
eXtreme-DataCloud Documentation

eXtreme-DataCloud

May 20, 2019

Contents

1	eXtreme-DataCloud releases	3
2	Release repositories	27
3	Release schedule	29
4	Supported platforms	31

You'll find here usefull information regarding the eXtreme-DataCloud services and components **releases**, their schedules, documentation and support.

Contents

- *Welcome to eXtreme-DataCloud releases Documentation*
 - *eXtreme-DataCloud releases*
 - *Release repositories*
 - * *Source Code repositories*
 - * *Artefacts repositories*
 - *Release schedule*
 - * *Support Model*
 - *Supported platforms*
 - * *Supported artifacts & packaging formats*

1.1 XDC-1 (Pulsar)

The eXtreme-DataCloud project is pleased to announce the general availability of its **first public software release**, codenamed **Pulsar**

1.1.1 Included components

CachingOnDemand

Table of Contents

- *CachingOnDemand*
 - *Release Notes*
 - *Documentation*
 - *Support*

The **CachingOnDemand** system provides recipes and PaaS description templates for an end to end deployment of an XCache cluster.

Release Notes

v1.0.0

What's new

This is the first version of the CachingOnDemand - providing recipes and PaaS description templates for an end to end deployment of an XCache cluster.

List of RfCs

- N/A

Known Issues

- None

List of Artifacts

- Ansible roles: * <https://github.com/indigo-dc/CachingOnDemand/tree/master/ansible>
- **Docker Container:**
 - `indigodatacloud/cachingondemand:DEEP-1` (signed)
 - `indigodatacloud/cachingondemand:1.0.0`

Documentation

Detailed documentation is available at:

- [GitHub.io](https://github.com)

Support

- GitHub issues: <https://github.com/Cloud-PG/CachingOnDemand/issues>
- XDC project's internal support ticketing system: <http://jira.extreme-datacloud.eu/>

dCache

Table of Contents

- *dCache*
 - *Release Notes*
 - *Documentation*
 - *Support*

dCache is a distributed storage system proven to scale to hundreds of Petabytes. Originally conceived as a disk cache (hence the name) in front of a tertiary storage to provide efficient data access for data intensive scientific experiments in the field of High Energy Physics (HEP) it has evolved into highly scalable general-purpose open source storage solution.

Release Notes

vXDC-1

What's new

Upstream corresponding version: v. 5.0.0

Highlights

- dCache now supports Java 11 as its platform
- Documentation, especially the dCache Book, received a major revision and will remain in focus
- HTTP 3rd-party-copying has matured to a feature-rich, well tested state
- Pinboard includes timestamps
- updated external dependencies

Incompatibilities

- This release breaks compatibility with pools running dcache version 3.2 or earlier.

Acknowledgments

- We thank HTW Berlin students Marcel Munce, Ferdinand Wolff and MaKrHTW (???) as well as Onno Zweers from surfSARA for their contributions.

List of RfCs

Release 5.0.0

- **Admin**
 - A new property in the *frontend frontend.authz.unlimited-operation-visibility* now controls visibility of operations exposing file metadata. The default is *false*, meaning non-admin users can only see file operations for files which they own or which are anonymous. Setting it to *true* allows everyone access. (267d937c79).
 - Monitoring information exposed through the HTTP GET method is now available to all users and not only admin role users. (32597dc77a).
 - The admin data fields like the lists of pools, groups, units, etc., are now sorted by default for the admin REST API. (4928eff71d).
 - The dCache admin ssh server now supports kerberos as an authentication mechanism (along with password and publickey).
 - cbab40a841 added the following property to configure admin ssh server authentication:

```
(any-of?kerberos|password|publickey) admin.ssh.authn.enabled = password,publickey
```

The keytab's location can be set under

```
admin.ssh.authn.kerberos.keytab-file = /etc/krb5.keytab
```

- **Alarms**
 - A bug impeding reception of email alarms when the XML database is used has been fixed.
- **DCAP**

- Improved features: when using dcap URL to create a file or a directory, they are created with dcap get desired file permissions.

• Frontend

- dCache now supports more scientific file formats: HDF4 and 5 files as well as ROOT files are now identified and treated as such.
- The new configuration property (*one-of>true|false*)`dcache.enable.authn.anonymous-fallback-on-failed-login = true` allows modifying the behaviour of the frontend in case of failed logins: dCache has a hard-coded “feature” where a user providing bad authentication (e.g., wrong password, expired OIDC access-token or macaroon) is treated as the anonymous user. This has proved counter-intuitive, as wrong/expired credentials often appear to succeed for some operations (e.g., directory listing), while failing others (upload/download). Providing the new property allows to set a fail-fast behaviour in those cases, providing a quicker response to users.
- To support inotify events, a new plugin for SSE is introduced. Clients can discover changes in dCache namespace using an interface modelled after the inotify(7) API (See dCache book for detail).
- dCache View is updated to a new version (v1.5), see `dcache-view` repository for new feature details.

• FTP

- Bug which have been fixed:
 - * The leaking server sockets issue , when a client aborts a proxied transfers with kafka enabled is now fixed. No further server sockets leaked when a proxy is being used, Kafka notification is enabled, and the client aborts the transfer.
- Improved features: Improve date value formatting when sending billing events via Kafka.

• gplazma

- The credential information (e.g., distinguished name) is now logged for x509 certification chain validation and FQAN extraction failures. ([9c39e149e0](#)).
- Large numerical value gids may be used to define roles fro groupid (gid). ([11b34011ae](#)).
- Wildcard match of FQANS is possible for the VO group (vo-group.json) gplazma plugin. ([173dca3a96](#)).
- A new role, “observer”, is defined and available for according read-only access to system or file information. ([4aa440ab2a](#)).
- The Storage AuthzDB file format is updated to accept an optional ‘max-upload=<value>’ element after the ‘read-write’ or ‘read-only’ value. The label is optional. If present, the value describes the maximum file size the user can upload. ([e3dce67083](#))
- As some newer authentication mechanisms embed usage limitations; i.e., a user may authenticate in a way that limits what that user can do (E.g. SciTokens) New authentication plugins have the possibility to specify a Restriction as part of the authentication process. Existing authentication plugins are supported as before. ([204024b9e8](#)).
- A new configuration option has been introduced to capture all information about an OpenID-Connect provider, which is some external service that dCache users can authenticate against. This configuration property is a map. Each entry of the map associates a nickname with information about that provider. The nickname is used when logging problems with the provider. The information is the URI of the issuer endpoint. This must be a valid URL that starts ‘https://’. ([bab4e635ac](#)).
 - * The following example associates the nickname ‘google’ with Google’s issuer endpoint.

```
{{ gplazma.oidc.provider!google = https://accounts.google.com/}}
```

• History

- Error handling in the history service was improved.

- **Info**

- The info service now publishes the time that information was collected along with the actual data. The timestamp is available via the last-updated attribute.
- Info clients (such as info-provider and storage-report) are now informed of the number of files stored in a space reservation.

- **NFS**

- When pNFS client uses flex_file layout IO errors with pool (data server) are reported to NFS door. The errors can be interpreted as:

```
{ { NFS4ERR_NXIO: The client was unable to establish any communication with the ↵
↵storage device.

NFS4ERR_*: The client was able to establish communication with the storage
device and is returning one of the allowed error codes.}}
```

- **PNFS Manager**

- A user with a macaroon that authorises them to upload data into a particular directory will now also be able to create parent directories to achieve uploading the data.
- A bug that prevented get file checksum from working in some cases was fixed.

- **Pool**

- Fixed pool repository space accounting leak on failed restores from tape (815ce3eb6a).
- Added Cross-Origin Resource Sharing (CORS) support for HTTP requests (049c87a814) required by dCacheView.
- Fixed HTTPS redirected transfers by returning pool canonical hostname in the redirected URLs. (7f81b8e79d).
- Fixed stopwath error to ensure that IO-statistics collecting is more robust, avoiding stack-traces with the message ‘This stopwatch is already stopped’ (86ede8a240).
- Better handling of HTTP 3-rd party transfers - improved logging of exceptions (a98d667c16),
- increased socket timeout for GET requests (845cfe0bda).
- Improved error logging in billing by using exception calss name if exception has null message (24de520285).
- Removed stack-trace logging of checked exceptions on P2P failures (7a570355fa).
- Fixed pool runtime configured size regression (f5ba0103ea).
- Updated HTTP 3-rd party copy to support retrying GET and HEAD requests for better ineroperability with DPM (d0a621c775).
- Updated FTP mover to log additional information if it detects partial transfers (e725f7b9e7).
- Dropped subject from StorageInfoMessage (0e60cdcaaa).
- Fixed regression when restoring files from tape (7cdf4e0a7).
- Fixed NullPointerException on flush when using Kafka to collect billing records (4e396b9234).
- Fixed protocol movers to handle out of disk/out of capacity eerrors.
- Added support for Content-MD5 request header (4d954e6b5f).

- Updated HTTP mover to report errors as HTTP status message phrase so that clients that log the status line now provide their users with more detailed information about what caused a transfer to fail (6fcaeca34c).
- Fixed regression that broke “dcache pool convert2 command (80461b2f9a) and “dcache pool convert” command (80461b2f9a).
- Introduced a retry loop to retry file attributes update in timeout to pnfs manager ([8c60877527]((<https://github.com/dCache/dcache/commit/8c60877527869095acf23dd95f424d1df1e5b790>)).

• Pool Manager

- Select Read Pool requests for which the user does not have enough permissions now do not affect other requests any more.
- Several smaller bugfixes for Pool Manager also went into this release.

• Resilience

- Bugs which have been fixed: * an error is no longer reported when trying to handle a broken file which has already been unlinked; * the entire pool scan no longer fails when one file in the list is not resilient or has no locations; * filters referencing invalid pool names no longer cause scan cancel to fail.
- Improved features: * command retry errors immediately reprocesses the most recent failed file operations; * the command pool ls now displays the number of file operation errors encountered during a given scan; * the list of pools is now sorted by STATE (RUNNING, WAITING, IDLE) and then by pool name in ascending lexicographic order; * the inaccessible command now has options to check the status of the job, to display the current

contents of the ‘inaccessible list’ file for that pool, and to clean up/delete that file;

- * ‘referring pool’ has been added to the inaccessible alarm to enable grep’ing the resilience log for a given scanned pool.

• SRM / SRM Manager

- Fixes in gridsite delegation storage handling - fixed querying validity of delegated credential stored on the gridsite end-point allowing clients like davix-cp to work (839604e45f) with dCache;
- fixed handling of delegated credential with VOMS AC that expires before the X.509 (54658383d1);
- improved error reporting (41976be12d);
- added add gridsite delegation interface access-log (5392271fcf).
- SRM client has been updated to support X509_CERT_DIR environmental variable (ed8b86e604).
- Fixed handling of duplicate SURLs by SRM client (36b9e0c7d6).

• WebDAV

- A lot of work has gone into making 3rd party copying functionality more robust and scalable.

• XRootD

- Third-party copy was introduced in 4.2, and continues to be improved. For further information on configuration, please refer to the documentation in The Book (5.0).
- Bug fixes and improvements:
 - * the correct error (kXR_NoSpace) is now returned to the client when there is no more disk space;
 - * xrootd now fails fast if the MaxUploadSize is supplied, and the file is too large;
 - * the xrootd door spring configuration no longer fails to load when kafka is not activated;
 - * the ‘stat’ request now supports both open file handles as well as paths, enabling use of the –zip option;

- * dCache no longer logs a stack trace when a client requests a file be created, the parent directory does not exist, and the make parent option is omitted;
- * a source path containing a query part on a mv request no longer causes the request to fail;
- * a potential race condition preventing directory listing now is correctly handled;
- * support for the 'tpc' query on the pools has been added in order to comply with the newer (4.9) XrootD clients;
- * it is now no longer necessary nor correct to list 'access-log' among the xrootd plugins; this log handler is added automatically as it is for other doors; (10) file handles and query strings are now included in the access log information;
- * logging of failed authentication is improved to include more useful information, like the DN;
- * it is now possible to identify all entries in the access-log from the same TCP connection via a session identifier.

Known Issues

- None

Documentation

Please find below notes on how to enable and exploit the new features introduced in this versions:

- **Quality of Service**

- **Users** will interact with this feature using the graphical UI dCache View or through the REST API. While switching between QoS levels in dCache View is intuitive, the REST API is dynamically documented: all RESTful services have been provided with basic annotations in order automatically to generate API documentation. A convenient web interface which allows exploration and testing of the API, describing paths, parameters, error codes and JSON output, now runs at: [[https://\[{}host{}\]:3880/api/v1](https://[{}host{}]:3880/api/v1)].
- **Administrators** will need to set up their pools with tape connection as usual, and the GUI and REST interfaces are by default enabled for systems where the administrators choose to activate the frontend service.

- **Events (Kafka and SSE)**

- **Users** can listen to the various events sent from a dCache using industry-standard tools for the respective messaging systems.
- **Administrators** need to enable messaging and configure topics and triggers. This will be described in detail as soon as the Book is published, on the subpage /kafkaproducer/. In short: Kafka and Zookeeper need to be installed and available for the dCache instance in question, and the following properties need to be configured

```
(one-of?true|false)dcache.enable.kafka = true
{{ {{ dcache.kafka.bootstrap-servers = localhost:9092}}}}
```

List of Artifacts

- **CentOS-7 RPMS**
 - dcache-5.0.0-1.xdc.noarch.rpm

- **Ubuntu 16.04 DEBS**
 - `dcache-5.0.0-1.xdc.noarch.rpm`

Documentation

Detailed documentation is available at:

- [dCache Manuals](#)
 - In particular [dCache Book 5.0](#)

Support

- GitHub issues: <https://github.com/dCache/dcache/issues>
- XDC project's internal support ticketing system: <http://jira.extreme-datacloud.eu/>

Dynafed

Table of Contents

- *Dynafed*
 - *Release Notes*
 - *Documentation*
 - *Support*

The **Dynamic Federations system** allows to expose via HTTP and WebDAV a very fast dynamic name space, built on the fly by merging and caching (in memory) metadata items taken from a number of (remote) endpoints.

One of its fundamental features is to be redirect GET/PUT requests to the site or cluster hosting the requested file that is closer to the client that requested it. The focus is on performance, scalability and realtime fault resilience with respect to sites that can go offline.

From the perspective of a normal user, HTTP and WebDAV clients can browse the Dynamic Federation as if it were a unique partially cached name space, which is able to redirect them to the right host when they ask for a file replica. Dynafed also supports writing. [\[more\]](#)

Release Notes

v1.3.3

What's new

Dynafed **1.3.3** allows configuration of multiple concurrent authentication systems via separate namespace prefixes. It supports now OpenID Connect, both as a Relying Party (redirecting a browser to an IdP) and Protected Resource (consuming oAuth access tokens for non-interactive access), facilitating in this way the integration with the XDC Orchestrator and allowing browser based access without X509 certificates. This is configured through Apache's `mod_auth_openidc`.

More information on the “Authentication and Authorisation” section of documentation

List of RfCs

- [XSD-31](#) - Enable oauth2.0 for accessing Dynafed

Known Issues

- N/A

List of Artifacts

- CentOS-7 RPMS - in group [applications/internet](#)
 - dynafed - Ultra-scalable dynamic system for federating HTTP-based storage resources
 - dynafed-dmlite-frontend - dmlite plugin for dynafed
 - dynafed-dmlite-plugin - dmlite plugin for dynafed
 - dynafed-http-plugin - Http and WebDav plugin for dynafed
 - dynafed-lfc-plugin - Logical File catalog (LFC) plugin for dynafed
 - dynafed-private-devel - Development files for dynafed

Documentation

Detailed documentation is available at:

- [Dynafed - The Dynamic Federation project](#)

Support

- through [GGUS](#) by using the [Dynafed Development Support Unit](#)
- Dynafed User Forum mailing-list: [dynafed-users-forum \(cern.ch\)](mailto:dynafed-users-forum@cern.ch)
- XDC project's internal support ticketing system: <http://jira.extreme-datacloud.eu/>

EOS

Table of Contents

- *EOS*
 - *Release Notes*
 - *Documentation*
 - *Deployment automation:*
 - *Support*

EOS is an open-source storage software solution to manage multi PB storage for the CERN Large Hadron Collider LHC. Core of the implementation is the XRootD framework providing a feature-rich remote access protocol.

Release Notes

v4.4

What's new

- Update of remote IO objects for S3, WebDav and XRootd.
 - Used by EOS to communicate with external endpoints
 - Allows EOS to create filesystems whose storage is residing at a remote endpoint, instead of a local disk.
- Implementation of logical path translation.
 - So far, the path of every EOS file (so called 'physical path') is generated from the File Id of that file within the namespace. This feature allows EOS to store files having their physical location the same as their logical location in the namespace. E.g.: /eos/test/myfile on filesystem 1 will be stored at /<filesystem1_basepath>/eos/test/myfile
 - The logical path information is stored as file extended attributes within the EOS namespace
 - Path translation utility class, to help deciphering paths for both cases
- Filesystem config options: s3credentials and logicalpath
 - Two new filesystem config options were added.
 - s3credentials=<accesskey>:<secretkey> Sets the S3 access credentials for this filesystem
 - logicalpath=1|0 Flag to indicate whether this filesystem uses logical path or not
 - S3 settings can be configured FST-wide, via environment variables or on a filesystem-by-filesystem basis, via *fs config*
- Implementation of file adoption
 - Allows scanning an endpoint accessible by a filesystem and import all the discovered files into the EOS namespace
 - Management of these files, upon importation, can be done fully through EOS
 - A new command to trigger the file adoption is available, called *fs import*. It offers the *fs import start* command to trigger the import process. When triggered, an UUID is associated with the operation. The *fs import query <id>* command can be used to check the status of an on-going import operation
- External storage tests * External storage test script is provided to test the file adoption functionality * A new Gitlab pipeline job is set up to run this script on every commit
- Documentation * Documentation for logical path and import functionality are provided on the eos-docs website

List of RfCs

- [XSD-61](#) - Identity Forwarding plugin between XCache and EOS
- [XSD-62](#) - VOMS integration for XCache testbed
- [XSD-63](#) - VOMS security extractor for XrdHttp

- [XSD-73](#) - File traversal over HTTP
- [XSD-74](#) - Mechanism to import a file into the EOS namespace
- [XSD-75](#) - Allow EOS to use logical path
- [XSD-76](#) - System tests for adoption of external storage/files
- [XSD-147](#) - Implement status query of import procedure

Known Issues

- N/A

List of Artifacts

- CentOS-7 RPMS
 - in group [applications/file](#)
 - * eos-archive - The EOS archive daemon
 - * eos-cleanup - The EOS test package
 - * eos-client - The EOS shell client
 - * eos-fuse - The new EOS fuse client
 - * eos-fuse-core - The EOS fuse client
 - * eos-fuse-sysv - The EOS fuse client
 - * eos-fusex - The new EOS fuse client
 - * eos-fusex-core - The new EOS fuse client
 - * eos-fusex-selinux - The new EOS fuse client selinux configuration
 - * eos-server - The EOS server installation
 - * eos-srm - The EOS srm script package for checksumming and space
 - * eos-test - The EOS test package
 - * eos-testkeytab - The EOS testkeytab package
 - in [applications/internet](#)
 - * eos-admin-gui - Web Interface for EOS administration
 - in [Development/Debug](#)
 - * eos-debuginfo - Debug information for package eos

Documentation

Detailed documentation is available at:

- [EOS - OpenStorage Documentation](#)
 - information specific to XDC features can be found under the sections:
 - * “Configuration” (setting a Filesystem to use logical path)

- * “Using EOS” (describing the adoption of storage/files process)
- * “Client Commands” (the command to trigger the import [adoption] procedure).

Deployment automation:

Puppt modules available for:

- EOS server module: <https://gitlab.cern.ch/ai/it-puppet-module-eosserver/>
- Example files for setting up the namespace node and the storage node:
 - <https://gitlab.cern.ch/ai/it-puppet-hostgroup-eos/blob/qa/data/hostgroup/eos/home/i00/ns.yaml>
 - <https://gitlab.cern.ch/ai/it-puppet-hostgroup-eos/blob/qa/data/hostgroup/eos/home/i00/storage.yaml>

Support

- EOS tracking: <https://its.cern.ch/jira/projects/EOS/issues/>
- XDC project’s internal support ticketing system: <http://jira.extreme-datacloud.eu/>

FTS

Table of Contents

- *FTS*
 - *Release Notes*
 - *Documentation*
 - *Support*

FTS3 FTS3 is a bulk data mover, created to distribute globally the multiple petabytes of data from the LHC (Large Hadron Collider) at CERN.

Its purpose is to efficiently schedule data transfers, maximising the use of available network and storage resources while ensuring that any policy limits are respected.

Release Notes

v3.9.0

What’s new

This release adds support for OpenIDConnect and offers initial, limited support for QoS.

- **OpenIDConnect**
 - FTS can be configured to authorise users on the basis of OAuth2 tokens issued by an Authorisation Server such as IAM. The token can be used to subsequently drive file transfers.
- **QoS**

- FTS can now be given a target QoS as part of a transfer submission. With the present release, this will result in a query, via CDMI, of the QoS capabilities of the destination. This lays the foundation for the full support of managed QoS transitions in a subsequent XDC release.

List of RfCs

- [XSD-64](#) - Low level QoS client
- [XSD-65](#) - QoS job submission
- [XSD-67](#) - Basic OIC/OAuth support
- [XSD-68](#) - Support for refresh tokens

Known Issues

- N/A

List of Artifacts

- CentOS-7 RPMS
 - in group `applications/internet`
 - * `fts-client` - File Transfer Service version 3 client
 - * `fts-mysql` - File Transfer Service V3 mysql plug-in
 - * `fts-rest` - FTS3 Rest Interface
 - * `fts-rest-cli` - FTS3 Rest Interface CLI
 - * `fts-rest-cloud-storage` - FTS3 Rest Cloud Storage extensions
 - * `fts-rest-oauth2` - FTS3 Rest OAuth2 provider
 - * `fts-rest-selinux` - SELinux support for `fts-rest`
 - * `fts-server-selinux` - SELinux support for `fts-server`
 - * `gfal2-all` - Meta package for GFAL 2.0 install
 - * `gfal2-devel` - Development files of `gfal2`
 - * `gfal2-plugin-dcap` - Provides the support access for `gfal2`
 - * `gfal2-plugin-file` - Provides file support for `gfal2`
 - * `gfal2-plugin-gridftp` - Provides the gridftp support for `gfal2`
 - * `gfal2-plugin-http` - Provides the HTTP/DAV support for `gfal2`
 - * `gfal2-plugin-lfc` - Provides the lfc support for `gfal2`
 - * `gfal2-plugin-mock` - Provides a Mock dummy protocol for `gfal2`
 - * `gfal2-plugin-rfio` - Provides the rfio support for `gfal2`
 - * `gfal2-plugin-sftp` - Provide sftp support for GFAL2
 - * `gfal2-plugin-srm` - Provides the srm access for `gfal2`
 - * `gfal2-plugin-xrootd` - Provide xrootd support for GFAL2

- * gfal2-python - Python bindings for gfal 2
- * gfal2-python-doc - Documentation for gfal2-python
- * gfal2-python3 - gfal2 python bindings for Python 3
- * gfal2-tests - gfal2 tests
- * python-fts - FTS3 database model
- in [Development/Debug](#)
 - * fts-debuginfo - Debug information for package fts
 - * gfal2-debuginfo - Debug information for package gfal2
 - * gfal2-python-debuginfo - Debug information for package gfal2-python
- in [Documentation](#)
 - * gfal2-doc - Documentation for gfal2
- in [System environment/Daemons](#)
 - * fts-infosys - File Transfer Service version 3 infosys integration
 - * fts-msg - File Transfer Service version 3 messaging integration
 - * fts-server - File Transfer Service version 3 server
- in [System environment/Libraries](#)
 - fts-libs - File Transfer Service version 3 libraries

Documentation

Detailed documentation is available at:

- [FTS3 Documentation](#)
 - in particular the [Features](#) section with information on OIDC and QoS support
- [FTS3 at CERN](#)

Support

- through [GGUS](#) by using the [FTS Development Support Unit](#)
- FTS User Support mailing-list: fts-support@cern.ch
- **'FTS JIRA issue tracker <<https://its.cern.ch/jira/projects/FTS/issue>>'** _
- XDC project's internal support ticketing system: <http://jira.extreme-datacloud.eu/>

Onedata

Table of Contents

- *Onedata*
 - *Release Notes*

- *Documentation*
- *Support*

Onedata is a global data management system, providing easy access to distributed storage resources, supporting wide range of use cases from personal data management to data-intensive scientific computations.

Release Notes

v18.02.0-rc13

What's new

The new versions provides many improvements and new features, from which the most important are:

- improvements in the metadata extraction, ingestion and manipulation
- WebDAV support
- improved data monitoring API allowing to monitor the changes in data and depending on the type of the change invoking actions, ie. data replication
- improved speed of indexing of large datasets
- implementation of user defined data indexes, used for querring the data based on specific metadata

List of RfCs

- Oneclient:
 - VFS-4902 Added maximum upload size and connection pool size params
 - VFS-4902 Added WebDAV helper
 - VFS-4843 Adjusted default prefetch evaluation frequency
 - VFS-4843 Optimized random read prefetch calculation
 - VFS-4804 Fixed macaroon error handling
 - VFS-4804 Fixed handshake error handling
 - VFS-4804 Fixed reconnect
 - VFS-4804 Removed rest based full file prefetch
 - VFS-4741 Added sync block prefetch option
 - VFS-4809 Added prefetch skipping for prefetched offsets
 - VFS-4800 Fixed prefetch offset cache
 - VFS-4800 Added block aligned prefetch offset cache
 - VFS-4772 Align block prefetch offsets to cluster window size
 - VFS-4660 Added synchronize block priority handling
 - VFS-4656 Added cephados helper
- Oneprovider:

- VFS-4936 Add function to change space support size
- VFS-4830 add add_reduce function
- Updating GUI, including: VFS-4454-login-loader-close * VFS-4454 Fix hanging authorization loader spinner
- Upgrade rtransfer_link.
- Added support for webdav
- VFS-4936 Check if space exists when handling REST call
- VFS-4936 Add space support resizing
- VFS-4656 Added cephados helper
- VFS-4463 Showing storage in GUI on storages view
- VFS-4560 Detect existing Let's Encrypt certificates
- VFS-4029 Better certificate hostname verification
- VFS-4504 Set min and max port for distributed erlang
- Onezone:
 - VFS-4623 Add defaults to dns config envs
 - Updating GUI, including: VFS-4702-auth-icons-config * VFS-4702 Support for customizable authorization providers icons, colors and names
 - VFS-4614 New Universal auth.config for OIDC/SAML Identity Providers
 - VFS-4633 Tokens are not consumed upon failed operation
 - VFS-4029 Better certificate hostname verification
 - VFS-4029 Support http Let's Encrypt challenge in OZ and OP
 - VFS-4560 Detect existing Let's Encrypt certificates
 - VFS-4504 Set min and max port for distributed erlang

Known Issues

- N/A

List of Artifacts

- CentOS-7 RPMS
 - cluster-manager-18.02.0.rc13-1.el7.centos.x86_64.rpm
 - oneclient-18.02.0.rc13-1.el7.centos.x86_64.rpm
 - oneprovider-18.02.0.rc13-1.el7.centos.x86_64.rpm
 - onezone-18.02.0.rc13-1.el7.centos.x86_64.rpm
 - op-panel-18.02.0.rc13-1.el7.centos.x86_64.rpm
 - op-worker-18.02.0.rc13-1.el7.centos.x86_64.rpm
 - oz-panel-18.02.0.rc13-1.el7.centos.x86_64.rpm

- oz-worker-18.02.0.rc13-1.el7.centos.x86_64.rpm
- Ubuntu 16.04
 - cluster-manager_18.02.0.rc13-1_amd64.deb
 - oneclient_18.02.0.rc13-1_amd64.deb
 - oneprovider_18.02.0.rc13-1_amd64.deb
 - onezone_18.02.0.rc13-1_amd64.deb
 - op-panel_18.02.0.rc13-1_amd64.deb
 - op-worker_18.02.0.rc13-1_amd64.deb
 - oz-panel_18.02.0.rc13-1_amd64.deb
 - oz-worker_18.02.0.rc13-1_amd64.deb
- Docker Container
 - indigodatacloud/onezone:XDC-1 (signed)
 - indigodatacloud/onezone:18.02.0-rc13
 - indigodatacloud/oneprovider:XDC-1 (signed)
 - indigodatacloud/oneprovider:18.02.0-rc13
 - indigodatacloud/oneclient:XDC-1 (signed)
 - indigodatacloud/oneclient:18.02.0-rc13

Documentation

Detailed Documentation is available at:

- [OneData Guides](#)

Support

- GitHub issues: <https://github.com/onedata/onedata/issues>
- XDC project's internal support ticketing system: <http://jira.extreme-datacloud.eu/>

PaaS Orchestrator plugin

Table of Contents

- *PaaS Orchestrator plugin*
 - *Release Notes*
 - *Support*

The **INDIGO PaaS Orchestrator** is a component of the PaaS layer that allows to instantiate resources on Cloud Management Frameworks (like [OpenStack](#) and [OpenNebula](#)) and [Mesos](#) clusters.

It takes the deployment requests, expressed through templates written in [TOSCA YAML Simple Profile v1.0](#), and deploys them on the best cloud site available. In order to do that

- it gathers SLAs, monitoring info and other data from other platform services,
- it asks to the cloud provider ranker for a list of the best cloud sites.

Release Notes

v2.1.1-FINAL

What's new

This version provides a number of new features, improvements and bug fixes

List of RfCs

- New features:
 - Get Onedata access token using the user IAM token (fixes #294 XSD-27)
 - Describe Dynafed resources in TOSCA (fixes #295 XSD-30)
 - Extend the scheduling strategy using Dynafed (fixes #296 XSD-32)
 - Implement Cloud Providers Retry Logic (#288 DPD-121)
 - Retrieve of information from CMDB (fixes #290 DPD-124)
 - AuthN through OIDC (fixes #291 DPD-125)
 - Add support for mesos tasks with GPUs (fixes #292 DPD-126)
 - Support IM async infrastructure creation #289
 - Allow to import custom CAs in the java truststore #238
- Improvements
 - Improve OneData Spaces definition through templates (fixes #293 XSD-28)
- Bug fixed:
 - Return 409 status code for SQLTransientExceptions #285
 - (MySQL) Timestamp columns have implicit default value assigned #283

Known Issues

The PaaS Orchestrator v.2.1.1-FINAL has the following services dependencies

- CMDB release v0.4 (docker image indigodatacloud/cmdb:indigo_2)
- SLAM release v2.0.0 (docker image indigodatacloud/slam:v2.0.0)
- CPR release v0.6.0 (docker image indigodatacloud/cloudproviderranker:indigo_2)
- Monitoring - Zabbix Wrapper 1.0.2 (docker image indigodatacloud/zabbix-wrapper:indigo_2)
- IM release >= 1.7.4 * (Optional) Onedata v18.0.2-rc[11,12]

- `tosca-types v3.0.0`

List of Artifacts

- **Docker Container:**
 - `indigodatacloud/orchestrator:XDC-1 (signed)`
 - `indigodatacloud/orchestrator:2.1.1-FINAL`

Detailed documentation is available at:

- [GitHub README](#)
- [GitBook Guides](#)

Support

- GitHub issues: <https://github.com/indigo-dc/orchestrator/issues>
- XDC project's internal support ticketing system: <http://jira.deep-hybrid-datacloud.eu/>

TOSCA types & templates plugin

Table of Contents

- *TOSCA types & templates plugin*
 - *Release Notes*
 - *Documentation*
 - *Support*

The **TOSCA types** repository shows a YAML description of new types added first in the INDIGO-DataCloud project, and afterwards in the DEEP-HybridDataCloud (DEEP) and eXtreme DataCloud (XDC) projects, to extend TOSCA Simple Profile in YAML Version 1.0 to add high level entities. In the examples directory there are a set of TOSCA documents using these types that will be supported by the INDIGO, DEEP and XDC components.

The **TOSCA Templates** repository contain templates Supporting the Use Cases for INDIGO-DataCloud, DEEP-HybridDataCloud and eXtreme DataCloud projects.

Release Notes

v.3.0.0

What's new

This version provides many new features, improvements and bug fixes

List of RfCs

- New features:
 - Added new types for Onedata and Dynafed storage resources
 - Updated example templates for deploying Chronos dockerized jobs that use Onedata for managing input/output data
 - Updated example template for deploying a Mesos cluster
 - Added example template for deploying a Mesos cluster with GPU support and a tensorflow container on top of it that uses GPU(s)
 - Added GPU support for compute nodes and dockerized jobs (chronos) and apps (marathon)
 - Added preemptible_instance property to support “spot” instances
 - Added Onedata support in DODAS template
 - Added new types for describing a Kubernetes cluster; new Ansible roles implemented.
 - Added new type for describing a JupyterHub node; new Ansible role implemented.
- Bug fixes for Galaxy on cloud:
 - Update support to Galaxy release_18.05 in indigo-dc.galaxycloud
 - Fix proftpd in indigo-dc.galaxycloud
 - New ansible role for galaxy tools installation, named indigo-dc.galaxycloud-tools
 - Update toasca.nodes.indigo.GalaxyShedTool with new ansible role
 - Fix CERN-VM FS reference data mount on cluster worker nodes on galaxy artifacts
 - Reworked ansible role (indigo-dc.galaxycloud-fastconfig) to reconfigure an image with Galaxy already installed
 - Reworked storage encryption script on indigo-dc.galaxycloud-os
 - Fix tasks order in Galaxy elastic cluster toasca template: now galaxy user is created before slurm configuration.

Known Issues

- N/A

List of Artifacts

- **Tarballs:**
 - toasca-types-3.0.0.tar.gz
 - toasca-templates-3.0.0.tar.gz

Documentation

Detailed documentation can be found at:

- [TOSCA types README](#)

- TOSCA templates README
- Brief guide to TOSCA
- How to deploy a TOSCA Template

Support

- **GitHub issues:**
 - <https://github.com/indigo-dc/tosca-types/issues>
 - <https://github.com/indigo-dc/tosca-templates/issues>
- XDC project's internal support ticketing system: <http://jira.extreme-datacloud.eu/>

1.1.2 Highlights

Key technical highlights:

- CachingOnDemand
 - deployment receipts for geographically distributed caches (via xcache)
 - deployment receipts for scalable local caches (via xcache and http)
- dCache
 - new QoS types integration, aggregated QoS for storage federations
 - OpenIDConnect support in dcache_view
 - dcache storage events (SSE notifications): Allow non-dCache agent to get notified that something of interest happen inside dCache
- Dynafed
 - Integration of OIDC authentication
- EOS
 - Caching with xcache for geographic deployment: Xcache deployed at a remote centre to accelerate its local CPU
 - External storage adoption (Through an S3 or a WebDAV interface)
 - External data adoption (Data already present on a system described above can be incorporated into EOS)
- FTS & GFAL
 - QoS support: can now accept a QoS job
 - OpenIDConnect support
 - QoS in gfal (gfal with basic cdmi client) – python bindings available
- PaaS Orchestrator
- Implementation of Dynafed plugin
- Interaction via INDIGO IAM OAUTH2 token
- Enhancement of ONEDATA plugin
- Onedata

- Performance and stability improvements
- Support for groups and roles
- New RADOS driver

1.1.3 Release Notes

The XDC-1 (Pulsar) release consists in X Products and a number of technical guides:

- 113 OS packages
 - 100 RPMS & SRPMS, and tarballs
 - 13 binary & source DEBS
- 6 Docker containers

The release is fully supported - on the following Operating Systems platforms:

- CentOS 7
- Ubuntu 16.04
- Optionally PTs support also other OS platforms. You can find more information in the individual products documentation.

You can find in the later sections the full list of products, with detailed release notes and instructions for their installation/configuration.

1.1.4 Generic Installation Notes

All eXtreme-DataCloud products are distributed from standard OperatingSystems (OS) repositories and DockerHub registry.

The packages repositories have the following structure:

- **XDC production** (stable):
 - `xdc/production/{1,2}/centos7/x86_64/{baselupdates}`
 - `xdc/production/{1,2}/ubuntu/dists/xenial/main/{binary-amd64,source}`
 - * containing signed, well tested software components
 - **third-party**:
 - * `xdc/production/{1,2}/centos7/x86_64/third-party`
 - * `xdc/production/{1,2}/ubuntu/dists/xenial/third-party{binary-amd64,source}`
 - containing packages that are not part of eXtreme DataCloud, or not part of the base OS or EPEL, but used as dependencies by other eXtreme DataCloud components
- **XDC testing**: `xdc/testing/{1,2}/{centos7,ubuntu}/`
 - containing packages that will become part of the next stable distribution; in the certification and validation phase.
- **XDC preview**: `xdc/preview/{1,2}/{centos7,ubuntu}/`
 - containing signed packages that will become part of the next stable update, available for technical-previews

All signed packages use the INDIGO - DataCloud gpg key. The public key can be downloaded from [here](#), and the fingerprint from [here](#).

It is strongly recommended the use of the latest version of the xdc-release packages containing the public key and the YUM and APT repository files.

On the [DockerHub Registry](#), eXtreme-DataCloud uses the INDIGO - DataCloud and XDC Organizations:

- [indigodatacloud](#),
- [xdc](#)

Containers present in those repositories and released in XDC-1 are tagged with “XDC-1” tag and signed, leveraging the Docker’s trust features so that users can pull trusted images.

To understand how to install and configure XDC-1/Pulsar services and components either refer to the *Generic Installation Notes* chapter or to each individual product documentation.

1.1.5 Software

XDC-1 software can be downloaded from [eXtreme-DataCloud repositories](#).

1.1.6 Documentation

Please find XDC-1 documentation [here](#).

1.1.7 Support

Most complex software contains bugs, we are not an exception. One of the features of free and open source software is the ability to report bugs, helping to fix or improve the software you use.

eXtreme-DataCloud project uses the [INDIGO Catch-All GGUS - Support Unit](#) and the support@extreme-datacloud.eu for general support requests. More details regarding each product support channels are provided in the respective products release pages.

Developers, researchers and IT enthusiasts: feel free to write to info@extreme-datacloud.eu to ask for more information on how to use XDC solutions for your work. For automatic notifications you can register to the eXtreme-DataCloud release RSS feed or subscribe to the eXtreme-DataCloud Announce Mailing list. You can also socialize with us via [Twitter](#), [Facebook](#) and join our [LinkedIn group](#).

2.1 Source Code repositories

Source code repositories are available on [GitHub](#) under the “indigo-dc” and “extreme-datacloud” organizations

2.2 Artefacts repositories

eXtreme-DataCloud **production** (stable) repositories:

- `xdc/production/{1,2}/centos7/x86_64/{base | updates}`
- `xdc/production/{1,2}/ubuntu/dists/xenial/main/{binary-amd64, source}` * containing signed, well tested software components
- `third-party: * xdc/production/{1,2}/centos7/x86_64/third-party * xdc/production/{1,2}/ubuntu/dists/xenial/third-party{binary-amd64, 'source <http://repo.indigo-datacloud.eu/repository/xdc/production/1/ubuntu/dists/xenial/third-party/source>' _}`
 - containing packages that are not part of eXtreme-DataCloud project, or not part of the base OS or EPEL, but used as dependencies by other eXtreme-DataCloud components

YUM & APT **configuration files** are available [here](#) or use the **xdc-release** package to install eXtreme-DataCloud repositories

Release schedule

- **Time-based releases**
 - **projects' Major releases** - the eXtreme-DataCloud project foresees two major releases, distributions, during its lifetime, at around 10 months since the start of the project.
- **As-soon-as-available**
 - **components' Minor/Revision releases** - in a project Major release, Development teams (aka Product Teams) can release updated versions of their components as soon as the XDC software quality criteria are met. Thought the project Continuous Integration and Delivery System tests are continuously run giving feedback on the status of the components.

3.1 Support Model

- in a Major Release for each component or service only the latest revision released is supported.
- for each component or service a (major) release is supported at least for the lifetime of the projects' major release in which this version was released the first time.

Supported platforms

- eXtreme-DataCloud releases are supported on the following platforms:
 - CentOS7 & Ubuntu 16.04
 - * for the products distributed through rpms and deb packages
 - all platforms supporting Docker containers
 - * for the products distributed as docker images

4.1 Supported artifacts & packaging formats

- Packages:
 - Binaries: executable packages
 - Sources: when available, package files that contain all of the necessary files to compile/build the respective piece of software
 - Tarballs: clients are distributed as tarballs for all the platforms
- Containers: Docker images are made available for some of the project software