

What to expect in 2.16

- or what's new since last year

Gerd Behrmann

10th dCache User Workshop



norden

NordForsk



Nordic e-Infrastructure
Collaboration

What to remember from this talk

- dCache 2.16 Released June 2016
- dCache 2.10 EOL December 2016
- dCache 2.13 EOL June 2017
- We need volunteers to write a 2.13 to 2.16 upgrade guide.

JGlobus is gone - caNI is on

- It's actually maintained.
- Faster, more configurable, integrates better with dCache.
- Was already used by the VOMS library.
- Reduced overhead for checking whether key material has changed, thus defaults are significantly reduced.
- Support for EUGRIDPMA name space verification (.namespace files).
- Support for OCSP.
- Better error messages.

Custom GSI implementation

- caNI lacks GSI implementation, so we rolled our own (with better performance).
- Caching of client certificate chain validation.
- Disabled RC4.
- Strict RFC 2818 host name verification (request host certs with SANs!) in clients.

```
$ openssl x509 -in hostcert.pem -text|grep -A1 Alternative  
X509v3 Subject Alternative Name:  
DNS:piggy.ndgf.org, DNS:dcache.ndgf.org,  
DNS:dav.ndgf.org, DNS:srm.ndgf.org
```

FTP

- FTP client was part of JGlobus. Now forked and ported to caNL.
- Identifies itself as dCache rather than Globus.
- New in the FTP door:

```
UberFTP (2.8)> chgrp 3000 test-1459847050-9  
UberFTP (2.8)> symlink foobar test-1459847050-9
```

Alternation lists in glob patterns

```
[piggy] (local) admin > \l *_{fi,org}*_  
csc_fi_025  
csc_fi_026  
csc_fi_027  
csc_fi_028  
csc_fi_029  
csc_fi_030  
csc_fi_031  
csc_fi_032  
csc_fi_033  
csc_fi_034  
csc_fi_035  
csc_fi_036  
ore_ndgf_org_002  
ore_ndgf_org_003
```

Pool group patterns in globs

```
[piggy] (local) admin > \l /ops_disk  
ore_ndgf_org_002  
ore_ndgf_org_003
```

```
[piggy] (local) admin > \l *_dk_*/atlas_tape_write  
atlas_hpc_ku_dk_w03
```

Lots of documentation updates

[Gerds-MacBook-Pro] (System@dCacheDomain) admin > \h version
NAME

version -- query jar file metadata

SYNOPSIS

version [packagename]

DESCRIPTION

Information about the implementation-title, -vendor and -version, as described within some jar. The jar file is the one that provides some specific Java package. If the jar file implementation-vendor is 'dCache.org' then the implementation-version is the dCache version.

ARGUMENTS

packagename

The package used to select the jar file. Defaults to **dmg.cells.nucleus**.

Primary key change on t_tags_inodes

chimera=# \d t_tags_inodes

		Table "public.t_tags_inodes"
Column	Type	Modifiers
itagid	bigint	not null default nextval('t_tags_inodes2_itagid_seq'::regclass)
imode	integer	not null
inlink	integer	not null
iuid	integer	not null
igid	integer	not null
isize	integer	not null
ictime	timestamp without time zone	not null
iatime	timestamp without time zone	not null
imtime	timestamp without time zone	not null
ivalue	bytea	

Indexes:

"t_tags_inodes_pkey" PRIMARY KEY, btree (itagid)

Referenced by:

TABLE "t_tags" CONSTRAINT "t_tags_itagid_fkey" FOREIGN KEY (itagid)
 REFERENCES t_tags_inodes(itagid) ON UPDATE CASCADE

Primary key change on t_inodes

chimera=# \d t_inodes

Table "public.t_inodes"			
Column	Type	Modifiers	
ipnfsid	character varying(36)	not null	
itype	integer	not null	
imode	integer	not null	
inlink	integer	not null	
iuid	integer	not null	
igid	integer	not null	
isize	bigint	not null	
iio	integer	not null	
ictime	timestamp with time zone	not null	
iatime	timestamp with time zone	not null	
imtime	timestamp with time zone	not null	
icrtime	timestamp with time zone	not null	
igeneration	bigint	not null	default 0
iaccess_latency	smallint		
iretention_policy	smallint		
inumber	bigint	not null	default nextval('t_inodes_inumber_seq'::regclass)

Indexes:

"t_inodes_pkey" PRIMARY KEY, btree (ipnfsid) CLUSTER

"i_inodes_inumber" UNIQUE CONSTRAINT, btree (inumber)

Other changes in Chimera

- . and .. is no longer stored in t_dirs.
- Lots of changes to triggers for tags, ACLs and inode creation.
- Lots of changes to stored procedures e.g. **inumber2path**, **path2inumber**, **pnfsid2inumber**, **inumber2pnfsid**.
- New PostgreSQL 9.5 specific driver.
- Radically changed transaction demarcation results in better isolation and better fault recovery.
- Lots of corner cases fixed.
- Custom scripts, views, schema modifications require manual work upon upgrade!

Chimera cleanup upon upgrade

- We used to have a bug that caused `t_tags_inodes` entries to be leaked upon directory deletion.
- Fixed in 2.10+, but already leaked entries are not deleted until upgrade to 2.13.
- Upon upgrade to 2.15 a one-time `fsck` like procedure runs and creates a `lost+found` directory in your namespace root.

PnfsManager Chimera stats

ChimeraNameSpaceProvider	average±stderr(ms)		min(ms)	max(ms)	STD(ms)	Samples
addInodeLocation(FsInode int Strin	0.96±	0.00	0	667	1.39	5,367,431
clearInodeLocation(FsInode int Str	0.63±	0.00	0	3,537	3.76	4,518,162
createFile(FsInode String int int	1.86±	0.01	0	1,317	4.25	307,799
getACL(FsInode)	0.44±	0.00	0	519	0.78	65,559,466
getAllTags(FsInode)	0.75±	0.00	0	1,021	1.06	8,627,915
getInfo()	0.12±	0.01	0	439	2.50	32,041
getInodeChecksums(FsInode)	0.56±	0.00	0	3,976	2.11	10,523,929
getInodeLocations(FsInode)	0.48±	0.00	0	2,600	1.75	8,322,064
getInodeLocations(FsInode int)	0.59±	0.00	0	3,256	2.28	6,410,034
getParentOf(FsInode)	0.44±	0.00	0	463	0.77	5,184,057
getStorageInfo(FsInode)	0.52±	0.00	0	409	0.76	1,181,207
id2inode(String StatCacheOption)	0.42±	0.00	0	3,195	1.86	29,002,323
inode2id(FsInode)	0.01±	0.00	0	9	0.08	434,135
inode2path(FsInode)	1.20±	0.01	0	19	0.92	9,241
inodeOf(FsInode String StatCacheOp	0.62±	0.00	0	981	1.47	2,103,859
mkdir(FsInode String int int int)	3.48±	0.08	1	1,611	11.31	22,266
mkdir(FsInode String int int int L	7.89±	0.01	2	1,352	7.77	303,986
newDirectoryStream(FsInode)	2.11±	0.00	0	134	1.13	132,502
path2inode(String)	1.22±	0.00	0	825	1.96	1,344,023
path2inodes(String)	2.14±	0.00	0	3,538	4.59	6,138,865
read(FsInode int long byte[] int i	0.39±	0.00	0	266	0.73	2,998,864
remove(FsInode String FsInode)	4.36±	0.00	0	1,362	2.59	1,035,761
rename(FsInode FsInode String FsIn	2.77±	0.01	0	829	3.05	302,220
setInodeAttributes(FsInode int Sta	0.60±	0.00	0	3,254	5.03	2,925,121
setInodeChecksum(FsInode int Strin	0.67±	0.00	0	108	0.75	307,097
setStorageInfo(FsInode InodeStorag	1.21±	0.01	0	28	0.76	5,896
stat(FsInode int)	0.50±	0.00	0	3,000	1.36	13,121,645
write(FsInode int long byte[] int	1.30±	0.01	0	3,253	6.03	307,407

Single request queue in PnfsManager

```
[piggy] (PnfsManager@dCacheDomain2) admin > infox pnfs-manager  
atime precision: Disabled
```

```
List queues (1)  
[0] 0
```

```
Message queues (1)  
[0] 0
```

gPlazma plugin instantiation changed

```
$ cat /etc/dcache/gplazma.conf
auth      optional  voms
auth      optional  x509
map       optional  vorolemap
map       optional  gridmap
map       requisite authzdb gplazma.authzdb.uid=uid,login,group,user \
                        gplazma.authzdb.gid=gid,login,group,user
session   requisite authzdb
```

Other gPlazma changes

- Concurrent request processing provides much better scalability
- DCAP doors rely on gPlazma for VOMS validation
- x509 plugin enforces client source IP policy constraints embedded in the proxy certificate.

gPlazma VOMS

Login attempt failed; detailed explanation follows:

LOGIN FAIL

in: Origin[127.0.0.1]

X509 Certificate chain:

+---CN=849220543,CN=Gerd Behrmann,OU=ndgf.org,O=NorduGrid,O=Grid [6464]

+---Issuer: CN=Gerd Behrmann,OU=ndgf.org,O=NorduGrid,O=Grid

+---Validity: OK for 11 hours, 59 minutes and 48,5 seconds

+---Algorithm: SHA-512 with RSA

+---Public key: RSA 1024 bits

+---Attribute certificates:

+---CN=host/voms.ndgf.org,O=NorduGrid,O=Grid

+---Validity: OK for 11 hours, 59 minutes and 15,4 seconds

+---Extensions:

+---Issuer: CN=NorduGrid Certification Authority 2015,O=NorduGrid,O=Grid

+---No revocation info

+---Authority key identifier

+---Algorithm: SHA-512 with RSA

+---FQANs: /ops.ndgf.org, /ops.ndgf.org/monitoring, ...

+---Key usage: digital signature, key encipherment

+---CN=Gerd Behrmann,OU=ndgf.org,O=NorduGrid,O=Grid [6464]

+---Issuer: CN=NorduGrid Certification Authority,O=NorduGrid,O=Grid

+---Validity: OK for 31 days, 7 hours, 24 minutes and 32,4 seconds

+---Algorithm: SHA-512 with RSA

+---Public key: RSA 2048 bits

+---Subject alternative names: email: behrmann@ndgf.org

+---Key usage: digital signature, non-repudiation, key encipherment



```
$ curl -X PROPFIND -H Depth:0 http://localhost:2880/disk/test-1447231167-1 \
--data '<?xml version="1.0" encoding="utf-8"?>
  <D:propfind xmlns:D="DAV:">
    <D:prop xmlns:R="http://www.dcache.org/2013/webdav"
      xmlns:S="http://srm.lbl.gov/StorageResourceManager">
      <R:Checksums/>
      <S:AccessLatency/>
      <S:RetentionPolicy/>
      <S:FileLocality/>
    </D:prop>
  </D:propfind>'

<?xml version="1.0" encoding="utf-8" ?>
<d:multistatus xmlns:cal="urn:ietf:params:xml:ns:caldav"
  xmlns:cs="http://calendarserver.org/ns/"
  xmlns:card="urn:ietf:params:xml:ns:carddav"
  xmlns:ns2="http://www.dcache.org/2013/webdav"
  xmlns:ns1="http://srm.lbl.gov/StorageResourceManager"
  xmlns:d="DAV:">
  <d:response>
    <d:href>/disk/test-1447231167-1</d:href>
    <d:propstat>
      <d:prop>
        <ns1:FileLocality>ONLINE</ns1:FileLocality>
        <ns1:RetentionPolicy>REPLICA</ns1:RetentionPolicy>
        <ns2:Checksums>adler32=3b8711d6</ns2:Checksums>
        <ns1:AccessLatency>ONLINE</ns1:AccessLatency>
      </d:prop>
      <d:status>HTTP/1.1 200 OK</d:status>
    </d:propstat>
  </d:response>
</d:multistatus>
```


HTTP Third Party Copy

```
$ curl -L -v -X COPY -k -E "Gerd Behrmann" \  
-H "Destination:https://localhost:2881/disk/test1" https://localhost:2881/disk/test  
* Trying 127.0.0.1...  
* Connected to localhost (127.0.0.1) port 2881 (#0)  
* Client certificate: Gerd Behrmann  
* TLS 1.2 connection using TLS_RSA_WITH_AES_128_CBC_SHA256  
* Server certificate: host/localhost  
> COPY /disk/test HTTP/1.1  
> Host: localhost:2881  
> User-Agent: curl/7.43.0  
> Accept: */*  
> Destination:https://localhost:2881/disk/test1  
>  
< HTTP/1.1 200 OK  
< Date: Fri, 08 Apr 2016 06:58:16 GMT  
< Server: dCache/2.16.0-SNAPSHOT  
< Transfer-Encoding: chunked  
<  
Perf Marker  
    Timestamp: 1460098696  
    State: 1  
    State description: querying file metadata  
    Stripe Index: 0  
    Total Stripe Count: 1  
End  
success: Created
```


HTTP Third Party Copy

```
$ curl -L -v -X COPY -k -E "Gerd Behrmann" \  
-H "Source:http://localhost:2880/disk/test1" https://localhost:2881/disk/test2  
* Trying 127.0.0.1...  
* Connected to localhost (127.0.0.1) port 2881 (#0)  
* Client certificate: Gerd Behrmann  
* TLS 1.2 connection using TLS_RSA_WITH_AES_128_CBC_SHA256  
* Server certificate: host/localhost  
> COPY /disk/test2 HTTP/1.1  
> Host: localhost:2881  
> User-Agent: curl/7.43.0  
> Accept: */*  
> Source:http://localhost:2880/disk/test1  
>  
< HTTP/1.1 200 OK  
< Date: Fri, 08 Apr 2016 07:00:45 GMT  
< Server: dCache/2.16.0-SNAPSHOT  
< Transfer-Encoding: chunked  
<  
Perf Marker  
    Timestamp: 1460098845  
    State: 5  
    State description: creating namespace entry  
    Stripe Index: 0  
    Total Stripe Count: 1  
End  
success: Created
```

```
[dCacheDomain/webdav]
webdav.cell.name=WebDAV-${host.name}
webdav.net.port=2880
webdav.authz.anonymous-operations=READONLY
webdav.authn.basic=true
webdav.authn.protocol=http
webdav.custom-response-header!Site=NDGF
```



```
$ curl -I http://localhost:2880/disk/test4
HTTP/1.1 200 OK
Date: Fri, 08 Apr 2016 07:18:59 GMT
Server: dCache/2.16.0-SNAPSHOT
Site: NDGF
Accept-Ranges: bytes
ETag: "0000272B914853F04518886E8981184A5CD1_-189838237"
Last-Modified: Fri, 08 Apr 2016 07:04:02 GMT
Content-Length: 628496
```




HTTP Third Party Copy

```
$ srmcp srm://localhost/disk/test2 https://localhost:2881/disk/test3  
$ srmcp http://localhost:2880/disk/test3 srm://localhost/disk/test4
```

SRM Database

```
$ dcache database ls
```

DOMAIN	CELL	DATABASE	HOST	USER	MIN-	MAX-CONNS	MANAGEABLE	AUTO
dCacheDomain	PnfsManager	chimera	localhost	dcache		30	Yes	Yes
dCacheDomain	replicaManager	replica	localhost	dcache			Yes	No
dCacheDomain	billing	billing	localhost	dcache		30	Yes	Yes
dCacheDomain	SpaceManager	spacemanager	localhost	dcache		30	Yes	Yes
dCacheDomain	PinManager	pinmanager	localhost	dcache		50	Yes	Yes
dCacheDomain	cleaner	chimera	localhost	dcache			No	No
dCacheDomain	RemoteTransferManager	transfermanagers	localhost	dcache			No	Yes
dCacheDomain	SRM-Gerds-MacBook-Pro	srm	localhost	dcache			Yes	Yes
dCacheDomain	NFS-Gerds-MacBook-Pro	chimera	localhost	dcache		30	Yes	No
dCacheDomain	Resilience	chimera	localhost	dcache		120	Yes	Yes
TOTAL					0	340		



SRM Database

```
[srm=# \d srmuser
```

```
Table "public.srmuser"
```

Column	Type	Modifiers
id	bigint	not null
type	character varying(8)	
encoded	bytea	

```
Indexes:
```

```
"srmuser_pkey" PRIMARY KEY, btree (id)
```

```
[srm=# select id,type,substring(encoded for 30) from srmuser limit 10;
```

id	type	substring
-4807410496433933260	Login	\xaced000577040000000d7372002a6f72672e676c6f6275732e6773692e67
7395220753401226756	Login	\xaced00057704000000097372002a6f72672e676c6f6275732e6773692e67
-4067826650631570841	Login	\xaced000577040000000c7372002a6f72672e676c6f6275732e6773692e67
8647431496091489806	Login	\xaced00057704000000097372002a6f72672e676c6f6275732e6773692e67
7988394277470390332	Login	\xaced000577040000000b7372002a6f72672e676c6f6275732e6773692e67
5885387840006900736	Login	\xaced00057704000000097372002a6f72672e676c6f6275732e6773692e67
-124735714123320394	Login	\xaced000577040000000d7372002a6f72672e676c6f6275732e6773692e67
-4506751710214006139	Login	\xaced000577040000000d7372002a6f72672e676c6f6275732e6773692e67
-7810150757120161484	Login	\xaced000577040000000a7372002a6f72672e676c6f6275732e6773692e67
-2411390742379799377	Login	\xaced000577040000000a7372002a6f72672e676c6f6275732e6773692e67

(10 rows)


```
level=INFO
ts=2016-04-05T19:27:19.150+0200
event=org.dcache.srm.request
session=Tw:1906734:srm2:prepareToPut
socket.remote=130.246.182.186:36180
request.method=srmPrepareToPut
user.dn="/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=ddmadmin/..."
user.mapped=5001:5001,5000,5002,5003,5004
request.token=-1173404901
status.code=SRM_SUCCESS
client-info=job-id=3f9c8e5a-f280-52ad-909f-9008d4687100;file-id=745545721;retry=0
user-agent="fts_url_copy/3.4.2 gfal2/2.10.3 srm-ifce/1.23.3 gSOAP/2.7"
```





Request lifetime workaround

```
# Maximum client-assumed average bandwidth  
#  
# A value of 0 switches off this feature.  
#  
srm.request.maximum-client-assumed-bandwidth = 0
```

```
level=INFO
ts=2016-04-05T19:27:18.837+0200
event=org.dcache.srm.request
session=Tw:1906729:srm2:getSpaceTokens
socket.remote=130.246.182.187:55023
request.method=srmGetSpaceTokens
user.dn="/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=ddmadmin/..."
user.mapped=5001:5001,5000,5002,5003,5004
status.code=SRM_SUCCESS
client-info=job-id=32009c70-3d6f-5869-94d2-7a465008ec04;file-id=743688994;retry=0
user-agent="fts_url_copy/3.4.2 gfal2/2.10.3 srm-ifce/1.23.3 gSOAP/2.7"
```

```
level=INFO
ts=2016-04-05T01:50:33.796+0200
event=org.dcache.srm.request
session=Tw:990999:srm2:getSpaceMetaData
socket.remote=128.142.154.39:50434
request.method=srmGetSpaceMetaData
user.dn="/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=ddmadmin/..."
user.mapped=5001:5001,5000,5002,5003,5004
status.code=SRM_SUCCESS
user-agent=gSOAP/2.7
```

--- storage (dCache plugin for SRM) ---

Custom reverse DNS lookup cache: CacheStats{hitCount=0, missCount=0, loadSuccessCount=0, loadExceptionCount=0, totalLoadTime=0, evictionCount=0}

Space token by owner cache: CacheStats{hitCount=133917, missCount=26463, loadSuccessCount=44859, loadExceptionCount=0, totalLoadTime=161699959455, evictionCount=19615}

Space by token cache: CacheStats{hitCount=159372, missCount=21, loadSuccessCount=5399, loadExceptionCount=0, totalLoadTime=18665408488, evictionCount=7}

```
--- scheduler-get (Scheduler for GET operations) ---
Queued ..... 0 [Queued]
In progress (max 1000) ..... 0 [InProgress]
Queued for transfer ..... 0 [RQueued]
Waiting for transfer (max 50000) .... 802 [Ready]
-----
Total requests (max 200000) ..... 802

Scheduling strategy          : inprogress-fair-share
Transfer strategy            : fair-share
```

Discriminator

```
# uid    The request submitter's UID.  
#  
# gid    The request submitter's primary GID.  
#  
# dn     The request submitter's distinguished name.  
#  
# fqan   The request submitter's primary FQAN.  
#  
# user   The request submitter's mapped user name.  
#  
# vo     The request submitter's VO association as  
#         extracted from the primary FQAN.  
#  
srm.plugins.discriminator = dn
```



Scheduling strategies

```
# fifo    First in, first out.
#
# lifo    Last in, first out.
#
# throughput-fair-share
#
#         Round robin scheduler providing fair sharing of the request rate.
#
# inprogress-fair-share
#
#         Least-inprogress-requests-first scheduler providing fair sharing
#         of the in progress slots.
#
srm.plugins.scheduler = inprogress-fair-share
```



Transfer strategies

```
# first-come-first-served
#
# The maximum number of requests in the READY state is limited by
# srm.request.*.max-transfers, but otherwise TURLs are handed out to
# whomever comes first.
#
# fair-share
#
# The maximum number of requests in the READY state is limited by
# srm.request.*.max-transfers, but slots are kept free such that each party with
# requests in RQUEUED will receive its fair share.
#
srm.plugins.transfer = fair-share
```



```
/**
 * A SPI for classes that can provide discriminating values for jobs. Typically used
 * by schedulers to group jobs.
 *
 * Implementations are discovered using Java's ServiceLoader mechanism.
 */
public interface JobDiscriminator
{
    /**
     * Given a job return a discriminating value.
     */
    @Nonnull
    String getDiscriminatingValue(Job job);

    /**
     * The key under which the discriminating value should be identified.
     */
    @Nonnull
    String getKey();
}
```




```
/**
 * SRM TURL request servicing policy.
 *
 * A transfer strategy provides provides a policy for servicing requests for TURLs.
 */
public interface TransferStrategy
{
    /**
     * Determines whether a given Job may enter the READY state.
     */
    boolean canTransfer(Job job);
}
```



```
public interface SchedulingStrategy
{
    /**
     * Adds a new job.
     */
    void add(Job job);

    /**
     * Removes and returns the ID of the next job. How the next job is
     * defined is the essence of any implementation of this interface.
     *
     * @return The ID of the next job to process, or null if there are no jobs to process.
     */
    Long remove();

    /**
     * Number of jobs in the queue.
     */
    int size();
}
```



```
public class FifoSchedulingStrategy implements SchedulingStrategy
{
    private final Queue<Long> queue = new ArrayDeque<>();

    @Override
    public synchronized void add(Job job)
    {
        queue.add(job.getId());
    }

    @Override
    public synchronized Long remove()
    {
        return queue.poll();
    }

    @Override
    public synchronized int size()
    {
        return queue.size();
    }
}
```





```
# ----- Custom kXR_Qconfig responses
#
# xrootd clients may query the server configuration using a kXR_Qconfig request.
# These key value pairs can be queried. Additional key value pairs may be added as
# needed, see the xrootd protocol specification at http://xrootd.org/ for details.
#
(prefix)xrootd.query-config = kXR_Qconfig responses
xrootd.query-config!version = dCache ${dcache.version}
xrootd.query-config!sitename = ${dcache.description}
xrootd.query-config!role = none
```

```
$ grep config packages/system-test/target/dcache/etc/layouts/system-test.conf
xrootd.query-config!manager=Mattias Wadenstein

$ xrd fs localhost query config manager
Mattias Wadenstein
```

Migration module

-cache=class

Only copy replicas with the given cache class. An empty string will match any replica without a cache class.

-meta-only

Only transfers meta data to an existing target replica. If a given file does not have any other replicas on any of the target pools, the file is skipped.


-sticky=[-]owner[, [-]owner]...

Only copy sticky replicas. Can optionally be limited to the list of owners. A sticky flag for each owner must be present for the replica to be selected. Presence of an owner may be negated by prefixing it with a minus sign; in that case the filter matches a file that does not have a sticky flag with the given owner.

Pool startup

- Now allows read only access immediately when a pool starts.
- Write access is allowed once the inventory is finished (as before).
- Performance improvements.
- Less lock contention.
- Runtime is written to the log file.
- Progress information is available in info output.

Early Discard



Other stuff

```
# ---- Default host of RDBMS used by various services.
#
# Various services need an RDBMS. Each service is recommended to use its
# own logical database, but these may or may not be hosted in the same RDBMS
# instance. This setting defines the host of the RDBMS used.
#
dcache.db.host = localhost

# ---- Default RDBMS user account by various services.
#
dcache.db.user = dcache

# ---- Default password of RDBMS used by various services.
#
dcache.db.password =

# ---- Database name
srm.db.name = srm

# ---- Pin Manager Database Name
#
pinmanager.db.name = pinmanager
```



norden

NordForsk



Nordic e-Infrastructure
Collaboration

New resilience manager

Ask Dmitry



norden

NordForsk



Nordic e-Infrastructure
Collaboration

Location manager, exotic animals on display and redundancy

Ask Gerd

What to remember from this talk

- dCache 2.16 Released June 2016
- dCache 2.10 EOL December 2016
- dCache 2.13 EOL June 2017
- We need volunteers to write a 2.13 to 2.16 upgrade guide.