Security recommendations for dCache

Paul Millar on behalf of ..

Antje, Dmitry, Christian, Gerd, Jan,
Patrick, Tanja, Tigran, Timur,
Thomas, Owen, Kai

dCache.org

dCache 1.9 releases

... along with the series support durations.

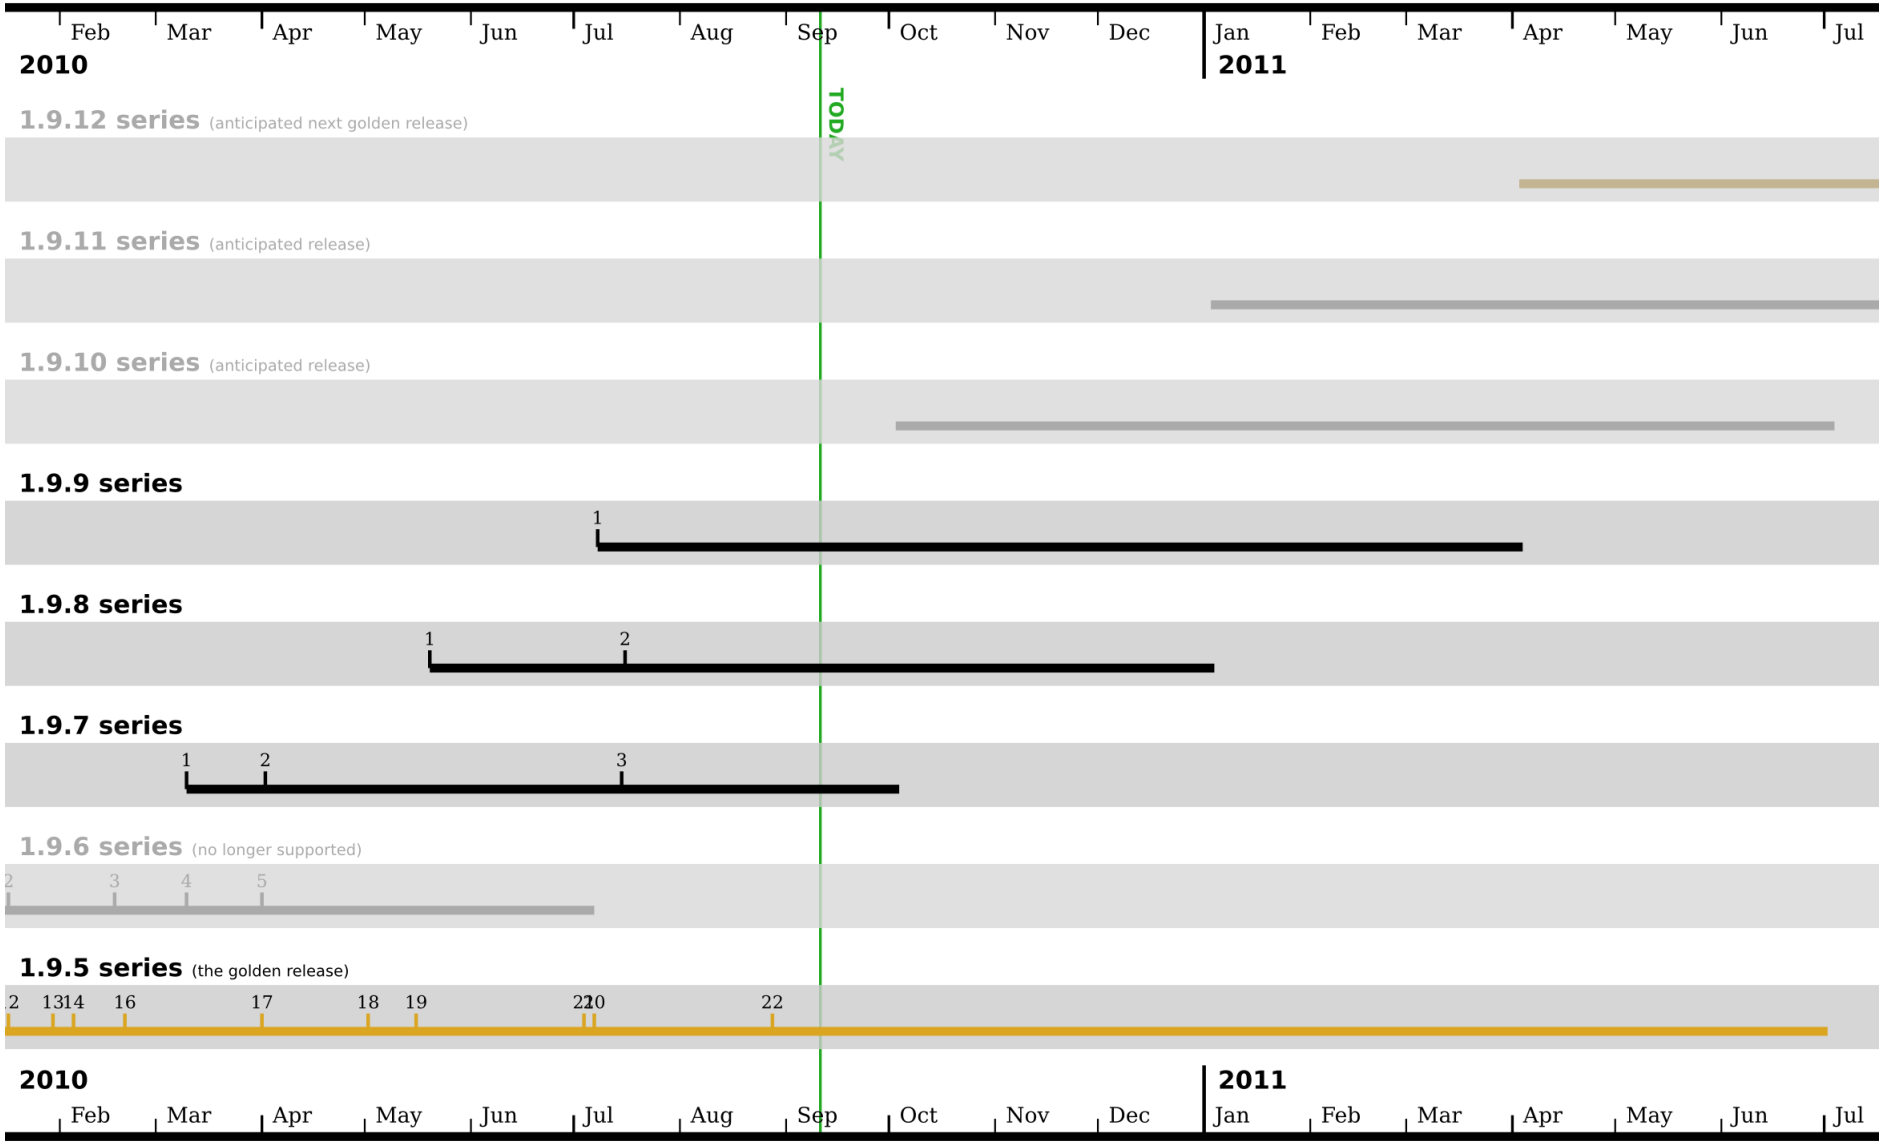

A long-lived **Golden release** and shorter supported **feature releases**. Recommendations that depend on which version is used are prefixed with "**[1.9.5]**" or "**[later]**"

Component overview

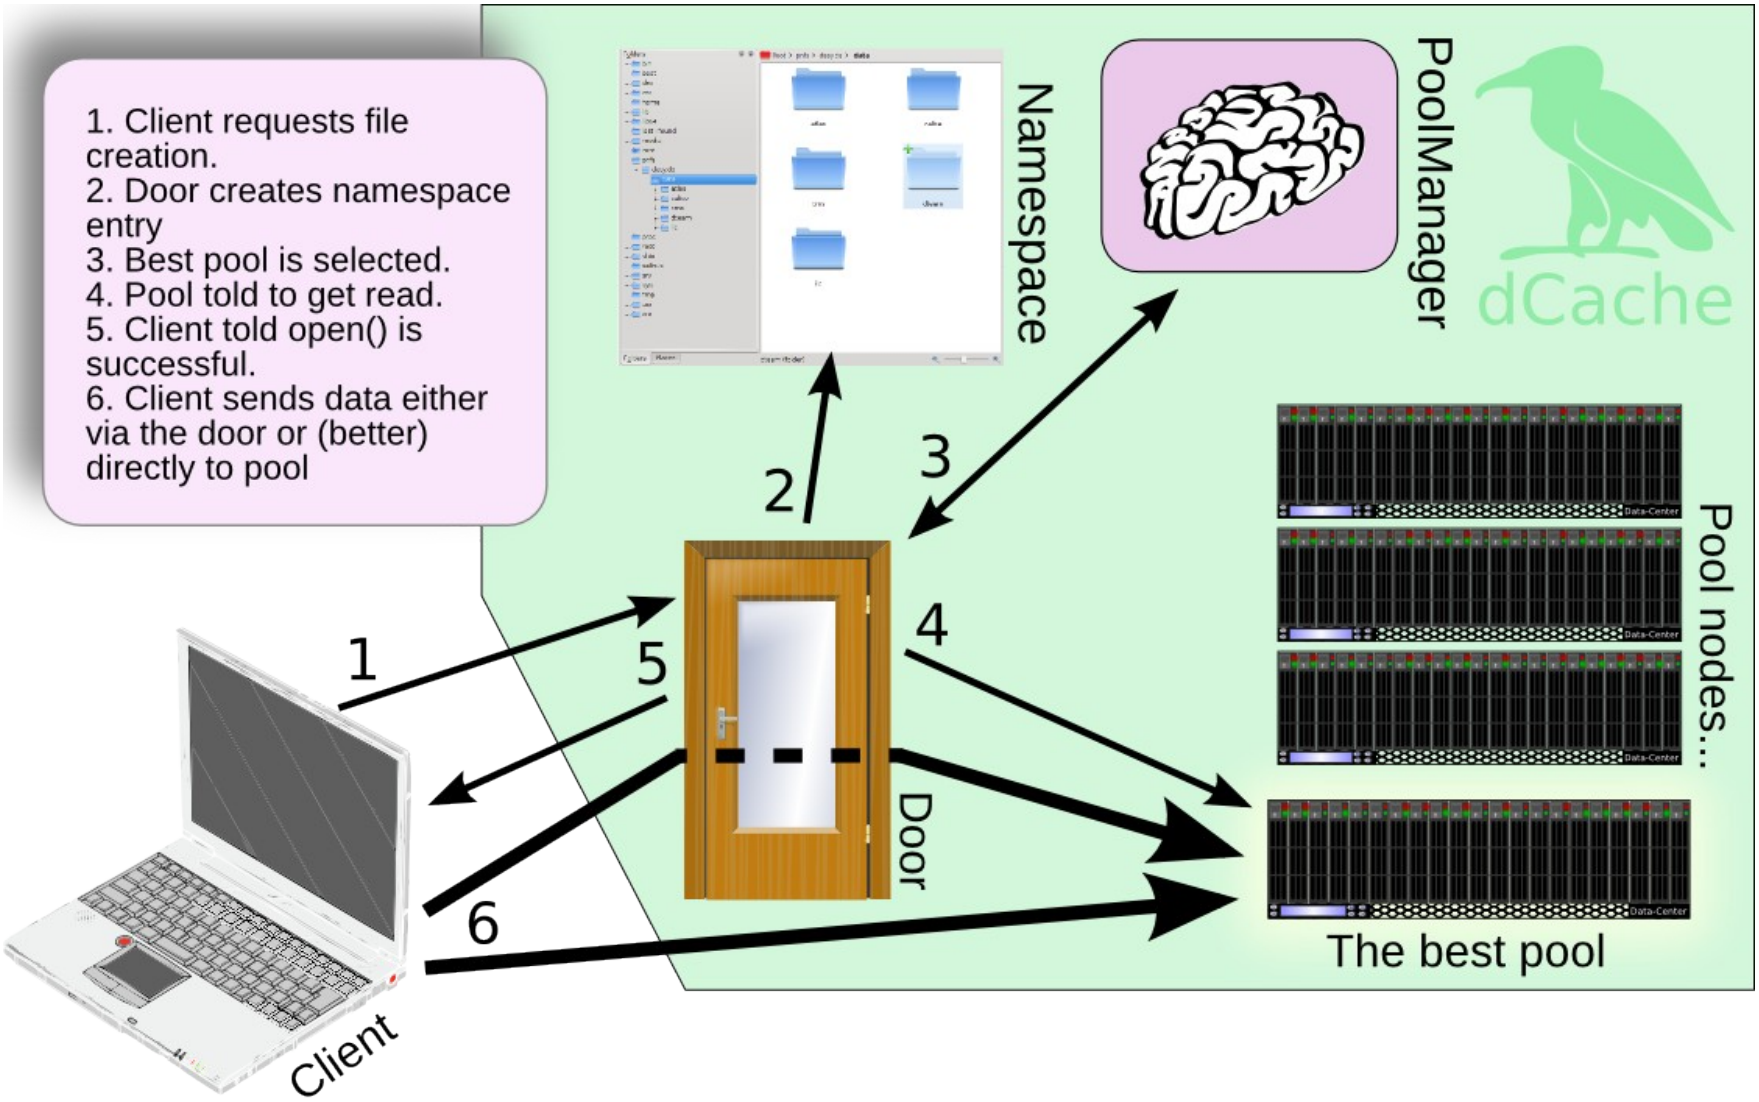

Additional, optional, component that are not shown here: SRM, SpaceManager, PinManager, TransferManager,

How to start/stop services

- dCache runs as one or more Java processes
- Starting and stopping all dCache processes:
  - /opt/d-cache/bin/dcache start
  - /opt/d-cache/bin/dcache stop
  - Chimera NFS server
    - **[1.9.5]** /opt/d-cache/libexec/chimera-nfs-run.sh
    - **[later]** NFS server is just another service.
- **[later]** can run dCache as non-root user
  - But there are some hurdles
    - Some services need to listen on <1024 port
    - Grid certificates may be readable only by root
    - Billing and statistics attempt to write to /opt/d-cache

Network usage - messaging

- dCache uses a message-passing system: cells and JMS.
- Cells:
  - Default on all current dCache releases.
  - Two steps: discovery, connectivity.
  - Discovery:
    - LB listens in dCacheDomain for UDP port 11111,
    - Remote domain:
      - sends packet to dCacheDomain,
      - receives reply which tunnels are needed to establish connectivity.
  - Connectivity:
    - dCacheDomain listens on a (by default) random port, but can be configured,
    - Remote domain establishes TCP connection,
- JMS
  - **[later]** JMS with embedded ActiveMQ: dCacheDomain listen on TCP 11112, others connect to this.

Network usage - transfers

- Doors and SRM listen on configurable ports:
  - LAN protocols:
    - NFS 2049, dcap 22125, Gsidcap 22128, Kerberos-dcap 22725
  - WAN protocols:
    - SRM 8443, GridFTP 2811, HTTP/WedDAV 2880,
    - KerberosFTP 22127, (insecure/normal) FTP 22126
- Ephemeral ports are taken from a configurable range:
  - LAN protocols use port-range (33115..33145)
  - WAN protocols use port-range (20000..25000)

Network usage – database, NFS

- Some components need access to a database
  - Those that need database access:
    - SRM, Pin Manager, space manager, replica manager
    - **[<1.9.9]** gPlazma cell.
    - If using Chimera: NFS server, PnfsManager
    - If using PNFS: PNFS (dbserver processes).
  - Restrict access to db:
    - Grant access to a username with a good password
- POSIX access to namespace (NFS mounting):
  - Only needed by some components:
    - PnfsManager (if using PNFS),
    - **[1.9.5]** SRM, dirDomain.
  - Results in NFS network traffic between namespace (PNFS and Chimera) and the node running the service.

User mapping

- gPlazma component to control user mapping.
- Several mechanisms:
  - XACML
  - SAML
  - Kpwd
  - GridMap,
  - Gplazmalite,
  - SAZ
- Any number can be enabled, site-admin supplies the order, first plugin to map wins.

How to prevent access

- Banning users is configured inside gPlazma
- Use gPlazmaLite (grid-vorolemap) with dash as username
  - This should be first in list of gPlazma plugin
- Ban a user outright:

```
"/O=ExampleGrid/OU=ExampleSite/CN=Joe Bloggs" -
```

- Ban a user as member of FQAN:

```
"/O=ExampleGrid/OU=ExampleSite/CN=Joe Bloggs" "/example.org" -
```

- Ban all members of an FQAN:
  - be aware of explicit DN mappings will not be banned.

```
"*" "/example.org/Role=rogue_user" -
```

Log files

- In general:
  - Services run within a “domain” (i.e., a JVM instance)
  - Each domain has a log file (in /var/log by default)
- SRM
  - **[1.9.5]** /opt/d-cache/libexec/apache-tomcat/logs/catalina.out
  - **[later]** normal logging for domain
- Chimera
  - **[1.9.5]** /tmp/himera.log
  - **[later]** normal logging for domain

Administrator access

- Two main means of sys-admin access: web and ssh
  - Web monitoring listens on port 2880
  - SSH admin interface listens on port 22223
- Recommend securing access to these two service from trusted hosts.
- For example:
  - Allow access to port 22223 only from localhost
  - Allow access access to port 2880 from site-admin's machine
  - Relay web access via an Apache instance listening on a secure port, require user certificate from web browser and authorise on that certificate.

Future plans (security)

- New gPlazma: more modular, more flexible
  - As part of this, support for Argus.
- Moving towards JMS:
  - Allows redundant communication,
  - Allows tunnelling messages through encrypted communication.
- New web admin interface:
  - Security model

Summary

- Run latest patch version of dCache,
- Secure network communication
  - Message,
  - Data transfers,
  - Mounted namespace (if needed).
- Know (practice) how to ban a user,
- “Know thy system”
  - To know when something is abnormal, you must
    - know what is normal.
    - know what is happening.