LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

GRIDKA SCHOOL 2012
KARLSRUHER INSTITUT FÜR TECHNOLOGIE
KARLSRUHE

August 30, 2012

dCache Introduction Course

Christoph Anton Mitterer
christoph.anton.mitterer@lmu.de

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

I. Introduction To dCache

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

Storage Management Systems

dCache is a so called *storage management system*. These systems manage files like "normal" filesystems (for example btrfs, ext4 or XFS) do, but usually on a much higher logical level.

Storage management systems are characterised by the following main attributes:
- Manage much higher amounts of data than normal filesystems usually do.
- Integrate many storage media that can be of different type into one system.
- Utilise **H**ierarchical **S**torage **M**anagement (for example hard disk ⇄ tape).
- Provide mechanisms to automatically increase the performance and balance the load (for example "auto-replication").
- Provide means for resilience and high availability.
- Provide advanced control systems to manage the data as well as the data flow.
- May provide interfaces to several access-protocols (for example POSIX or SRM).

As a component of a computing centre or a grid, an instance of a storage management system (including the actual data) is often referred to as "*S*torage *E*lement" (in contrast to the "*C*ompute *E*lement").

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

Overview And History

- dCache is a highly sophisticated storage management system written in Java.
- It is being actively developed at DESY since 2003.
  Major contributions come from FNAL and NDGF.
- It uses standard open source technologies in several places.
- It is released under the GNU Lesser General Public License version 3.
- It is one of the main storage management systems used within the European Middleware Initiative and the LHC Computing Grid and it is used in different other places.
- Support is available through the developers, dedicated support groups and mailing lists.

Major release-branches of dCache include:
- 1.9.12: Old stable branch.
- 2.2: Current stable branch.
- 2.3: Development branch.

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

FUNCTIONALITIES AND FEATURES

The main functionalities and features of dCache include:

- Management of enormous amounts of data (no hard limit, tested in the petabyte-range).
- **H**ierarchical **S**torage **M**anagement
- Scalability, reasonable performance, modular design and portability.
- Highly configurable and customisable, thus allowing very mighty setups.
- Some support for the usage of different database management systems.
- Use of "normal" filesystems (for example btrfs, ext4 or XFS) to store the actual data on the storage-nodes.
- X.509-based authentication through **G**rid **S**ecurity **I**nfrastructure and powerful authorisation systems with support for access control lists.
- Kerberos-based authentication.
- Firewall-friendly by utilising the concept of door-nodes.
- Access interfaces for NFS 4.1 (pNFS), HTTP and WebDAV, GridFTP ("gsiFTP"), xrootd ("proofd") and SRM (version 1 to 2.2) in addition to its native protocols DCAP and gsiDCAP.
- Scriptable administration interface with terminal-based and graphical front-ends.

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

Functionalities And Features

- Detailed logging and debugging as well as accounting and statistics.
- XML information system with detailed live information about the cluster.
- Web-interface with live summaries of the most important information.
- Checksumming of data.
- Easy migration and management of data via the "migration module".
- Resilience and high availability can be implemented on different ways by having multiple replicas of the same files.
- Powerful cost calculation systems that allows to control the data flow (from and to pools, between pools and also between pools and tape).
- Load balancing and performance tuning by "hot pool replication" (via cost calculation and replicas created by pool-to-pool-transfers).
- Garbage collection of replicas, depending on their flags, age, et cetera.
- Space management and support for space tokens.

This course does not cover the HSM-capabilities ("writing to tape") of dCache.

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

FUTURE AND PLANNED DEVELOPMENTS

dCache is under active development with the current focus on the following topics:
- Improving or rewriting parts of the existing code.
- Improving the NFS 4.1 (pNFS) interface, to support GSI, Kerberos and other features.
- Further deploying standardised technologies in many areas of dCache.
- Et cetera.

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

Organisations That Are Using dCache

Some organisations that are using dCache include:

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

Contact Addresses

You should note the following contact addresses:

- dCache Website
  `http://www.dcache.org/`
  (includes the "dCache-Book" and a dCache-Wiki)
- dCache Support
  `support@dcache.org`
- dCache User Mailing List
  `user-forum@dcache.org`
- Mass Storage Support Group of the Helmholtz-Gemeinschaft Deutscher Forschungszentren

  through the mailing list:
  `german-support@dcache.org`

  through the Global Grid User Support:
  `http://ggus.eu/`

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

dCache's Structure

dCache is a modularised system that consists of many different components, where each component fulfils a dedicated task. Depending on the specific task, a component may occur exactly once or multiple times per cluster.

A dCache cluster is structured by the following three classes of objects:
- Domains
  A domain consists of one or more cells, grouping them (typically logically) together.
  Each domain corresponds to exactly one Java VM and is thus bound to exactly one node of the cluster.
  Nodes can generally run multiple domains.
- Cells
  A cell is the actual entity that fulfils a specific task and thus corresponds roughly to one of the components mentioned above.
  Depending on the cell-type there may be multiple instances per cluster, possibly each running with a different set of settings.
  Every cell is always a member of exactly one domain.

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

DCACHE'S STRUCTURE

- Services
  A service is provided by one or more cells (possibly from different domains), grouping them semantically together.
  The cells may be mandatory or optional for its functionality.
  Depending on the service-type there may be multiple instances per cluster, possibly each running with a different set of settings.

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

Service- And Cell-Classes

The different types of services and cells are divided in the following classes:

- Core-services and -cells
  These implement basic or common tasks for the whole cluster, for example inter-cell-communication, access to the file hierarchy provider or accounting.
  Many types of core-services and -cells can exist only once per cluster.
- Door-services and -cells ("doors")
  These serve as gateways to the actual data and provide accessing clients with an interface speaking a given access-protocol (Depending on the protocol, its configuration and the client's request, the actual data flows either directly between the client and the pool or indirectly via the door.).
  The different types of doors can exist multiple times per cluster, typically placed on different nodes (so called "door-nodes").
- Pool-services and -cells ("pools")
  These manage the actual data on the storage-nodes (so called "pool-nodes") and can even communicate directly with the accessing clients.
  Of course, there may be many pools per cluster, but even per node.

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

IMPORTANT DOMAINS AND CORE-CELL-TYPES

The following describes some more important domains and core-cell-types:

- every domain

  The `RoutingMgr`- and `lm`-cells are used with dCache's native cell-messaging (and to some extent also with its successor JMS).

  They maintain the routes to domains and cells known to the specific "local" domain as well as network communication information.

  In the default configuration, `dCacheDomain`'s `RoutingMgr`-cell has knowledge on all routes within the cluster and its `lm`-cell acts as a server to all the other `lm`-(client-)cells.

- `dCacheDomain`

  The `PoolManager`-cell manages and configures the pools as well as the data flow ("pool selection") from and to pools, between pools and also between pools and tape.

  The `LoginBroker`-cell maintains a list of doors as well as their protocol, version, fully qualified domainname, port and some other data.

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

Important Domains And Core-Cell-Types

- `namespaceDomain`

  Contains the cells that are related to dCache's file hierarchy provider (usually "Chimera").

  The `PnfsManager`-cell (called so for historical reasons) is the general interface to the file hierarchy (and thus to its database), providing an interface to the other cells. It allows access to all the file meta data, including pathnames, file properties, ACLs and PNFS-IDs.

  The `acladmin`-cell allows managing ACLs from within dCache's administration interface.

- `dirDomain`

  The `dirLookupPool`-cell provides the listing of directories with DCAP and gsiDCAP.

- `adminDoorDomain`

  Contains several cells that compose dCache's administration interface.

- `nfsDomain`

  Typically used for a NFS 3- or NFS 4.1 (pNFS)-door.

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

Important Domains And Core-Cell-Types

- `httpdDomain`
  Contains several cells that compose dCache's web interface.
  Furthermore, the `billing`-cell (which is usually in this domain) implements dCache's accounting services.
- `infoDomain`
  Contains the cells that compose dCache's information system, which provides all kinds of live information to other services.
- `statisticsDomain`
  Contains the cells that compose dCache's statistics gathering system.
- `utilityDomain`
  Contains miscellaneous other cells.
  The `PinManager`-cell (which is usually in this domain) controls the pinning of files.
- `gPlazmaDomain`
  Contains the `gPlazma`-cell which provides functionalities for access control, that are used by doors in order to implement their access control system.

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

IMPORTANT DOMAINS AND CORE-CELL-TYPES

- `spacemanagerDomain`
  Contains the cells that compose the `spacemanager`-service, which is responsible for space management and space reservation.
- `webadminDomain`
  Contains the cells that compose the `webadmin`-service, which provides a web-based administration interface.

It should be noted, that all those domain- and cell-names as well as their "distribution" are actually freely configurable.
The above names and distribution are just some usual defaults or examples.

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

Access-Protocols And Door-Types

dCache provides accessing clients with interfaces (so called "doors") for a number of protocols. A door runs usually (but not necessarily) in one domain and all types of doors can have in principle multiple instances per cluster.

Currently, there are door-types for the following protocols:
- NFS 3
  NFS version 3 is a network filesystem, which allows rather limited access to the file hierarchy itself and the file-meta-data.
- NFS 4.1 (pNFS)
  NFS version 4.1 is a network filesystem defining parallel NFS, which allows "full" (and "direct") access to the files.
- WebDAV
  The WebDAV protocol, which is a superset of HTTP.
- DCAP
  The **dC**ache **A**ccess **P**rotocol is dCache's original and native access protocol.
- gsiDCAP
  The GSI-secured version of DCAP.

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

ACCESS-PROTOCOLS AND DOOR-TYPES

- FTP
  The **F**ile **T**ransfer **P**rotocol.
- GridFTP ("gsiFTP")
  A GSI-secured version of FTP targeted on grid-environments, which is very common in most grids.
- xrootd ("proofd")
  The protocol from `xrootd`, the **eX**tended **ROOT D**aemon, which is used within the ROOT and PROOF frameworks and in some other places.
- SRM
  Storage **R**esource **M**anager is a meta-protocol targeted at storage management systems, including many specific features like space tokens, lifetime and pinning.

  For the actual data transport another transfer-protocol (like GridFTP or gsiDCAP) is used.

Having multiple SRM-doors per cluster or more than one NFS 4.1 (pNFS)-door per node is possible, but a somewhat advanced setup. All other door-types can easily occur multiple times per cluster and node.

File Hierarchy Provider

In order to make access to the actual data possible trough a usual file hierarchy and to allow files being identified via a pathname, dCache makes use of a special service, namely the file hierarchy provider.

In the past, the separate "`pnfsd`" was used for this but it has been replaced by "Chimera", which provides much better performance and more advanced features.

Chimera keeps track of all pathname↔PNFS-ID-mappings, the pool-locations of each file's replicas as well as all the other file-meta-data (for example file-times, traditional POSIX file permission modes, ACLs and PNFS-tags).

All this information is stored within a database, managed by one of the supported database management systems.

dCache's components can access the file hierarchy and its meta-data via the previously mentioned PnfsManager.

It is furthermore possible to access both via NFS-mounts (typically at "`/pnfs/`").

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

PNFS-IDs

dCache itself uses in most places the so called PNFS-IDs to identify and manage the actual files.

Basically there are two types of PNFS-IDs:
- Legacy PNFS-IDs
  These are 96 bit numbers that were used with dCache's legacy file hierarchy provider `pnfsd`. They are still used in clusters that were migrated from `pnfsd` to Chimera.
- Current PNFS-IDs
  With Chimera a new number format was introduced, now having 144 bits.
  The name "PNFS-ID" was retained for historical reasons.

PNFS-IDs are usually written as hexadecimal numbers, for example:
                    `000100000000000152203B20` (legacy PNFS-ID)
`0000AF1AF0735E30448C87FAF373325DEDA0` (current PNFS-ID)

It is worth to note, that all replicas of a given file share the same PNFS-ID.

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

DATABASES

Apart from the actual data (which is stored as files in a filesystem), dCache uses a database in most cases to store its internal information (working- and state-data).

The following databases may be encountered:

- `dcache`
  Holds miscellaneous data, including those from the PinManager and the SpaceManager as well as any `gplazma`-service(s) and SRM-door(s).
- `chimera`
  Holds the data from dCache's file hierarchy provider Chimera.
- `replicas`
  Holds the data from dCache's ReplicaManager.
- `billing`
  Holds the data from dCache's billing system.
  This is only used when `billingToDb` is set to `yes` in the configuration.
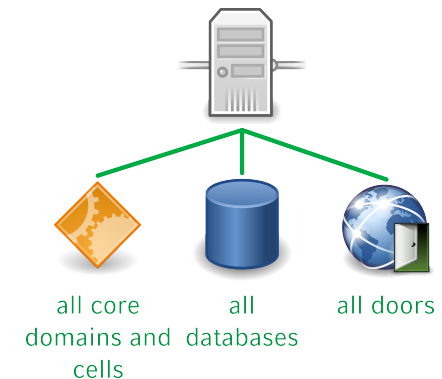
Currently, PostgreSQL is the only fully supported DBMS, but in the future generic database interfaces will be implemented that allow other DBMS being used.

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

Typical Cluster Setups – Core-Nodes

In a small cluster, the following setup might be used:
All core domains and cells, all databases and all used doors run on one single node.

+ Cheap as it requires only one node.
+ Might be easier to administrate.
- Puts a high load on this node.
- Definitely not suitable for "bigger" clusters and one could even consider whether a storage management system is the right choice.
- Single point of failure.
- Might be difficult to maintain, especially if more functionality or features shall be added.

all core domains and cells   all databases   all doors

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

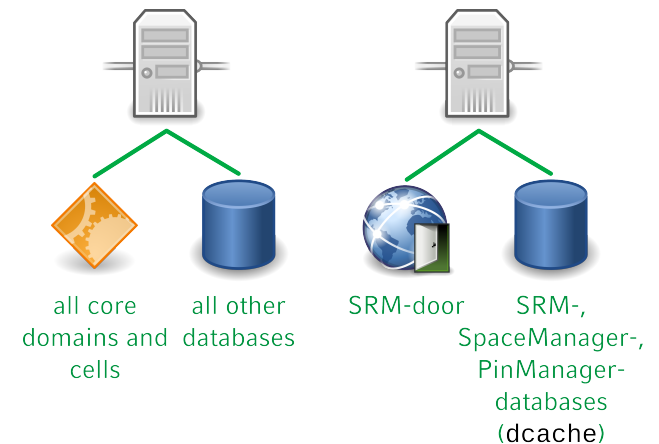DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

Typical Cluster Setups – Core-Nodes

In a medium cluster, the following setup might be used:
All core domains and cells and most databases run on one node, while the SRM-door and its related databases run on another.

+ "Reasonable" distribution of the load.
+ Great benefit for relatively few additionally required hardware-resources.
- There are still many components and probably large databases on the non-SRM-node.
- Only makes sense if SRM is heavily used.

Analogous, `namespaceDomain` and its related database (`chimera`) could be split off.

all core   all other   SRM-door   SRM-,
domains and   databases   SpaceManager-,
cells   PinManager-
                                               databases
                                               (dcache)

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
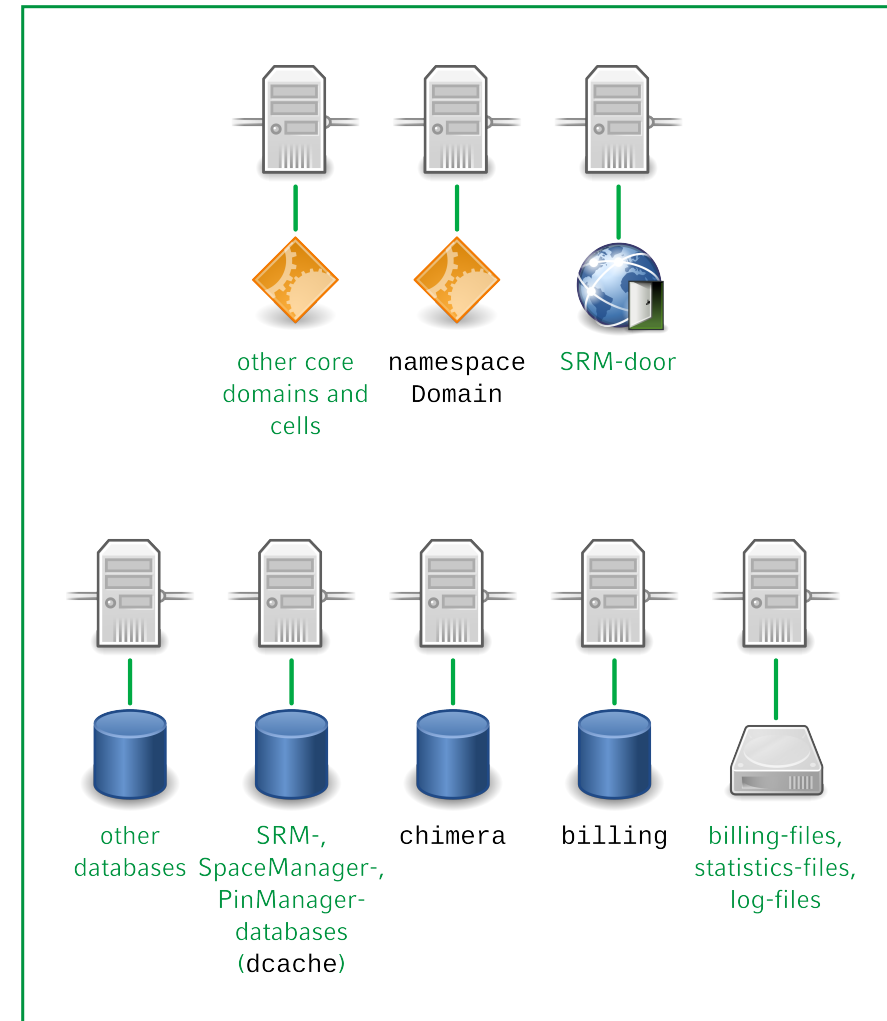INTRODUCTION TO DCACHE

Typical Cluster Setups – Core-Nodes

In a large cluster, the following setup might be used:
Most core domains and cells run on one node, but `namespaceDomain` has a dedicated node as well as the SRM-door.
The big databases have dedicated nodes, too.
The billing-files and statistics-files as well as the dCache-log-files from the miscellaneous nodes are written to a remote filesystem.

+ "Perfect" distribution of the load.
- Requires much hardware-resources "just" for the core-nodes.
- Might be difficult to administrate.
■ Splitting off the SRM-door and its related databases only makes sense if SRM is heavily used.

other core
domains and
cells

namespace
Domain

SRM-door

other
databases

SRM-,
SpaceManager-,
PinManager-
databases
(dcache)

chimera

billing

billing-files,
statistics-files,
log-files

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
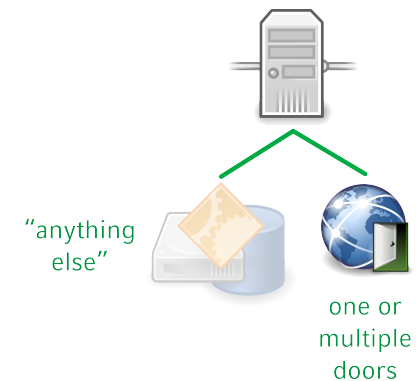INTRODUCTION TO DCACHE

Typical Cluster Setups – Door-Nodes

There are basically two ways to run a door:
- Using a dedicated node.
- Attaching it to "another" node, for example a core- or a pool-node.

In addition, it is possible in both cases to have either one or multiple doors per node.

The answer to these choices depends on at least the following factors:
- Is the specific protocol heavily used or not?
- Produces the door such a high load that it needs a dedicated node or can it coexist with the services from another node?
- Flows the actual data directly between the client and the pool or indirectly through the door (thus producing a high network load)?

one or multiple doors per node

"anything else"

one or multiple doors

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

Typical Cluster Setups – Pool-Nodes

By their nature, pools run usually on dedicated nodes (the storage-nodes).

It is however possible to run one or multiple pools per node.

The answer on the "best" pool-allocation depends on at least the following factors:

- Usually, there should be one pool per volume (for example a single hard disk or the storage exported by a RAID-controller).
- Pools should however be not to small and not to big, which is why it can make sense to unify or to split into logical volumes, each running separate pools.
- Administration of the cluster might be easier, if all pools have similar sizes.


one or multiple pools per node

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DCACHE INTRODUCTION COURSE
INTRODUCTION TO DCACHE

Data Access Model

The following is a very brief description of the access-process, the procedure that happens when a client reads or writes data to the cluster:

1. The client connects to a door with the desired protocol.
2. dCache determines the pool from which data should be read or where the data should be written to.
3. Depending on the protocol, its configuration and the client's request the actual data flows either directly between the client and the pool or indirectly via the door.
   As far as new connections are opened, these may be initiated either by the client or the server (which can be the door or the pool).
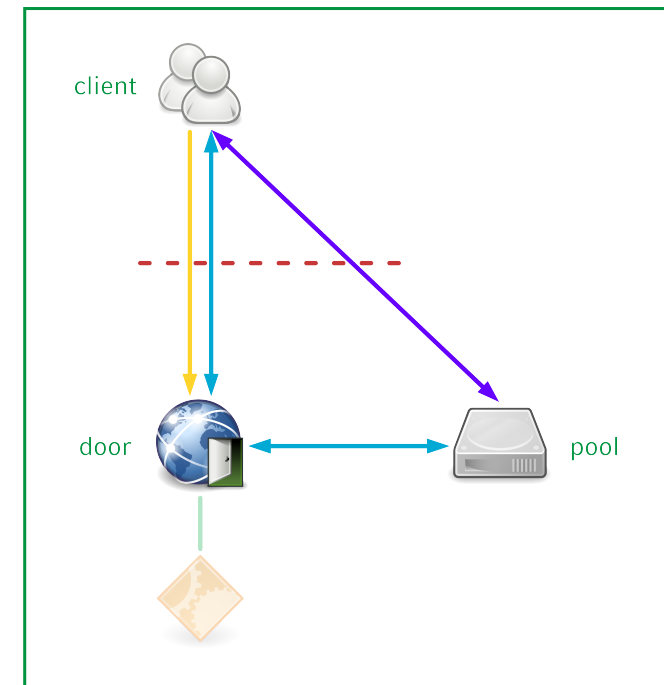
Of course there are many other steps like the authentication- and authorisation-process, accounting, et cetera.

Finis coronat opus.