. For details, refer to [Install BlueXP classification on a host that has internet access](#).

Automatic upgrades

If you deployed BlueXP classification on a system with internet access, the system upgrades automatically. Previously, the upgrade occurred after a specific time elapsed since the last user activity. With this release, BlueXP classification upgrades automatically if the local time is between 1:00 AM and 5:00 AM. If the local time is outside of these hours, the upgrade occurs after a specific time elapses since the last user activity. For details, refer to [Install on a Linux host with internet access](#).

If you deployed BlueXP classification without internet access, you'll need to upgrade manually. For details, refer to [Install BlueXP classification on a Linux host with no internet access](#).

4 March 2024 (version 1.29)

Now you can exclude scanning data that resides in certain data source directories

If you want BlueXP classification to exclude scanning data that resides in certain data source directories, you can add these directory names to a configuration file that BlueXP classification processes. This feature enables you to avoid scanning directories that are unnecessary, or that would result in returning false positive personal data results.

[Learn more](#).

Extra Large instance support is now qualified

If you need BlueXP classification to scan more than 250 million files, you can use an Extra Large instance in your cloud deployment or on-premises installation. This type of system can scan up to 500 million files.

[Learn more](#).

10 January 2024 (version 1.27)

Investigation page results now display the total size in addition to total number of items

The filtered results in the Investigation page now shows the total size of the items in addition to the total number of files. This can help when moving files, deleting files, and more.

Configure additional Group IDs as "Open to Organization"

Now you can configure Group IDs in NFS to be considered as "Open to Organization" directly from BlueXP classification if the group had not initially been set with that permission. Any files and folders that have these group IDs attached will show as "Open to Organization" in the Investigation Details page. See how to [add additional Group IDs as "open to organization"](#).

14 December 2023 (version 1.26.6)

This release included some minor enhancements.

The release also removed the following options:

- The option to activate audit log collection has been disabled.

- During the Directories investigation, the option to calculate the number of personal identifiable information (PII) data by Directories is not available. Refer to [Investigate the data stored in your organization](#).
- The option to integrate data using Azure Information Protection (AIP) labels has been disabled. Refer to [Organize your private data](#).

6 November 2023 (version 1.26.3)

The following issues have been fixed in this release

- Fixed an inconsistency when presenting the number of files scanned by the system in dashboards.
- Improved the scanning behavior by handling and reporting on files and directories with special characters in the name and metadata.

4 October 2023 (version 1.26)

Support for on-premises installations of BlueXP classification on RHEL version 9

Red Hat Enterprise Linux versions 8 and 9 do not support the Docker engine; which was required for the BlueXP classification installation. We now support BlueXP classification installation on RHEL 9.0, 9.1, and 9.2 using Podman version 4 or greater as the container infrastructure. If your environment requires using the newest versions of RHEL, now you can install BlueXP classification (version 1.26 or greater) when using Podman.

At this time we don't supported dark site installations or distributed scanning environments (using a master and remote scanner nodes) when using RHEL 9.x.

5 September 2023 (version 1.25)

Small and medium deployments temporarily unavailable

When you deploy an instance of BlueXP classification in AWS, the option to select **Deploy > Configuration** and choose a small or medium-sized instance is unavailable at this time. You can still deploy the instance using the large instance size by selecting **Deploy > Deploy**.

Apply tags on up to 100,000 items from the Investigation Results page

In the past you could only apply tags to a single page at a time in the Investigation Results page (20 items). Now you can select **all** items in the Investigation Results pages and apply tags to all the items - up to 100,000 items at a time. [See how](#).

Identify duplicated files with a minimum file size of 1 MB

BlueXP classification used to identify duplicated files only when files were 50 MB or larger. Now duplicated files starting with 1 MB can be identified. You can use the Investigation page filters "File Size" along with "Duplicates" to see which files of a certain size are duplicated in your environment.

17 July 2023 (version 1.24)

Two new types of German personal data are identified by BlueXP classification

BlueXP classification can identify and categorize files that contain the following types of data:

- German ID (Personalausweisnummer)
- German Social Security Number (Sozialversicherungsnummer)

[See all the types of personal data that BlueXP classification can identify in your data.](#)

BlueXP classification is fully supported in Restricted mode and Private mode

BlueXP classification is now fully supported in sites with no internet access (Private mode) and with limited outbound internet access (Restricted mode). [Learn more about BlueXP deployment modes for the Connector.](#)

Ability to skip versions when upgrading a Private mode installation of BlueXP classification

Now you can upgrade to a newer version of BlueXP classification even if it is not sequential. This means that the current limitation of upgrading BlueXP classification by one version at a time is no longer required. This feature is relevant starting from version 1.24 onwards.

The BlueXP classification API is now available

The BlueXP classification API enables you to perform actions, create queries, and export information about the data you are scanning. The interactive documentation is available using Swagger. The documentation is separated into multiple categories, including Investigation, Compliance, Governance, and Configuration. Each category is a reference to the tabs in the BlueXP classification UI.

[Learn more about the BlueXP classification APIs.](#)

6 June 2023 (version 1.23)

Japanese is now supported when searching for data subject names

Japanese names can now be entered when searching for a subject's name in response to a Data Subject Access Request (DSAR). You can generate a [Data Subject Access Request report](#) with the resulting information. You can also enter Japanese names in the "Data Subject" filter in the [Data Investigation page](#) to identify files that contain the subject's name.

Ubuntu is now a supported Linux distribution on which you can install BlueXP classification

Ubuntu 22.04 has been qualified as a supported operating system for BlueXP classification. You can install BlueXP classification on a Ubuntu Linux host in your network, or on a Linux host in the cloud when using version 1.23 of the installer. [See how to install BlueXP classification on a host with Ubuntu installed.](#)

Red Hat Enterprise Linux 8.6 and 8.7 are no longer supported with new BlueXP classification installations

These versions are not supported with new deployments because Red Hat no longer supports Docker, which is a prerequisite. If you have an existing BlueXP classification machine running on RHEL 8.6 or 8.7, NetApp will continue to support your configuration.

BlueXP classification can be configured as an FPolicy Collector to receive FPolicy events from ONTAP systems

You can enable file access audit logs to be collected on your BlueXP classification system for file access events detected on volumes in your working environments. BlueXP classification can capture the following types of FPolicy events and the users who performed the actions on your files: Create, Read, Write, Delete, Rename, Change owner/permissions, and Change SACL/DACL.

Data Sense BYOL licenses are now supported in dark sites

Now you can upload your Data Sense BYOL license into the BlueXP digital wallet in a dark site so that you are notified when your license is getting low. [See how to obtain and upload your Data Sense BYOL license.](#)

3 April 2023 (version 1.22)

New Data Discovery Assessment Report

The Data Discovery Assessment Report provides a high-level analysis of your scanned environment to highlight the system's findings and to show areas of concern and potential remediation steps. The goal of this report is to raise awareness of data governance concerns, data security exposures, and data compliance gaps of your data set. [See how to generate and use the Data Discovery Assessment Report.](#)

Ability to deploy BlueXP classification on smaller instances in the cloud

When deploying BlueXP classification from a BlueXP Connector in an AWS environment, now you can select from two smaller instance types than what is available with the default instance. If you are scanning a small environment this can help you save on cloud costs. However, there are some restrictions when using the smaller instance. [See the available instance types and limitations.](#)

Standalone script is now available to qualify your Linux system prior to BlueXP classification installation

If you would like to verify that your Linux system meets all prerequisites independently of running the BlueXP classification installation, there is a separate script you can download that only tests for the prerequisites. [See how to check if your Linux host is ready to install BlueXP classification.](#)

7 March 2023 (version 1.21)

New functionality to add your own custom categories from the BlueXP classification UI

BlueXP classification now enables you to add your own custom categories so that BlueXP classification will identify the files that fit into those categories. BlueXP classification has many [predefined categories](#), so this feature enables you to add custom categories to identify where information that is unique to your organization are found in your data.

[Learn more.](#)

Now you can add custom keywords from the BlueXP classification UI

BlueXP classification has had the ability to add custom keywords that BlueXP classification will identify in future scans for a while. However, you needed to log into the BlueXP classification Linux host and use a command line interface to add the keywords. In this release, the ability to add custom keywords is in the BlueXP classification UI, making it very easy to add and edit these keywords.

[Learn more about adding custom keywords from the BlueXP classification UI.](#)

Ability to have BlueXP classification not scan files when the "last access time" will be changed

By default, if BlueXP classification doesn't have adequate "write" permissions, the system won't scan files in your volumes because BlueXP classification can't revert the "last access time" to the original timestamp. However, if you don't care if the last access time is reset to the original time in your files, you can override this behavior in the Configuration page so that BlueXP classification will scan the volumes regardless of

permissions.

In conjunction with this capability, a new filter named "Scan Analysis Event" has been added so you can view the files that were not classified because BlueXP classification couldn't revert last accessed time, or the files that were classified even though BlueXP classification couldn't revert last accessed time.

[Learn more about the "Last access time timestamp" and the permissions BlueXP classification requires.](#)

Three new types of personal data are identified by BlueXP classification

BlueXP classification can identify and categorize files that contain the following types of data:

- Botswana Identity Card (Omang) Number
- Botswana Passport Number
- Singapore National Registration Identity Card (NRIC)

[See all the types of personal data that BlueXP classification can identify in your data.](#)

Updated functionality for directories

- The "Light CSV Report" option for Data Investigation Reports now includes information from directories.
- The "Last Accessed" time filter now shows the last accessed time for both files and directories.

Installation enhancements

- The BlueXP classification installer for sites without internet access (dark sites) now performs a pre-check to make sure your system and networking requirements are in place for a successful installation.
- Installation audit log files are saved now; they are written to `/ops/netapp/install_logs`.

5 February 2023 (version 1.20)

Ability to send Policy-based notification emails to any email address

In earlier versions of BlueXP classification you could send email alerts to the BlueXP users in your account when certain critical Policies return results. This feature enables you to get notifications to protect your data when you're not online. Now you can also send email alerts from Policies to any other users - up to 20 email addresses - who are not in your BlueXP account.

[Learn more about sending email alerts based on Policy results.](#)

Now you can add personal patterns from the BlueXP classification UI

BlueXP classification has had the ability to add custom "personal data" that BlueXP classification will identify in future scans for a while. However, you needed to log into the BlueXP classification Linux host and use a command line to add the custom patterns. In this release, the ability to add personal patterns using a regex is in the BlueXP classification UI, making it very easy to add and edit these custom patterns.

[Learn more about adding custom patterns from the BlueXP classification UI.](#)

Ability to move 15 million files using BlueXP classification

In the past you could have BlueXP classification move a maximum of 100,000 source files to any NFS share. Now you can move up to 15 million files at a time. [Learn more about moving source files using BlueXP](#)

[classification.](#)

Ability to see the number of users who have access to SharePoint Online files

The filter "Number of users with access" now supports files stored in SharePoint Online repositories. In the past only files on CIFS shares were supported. Note that SharePoint groups that are not active directory based will not be counted in this filter at this time.

New "Partial Success" status has been added to the Action Status panel

The new "Partial Success" status indicates that a BlueXP classification action is finished and some items failed and some items succeeded, for example, when you are moving or deleting 100 files. Additionally, the "Finished" status has been renamed to "Success". In the past, the "Finished" status might list actions that succeeded and that failed. Now the "Success" status means that all actions succeeded on all items. [See how to view the Actions Status panel.](#)

9 January 2023 (version 1.19)

Ability to view a chart of files that contain sensitive data and that are overly permissive

The Governance dashboard has added a new *Sensitive Data and Wide Permissions* area that provides a heatmap of files that contain sensitive data (including both sensitive and sensitive personal data) and that are overly permissive. This can help you to see where you may have some risks with sensitive data. [Learn more.](#)

Three new filters are available in the Data Investigation page

New filters are available to refine the results that display in the Data Investigation page:

- The "Number of users with access" filter shows which files and folders are open to a certain number of users. You can choose a number range to refine the results - for example, to see which files are accessible by 51-100 users.
- The "Created Time", "Discovered Time", "Last Modified", and "Last Accessed" filters now allow you to create a custom date range instead of just selecting a pre-defined range of days. For example, you can look for files with a "Created Time" "older than 6 months", or with a "Last Modified" date within the "last 10 days".
- The "File Path" filter now enables you to specify paths that you want to exclude from the filtered query results. If you enter paths to both include and exclude certain data, BlueXP classification finds all files in the included paths first, then it removes files from excluded paths, and then it displays the results.

[See the list of all the filters you can use to investigate your data.](#)

BlueXP classification can identify the Japanese Individual Number

BlueXP classification can identify and categorize files that contain the Japanese Individual Number (also known as My Number). This includes both the Personal and Corporate My Number. [See all the types of personal data that BlueXP classification can identify in your data.](#)

Known limitations

Known limitations identify functions that are not supported by this release of the product, or that do not interoperate correctly with it. Review these limitations carefully.

BlueXP classification release removed options

The December 2023 (version 1.26.6) release removed the following options:

- The option to activate audit log collection has been disabled.
- During the Directories investigation, the option to calculate the number of personal identifiable information (PII) data by Directories is not available.
- The option to integrate data using Azure Information Protection (AIP) labels has been disabled.

BlueXP classification scanning limitations

BlueXP classification scans only one share under a volume

If you have multiple file shares under a single volume, BlueXP classification scans the share with the highest hierarchy. For example, if you have shares like the following:

- /A
- /A/B
- /C
- /D/E

Then the data in /A will be scanned. The data in /C and /D will not be scanned.

Workaround

There is a workaround to make sure you are scanning data from all the shares in your volume. Follow these steps:

1. In the working environment, add the volume to be scanned.
2. After BlueXP classification has completed scanning the volume, go to the *Data Investigation* page and create a filter to see which share is being scanned:

You'll filter the data by "Working Environment Name" and "Directory Type = Share" to see which share is being scanned.

3. Get the complete list of shares that exist in the volume so you can see which shares are not being scanned.
4. [Add the remaining shares to a share group.](#)

You'll need to add all the shares individually, for example:

```
/C  
/D
```

5. Perform these steps for each volume in the working environment that has multiple shares.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.