



Hewlett Packard
Enterprise

HPE Content Component

Software Version: 11.4.0

Release Notes

Document Release Date: June 2017

Software Release Date: June 2017

Legal notices

Warranty

The only warranties for Hewlett Packard Enterprise Development LP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted rights legend

Confidential computer software. Valid license from HPE required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright notice

© Copyright 2017 Hewlett Packard Enterprise Development LP

Trademark notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

This product includes an interface of the 'zlib' general purpose compression library, which is Copyright © 1995-2002 Jean-loup Gailly and Mark Adler.

Documentation updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent software updates, go to <https://downloads.autonomy.com/productDownloads.jsp>.

To verify that you are using the most recent edition of a document, go to <https://softwaresupport.hpe.com/group/softwaresupport/search-result?doctype=online help>.

This site requires that you register for an HPE Passport and sign in. To register for an HPE Passport ID, go to <https://hpp12.passport.hpe.com/hppcf/login.do>.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HPE sales representative for details.

Support

Visit the HPE Software Support Online web site at <https://softwaresupport.hpe.com>.

This web site provides contact information and details about the products, services, and support that HPE Software offers.

HPE Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Access product documentation
- Manage support contracts
- Look up HPE support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HPE Passport user and sign in. Many also require a support contract.

To register for an HPE Passport ID, go to <https://hpp12.passport.hpe.com/hppcf/login.do>.

To find more information about access levels, go to <https://softwaresupport.hpe.com/web/softwaresupport/access-levels>.

To check for recent software updates, go to <https://downloads.autonomy.com/productDownloads.jsp>.

Contents

New in this Release	5
Resolved Issues	8
Documentation	9

New in this Release

This section lists the enhancements to HPE Content Component version 11.4.0.

- You can now enable Content to search for documents that have not yet been flushed to disk, by setting the `SearchUncommittedDocuments` parameter to `True` in the `[Server]` section. In this case, documents are available to query from the time that an index job reaches the `Pending Commit` state.
- The `Query` action `Text` parameter now accepts the `SYNONYM` operator, which allows you to specify a list of terms to consider as synonyms. The query returns results as though these terms have `OR` operators between them, but relevance scores are calculated for all the synonym terms as if they were a single term. For example:

```
Text=SYNONYM(cat moggy kitten) NEAR sale
```

- You can now control the weight contribution that a single term in a particular field has toward the result score in a query, by using the new `MaxTermContribution` field property. This option allows you more fine control on the weighting of terms in your queries, if required.
- You can now configure Content to use GSSAPI authorization on the index port. To enable it, you must set the `GSSServiceName` parameter in the `[Server]` section to the fully qualified GSSAPI service name that the Content component must use to identify itself. You must also set `RequireGSSAuth` to **True** in the `[IndexServer]` section to enable authorization.

NOTE:

This method provides an authentication requirement only. HPE recommends that you use it in conjunction with TLS/SSL to encrypt the authentication data.

- The `GSSServiceName` parameter has been added to the `DREEXPORTREMOTE` index action to allow you to export documents to a server that uses GSSAPI authentication.
- The Content component now supports refreshing of `SecurityInfo` tokens. When you configure the new `SecurityTokenIdleTime` parameter in Community 11.4, Content returns refresh tokens in a `responsedata/autn:securityinfo` element in the response for actions that accept the `SecurityInfo` parameter.
- The `MinDate` and `MaxDate` query parameters now accept the following ISO-8601 date specifiers.
 - `YYYY-MM-DD`
 - `YYYY-MM-DDTHH:NN:SSZ`
 - `YYYY-MM-DDTHH:NN:SS±HHMM`

In particular, this change allows you to specify time zone information. These new formats are also available for the `BIASDATE` and `RANGE` `FieldText` operators, the `FIXEDDATE` option in the `Ranges` parameter for `GetQueryTagValues`, the `BackupForRestoreTime` action parameter, and the `ArchiveCleanupCutOffTime` and `RestoreTime` index action parameters.

- The `DateFormatCSVs` parameter `ZZZZZ` option now matches a single `Z` character as UTC, and matches time zone values between `-1300` and `+1500`.
- The `GetQueryTagValues` action now provides additional information about the range of numeric data when you set `ValueDetails` to `True`. The response now includes the

<autn:valuepercentile> tag, with values for various percentile values. By default, it returns the 10th, 25th, 50th (median), 75th, and 90th percentiles. You can optionally use the new `ValuePercentiles` parameter to specify additional options.

NOTE:

If you send the `GetQueryTagValues` action to a DAH, the `ValueDetails` response always includes only the default value percentiles. However, the DAH requests additional percentiles from the child servers so that it can calculate the appropriate ranges for the whole dataset.

- The `Ranges` parameter for the `GetQueryTagValues` action now accepts a `FIXEDDATE` operator, which allows you to set ranges of dates for date fields. You can specify dates in the same formats as the `MinDate` and `MaxDate` parameters. For example:

```
FIXEDDATE{.,13/03/2017,12:00:00 13/03/2017,10/03/2017,.}:autn_date
```

- You can now use multiple field names in the `FIXED` and `FIXEDDATE` operators in the `Ranges` parameter for the `GetQueryTagValues` action. Separate multiple field names after the operator with colons to apply the set of ranges to multiple fields. For example:

```
FIXEDDATE{.,13/03/2017,12:00:00 13/03/2017,10/03/2017,.}:autn_date:NUMDATE2
```

- The `GetQueryTagValues` action `ValueDetails` response now includes date attributes for values from `DateType` or `NumericDateType` fields, which contain formatted date values. When you also set `DateOffset`, the response also includes `date_with_offset` attributes.
- You can now filter the values that the `GetQueryTagValues` action returns by a string value (which can include wildcard characters), by using the new `ValueRestriction` parameter. When you set this parameter, the action returns only values for the field that match one of the specified restrictions.
- The `GetQueryTagValues` response now includes an `end_date` attribute for any date range values returned (for example, when you use the `FIXEDDATE` option in `Ranges`).

NOTE:

If you use Content through a DAH or QMS component, you must also upgrade those components to 11.4.0 to return the new response tag.

- The `CostEstimate` parameter has been added to the `DocumentStats` action. This parameter returns an estimate of the cost of running a query with the specified `Text`, including the minimum and maximum documents that Content might be read, and an estimate of the amount of data that might be read.
- The generic mapped security module supports Boolean operators in ACL fields such as `AUTONOMYMETADATA`. This allows connectors and custom scripts to use more complex logic to specify who is permitted to view a document. This might be necessary if the repository has a complex security model that cannot be represented by a single ACL.
- You can now configure action authorization more flexibly. The `[AuthorizationRoles]` configuration section has been added. You can add subsections to create roles, which can use a combination of existing roles (equivalent to the existing `AdminClients`, `QueryClients`, and so on), or a specific set of actions. For each role, you can specify the client IPs and hosts, SSL identities, and GSS principals to use to identify users that have particular permissions to run actions.

If you want to use only SSL and GSS authorization, you can disable the client settings by setting the appropriate client configuration parameters to `""`. For example, `AdminClients=""` disables client

authorization for administrative actions, and ensures that users must meet the SSL or GSS requirements.

- You can now set `SSLCertificate` to be a chain certificate in PEM format (consisting of the end-entity certificate, any intermediate certificates, and ending with the root CA certificate). This option allows a complete certificate to be returned to the connected peer.
- You can now set `SSLCheckCertificate` to `False` even when `SSLCACertificate` or `SSLCACertificatePath` are set. This allows the component to fill in any chain required for the `SSLCertificate` by using the certificates that you specify in `SSLCACertificate` and `SSLCACertificatePath`, without requiring a certificate from the connected peer.
- The `GSSAPILibrary` configuration parameter has been added to the `[Paths]` section. You can set this parameter to the path to the GSSAPI shared library or DLL that the application uses. Depending on your system configuration, HPE Content Component attempts to detect the appropriate library to use. However, if you use Kerberos or GSSAPI security in your setup, HPE recommends that you set an explicit value for this parameter.

Resolved Issues

This section lists the resolved issues in HPE Content Component version 11.4.0.

- When restarting, Content could erroneously log an error Expected at least N stemming entries, but read only M from stemming hash file.
- The `GetQueryTagValues` action sometimes omitted the `<autn:total_values>` response tag for fields where no values were returned, even when `TotalValues` was set to `True`.
- The persistent term cache was not correctly updated with newly indexed data.
- The `QuerySummary` response `<autn:element>` tag `pdocs` and `poccs` attributes could count the occurrences of the summary element that were not exact phrase matches. For example, the counts for `cat dog` would also include occurrences of `dog cat` and `cat dogs`.
- The `QuerySummary` response `<autn:element>` tag `ids` attribute could contain the IDs of any documents that contained all the terms in the element, regardless of whether they appeared as a phrase.
- When sending a query that set `Highlight` to `Proximity`, highlight tags could be incorrectly placed in fields with restricted indexing of alphanumeric terms (that is when your configuration has `IndexNumbers` is set to a value other than `1`, or includes the `IndexNumbersMMaxLength` settings).
- Sending the `GetTagValues` action with `FieldName` set to a metadata field could result in an ACI thread getting stuck in a busy loop.
- The response for the `GetQueryTagValues` action with `AllowNonParametricFields` set to `True` did not use consistent case normalization for values from parametric and non-parametric fields. Now, values for both parametric and non-parametric fields are returned as upper case unless you have configured `CaseSensitiveParametricValues`.
- The `GetVersion` action could incorrectly report their operating system on Microsoft Windows 10 and Microsoft Windows Server 2016 platforms.

Documentation

The following documentation was updated for this release.

- *IDOL Expert*
- *IDOL Getting Started Guide*
- *IDOL Server Reference*
- *IDOL Content Component Reference*
- *IDOL Server Administration Guide*