



FOS XD

BEST OF FUSE

Andreas-Joachim Peters
IT-SD



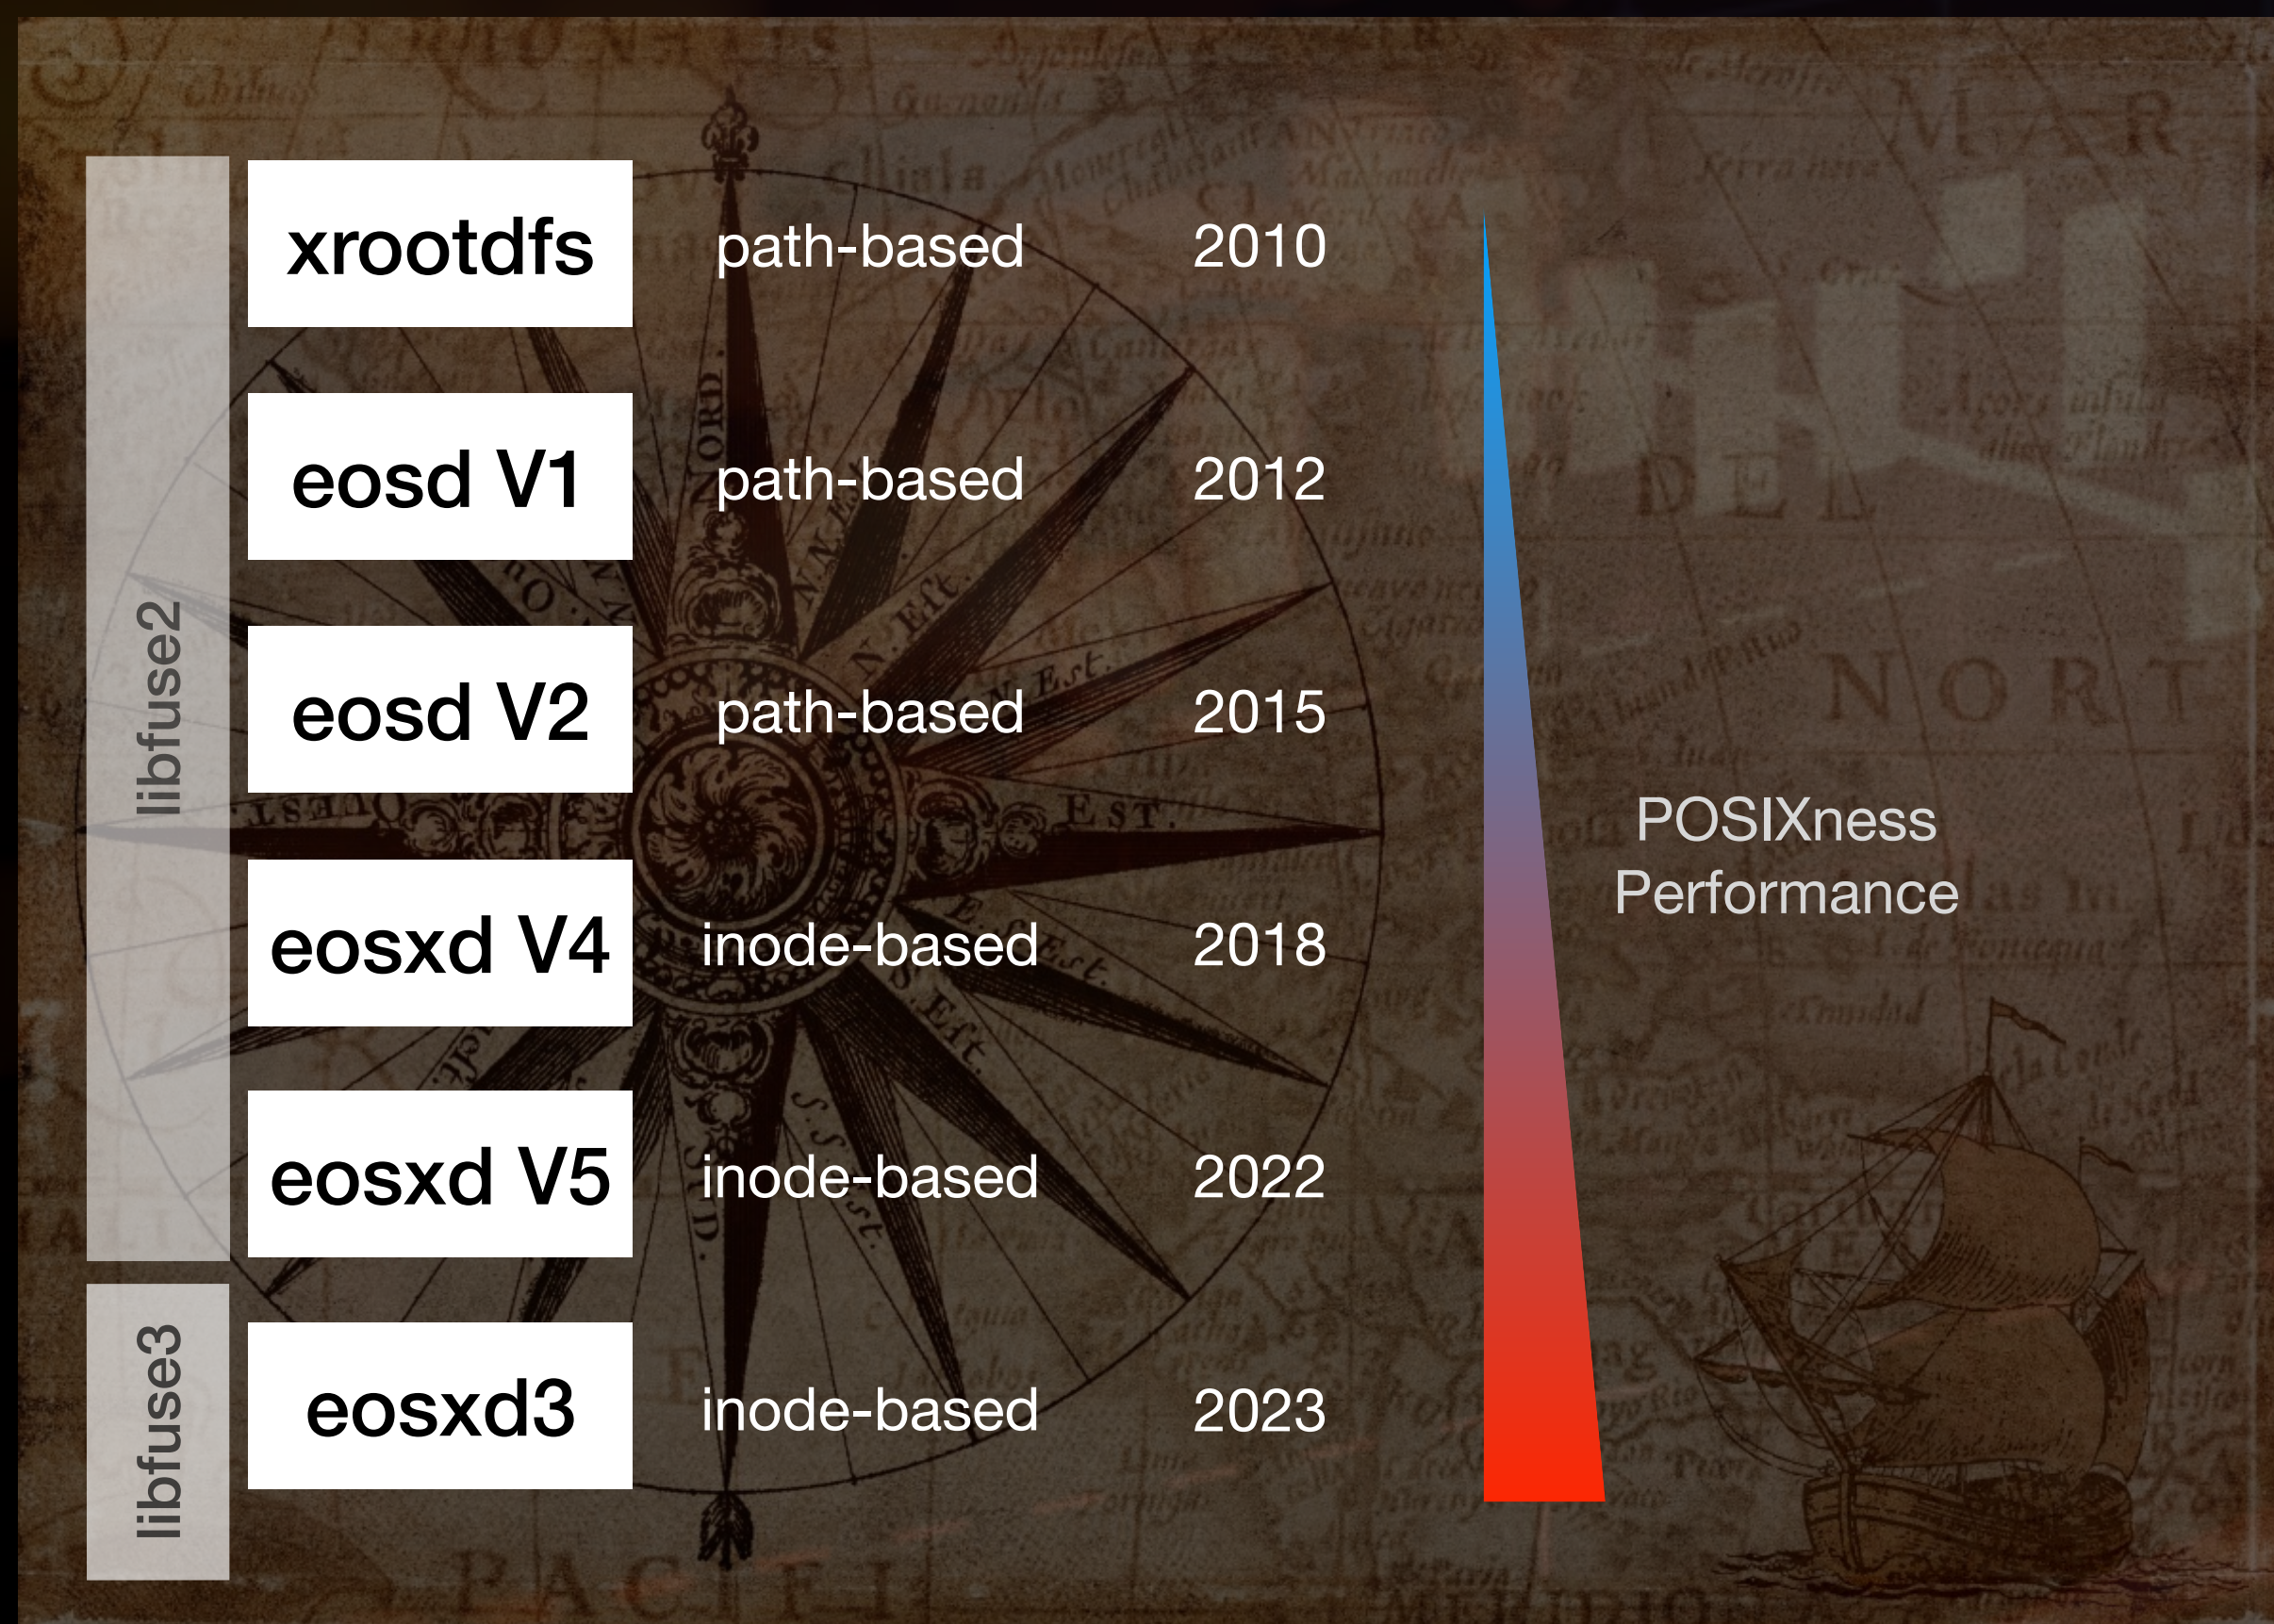
The EOS fuse filesystem

eosxd *genealogy*



The EOS fuse filesystem

eosxd genealogy




The EOS fuse filesystem

eosxd *genealogy*

FUSE is used everywhere, even in HPC computing ...
it is just too convenient, has worse reputation, than it
actually is ...

AZURE	- blobFUSE
INTEL	- daosFUSE
CEPHFS	- ceps-fuse
RCLONE	- various FUSE backends
RESTIC	- browse your backup



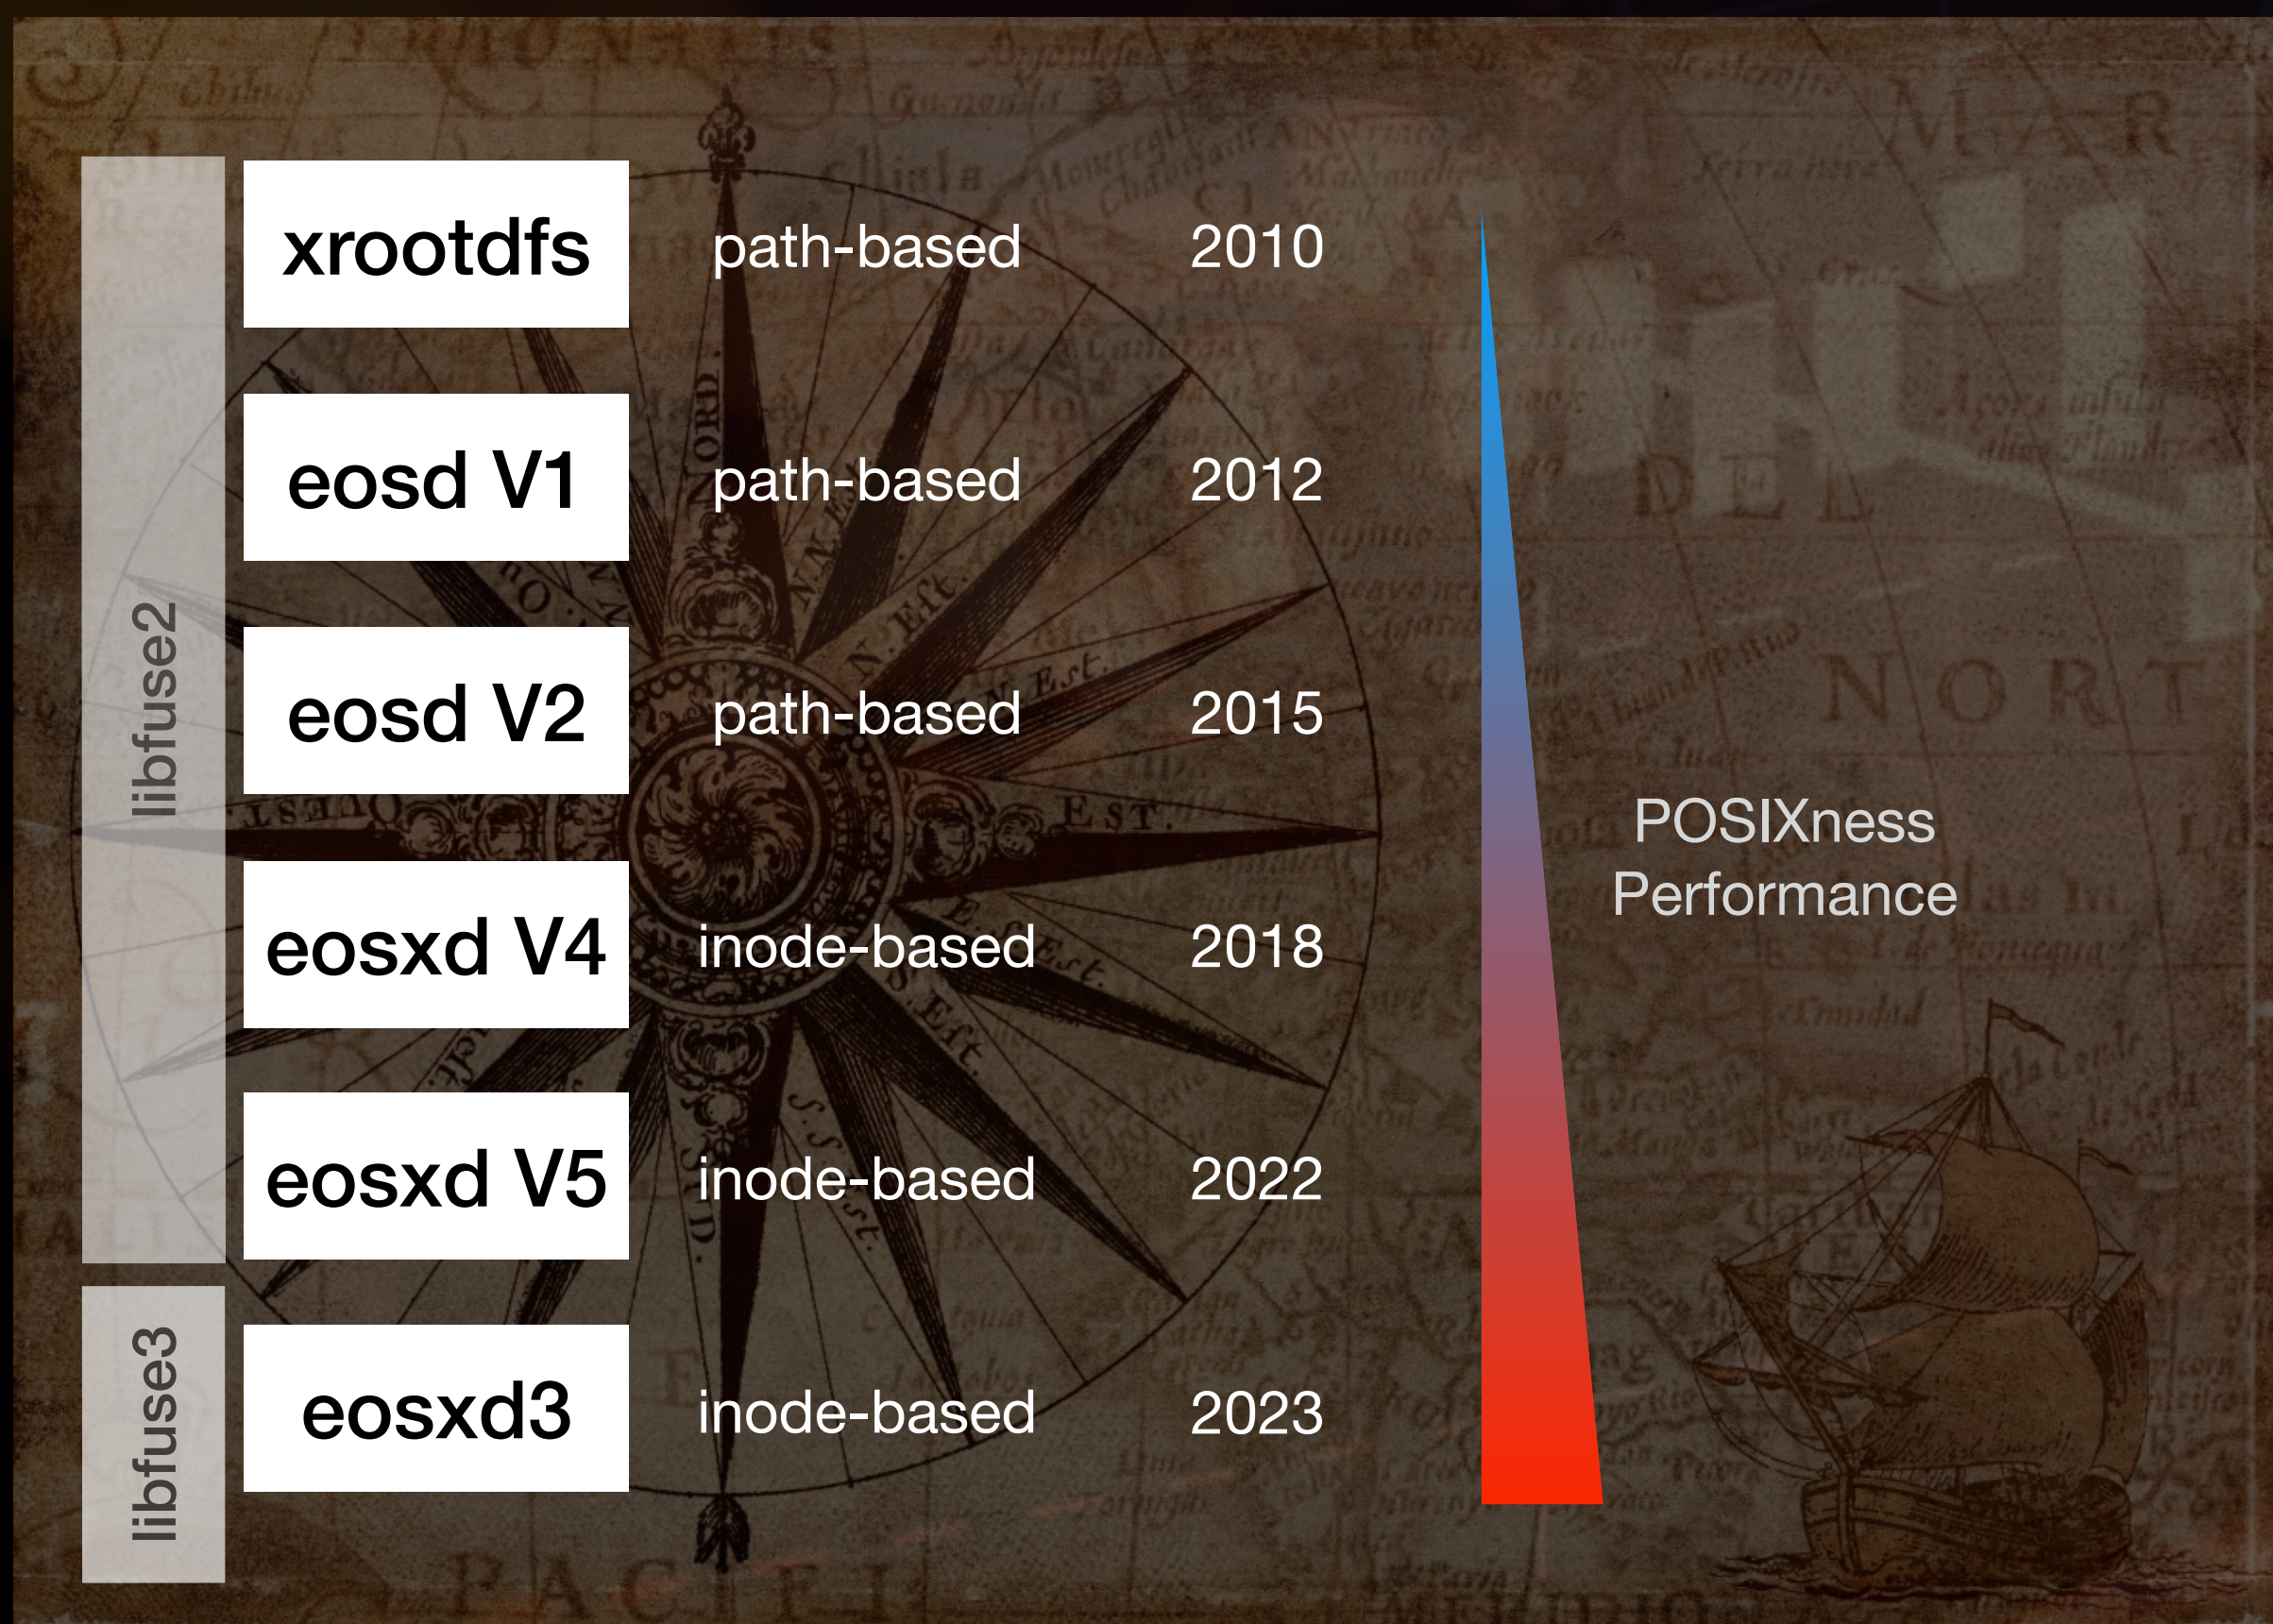
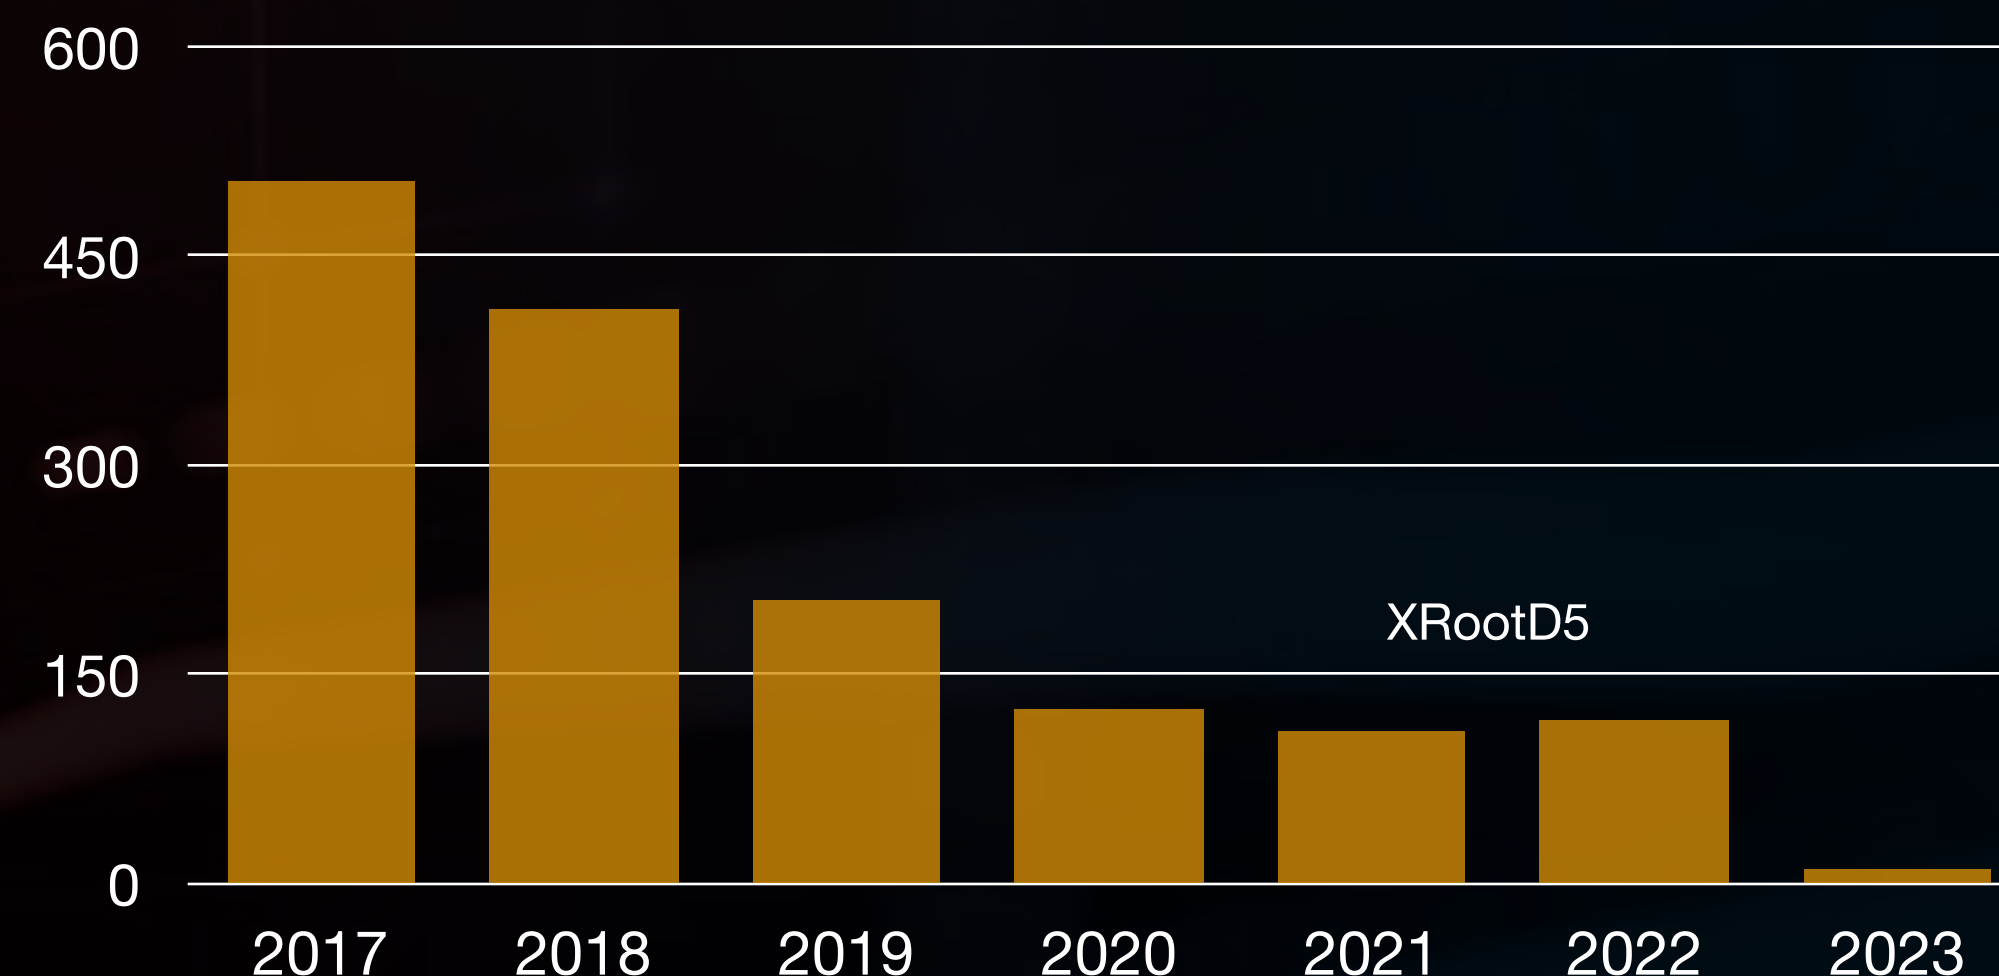
The EOS fuse filesystem

eosxd genealogy

FUSE is used everywhere, even in HPC computing ... it is just too convenient, has worse reputation, than it actually is ...

AZURE - blobFUSE
 INTEL - daosFUSE
 CEPHFS - ceps-fuse
 RCLONE - various FUSE backends GDRIVE ONEDRIVE
 RESTIC - browse your backup

EOS Fuse Dev - Commits per Year

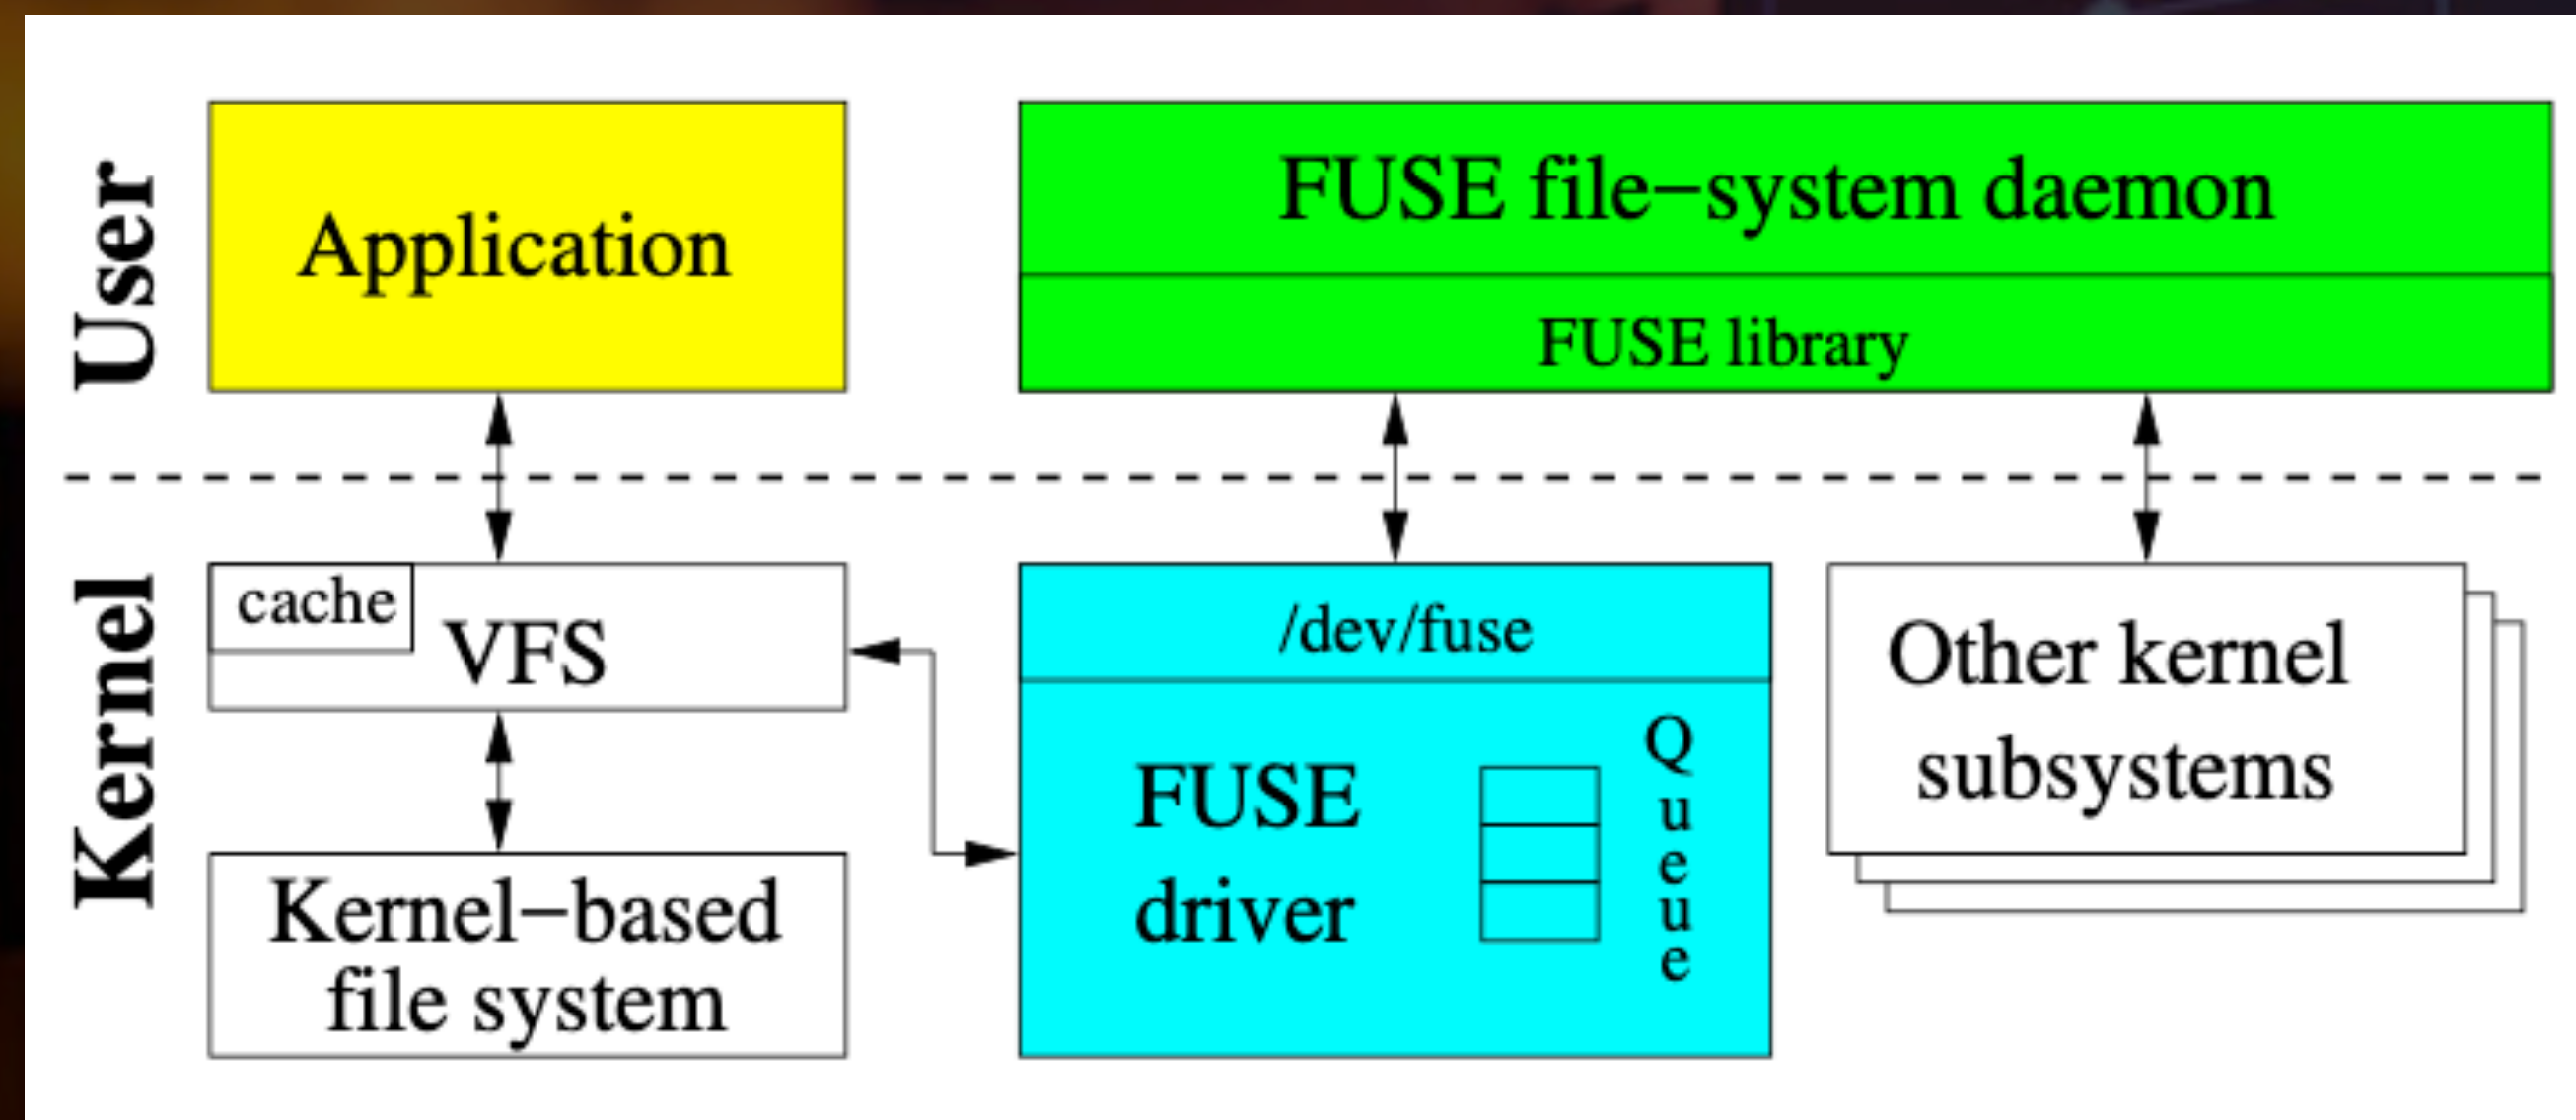


How FUSE works ...

FUSE *architecture*

How FUSE works ...

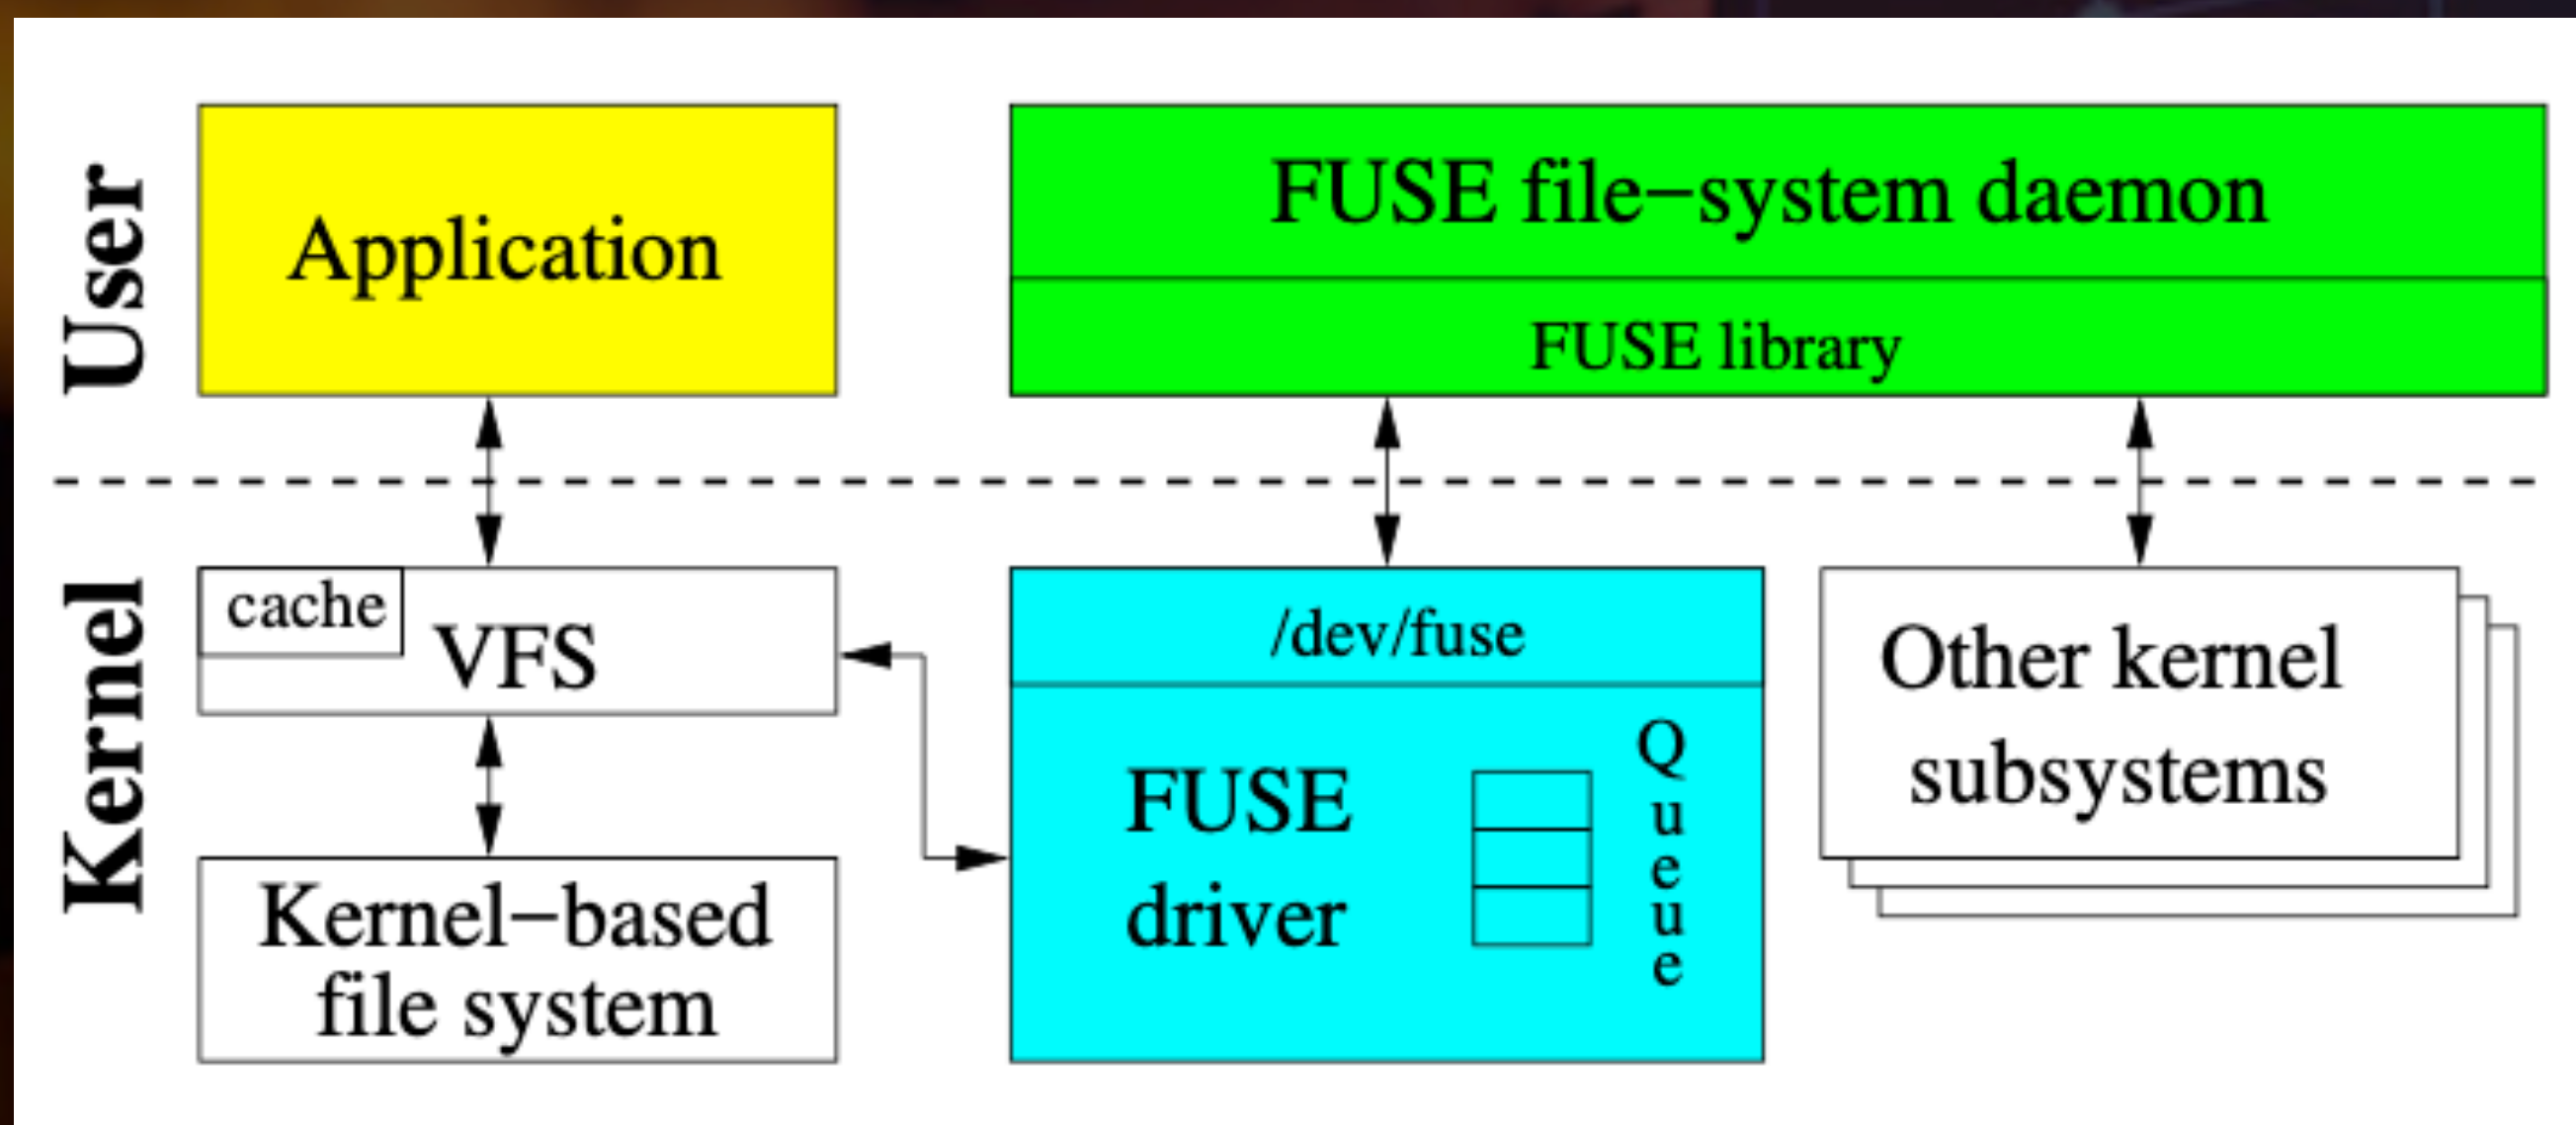
FUSE *architecture*



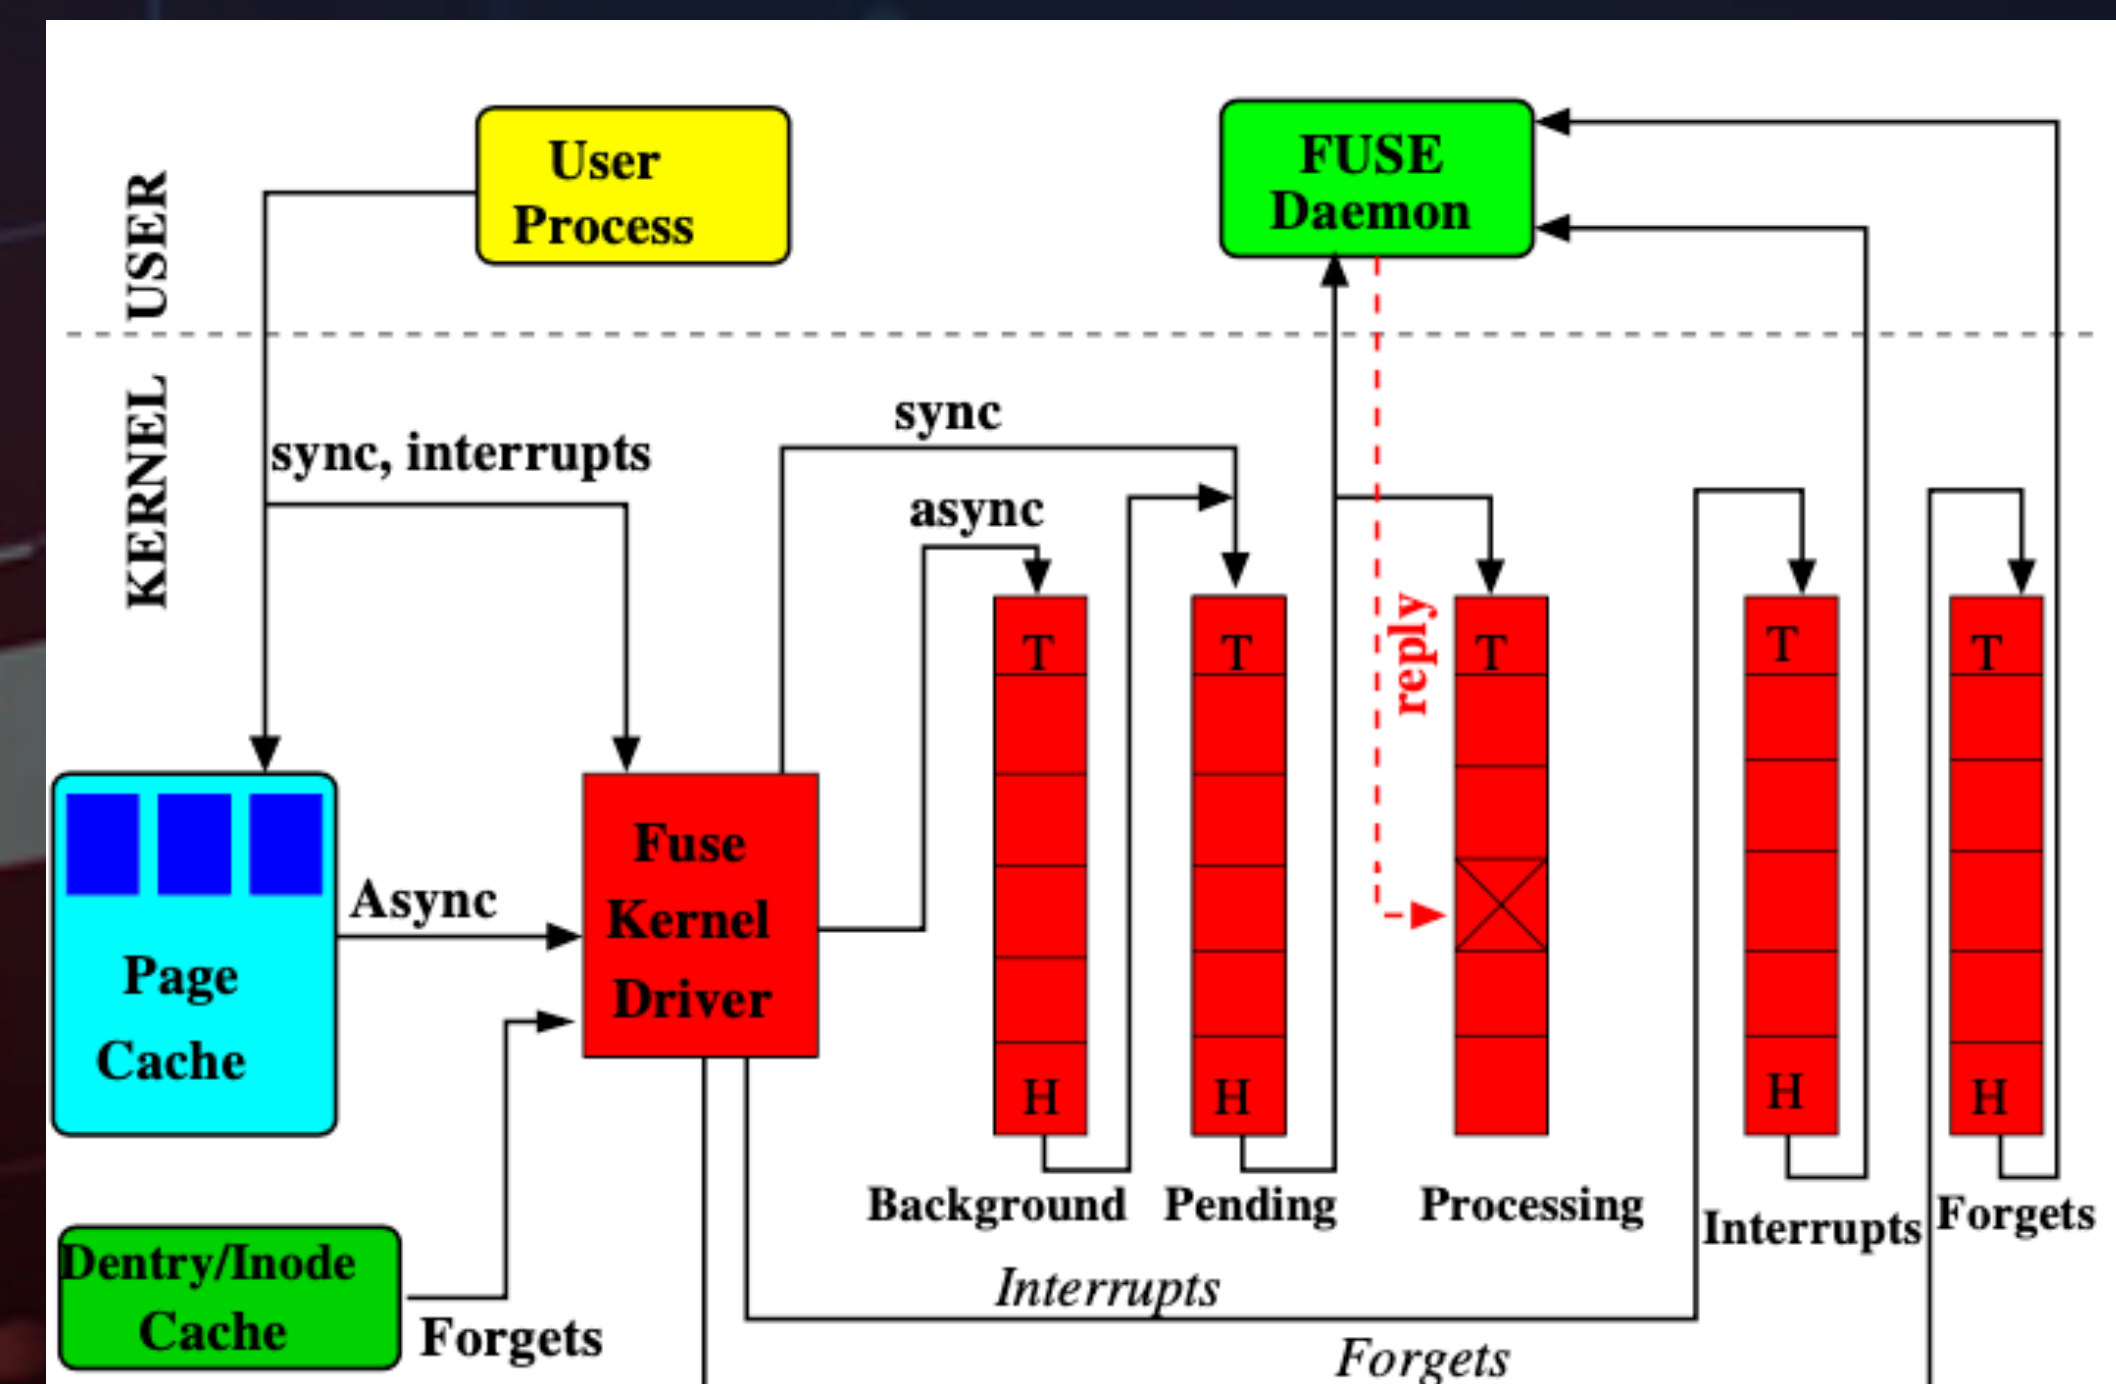
Vangoor Fast'17

How FUSE works ...

FUSE architecture



Vangoor Fast'17



Vangoor Fast'17

How FUSE works ...

FUSE *internals*

How FUSE works ...

FUSE *internals*

The main function categories in FUSE

How FUSE works ...

FUSE *internals*

The main function categories in FUSE

```
lookup(parent, name)
```


How FUSE works ...

FUSE *internals*

The main function categories in FUSE

lookup(parent, name)

Lookup FS hierarchy
using names and
associate an inode

How FUSE works ...

FUSE *internals*

The main function categories in FUSE

lookup(parent, name)

Lookup FS hierarchy
using names and
associate an inode

forget(parent, name)

How FUSE works ...

FUSE *internals*

The main function categories in FUSE

lookup(parent, name)

Lookup FS hierarchy
using names and
associate an inode

forget(parent, name)

Forget something
in the hierarchy

How FUSE works ...

FUSE *internals*

The main function categories in FUSE

lookup(parent, name)

Lookup FS hierarchy
using names and
associate an inode

forget(parent, name)

Forget something
in the hierarchy

func(inode[s])

How FUSE works ...

FUSE *internals*

The main function categories in FUSE

lookup(parent, name)

Lookup FS hierarchy
using names and
associate an inode

forget(parent, name)

Forget something
in the hierarchy

func(inode[s])

Do something
by inode(s)

How FUSE works ...

FUSE *internals*

When using the FUSE low-level API you have to implement these functions :

The main function categories in FUSE

lookup(parent, name)

Lookup FS hierarchy
using names and
associate an inode

forget(parent, name)

Forget something
in the hierarchy

func(inode[s])

Do something
by inode(s)

How FUSE works ...

FUSE *internals*

The main function categories in FUSE

lookup(parent, name)

Lookup FS hierarchy
using names and
associate an inode

forget(parent, name)

Forget something
in the hierarchy

func(inode[s])

Do something
by inode(s)

When using the FUSE low-level API you have to implement these functions :

Group (#)	Request Types
Special (3)	INIT, DESTROY, INTERRUPT
Metadata (14)	LOOKUP, FORGET, BATCH_FORGET, CREATE, UNLINK, LINK, RENAME, RENAME2, OPEN, RELEASE, STATFS, FSYNC, FLUSH, ACCESS
Data (2)	READ, WRITE
Attributes (2)	GETATTR, SETATTR
Extended Attributes (4)	SETXATTR, GETXATTR, LISTXATTR, REMOVEXATTR
Symlinks (2)	SYMLINK, READLINK
Directory (7)	MKDIR, RMDIR, OPENDIR, RE-LEASEDIR, REaddir, REaddirplus, FSYNCDIR
Locking (3)	GETLK, SETLK, SETLKW
Misc (6)	BMAP, FALLOCATE, MKNOD, IOCTL, POLL, NOTIFY_REPLY

Misc (6)

BOGГ' ИОЛІЕХ КЕБГХ

ВМВБ' ЕАГГОСАЛЕ' МКНОД' ІОСЛГ'

ГОСКІВ (2)

СЕЛК' 2ЕЛК' 2ЕЛКМ

Vangoor Fast'17

How EOS FUSE works ...

eosxd architecture



How EOS FUSE works ...

eosxd *architecture*

Low-Level FUSE API

libfuse (extern C)



How EOS FUSE works ...

eosxd *architecture*

Low-Level FUSE API

libfuse (extern C)

Timing/Tracking

tracks operations i.e. blocked IO

How EOS FUSE works ...

eosxd *architecture*

Low-Level FUSE API

libfuse (extern C)

Timing/Tracking

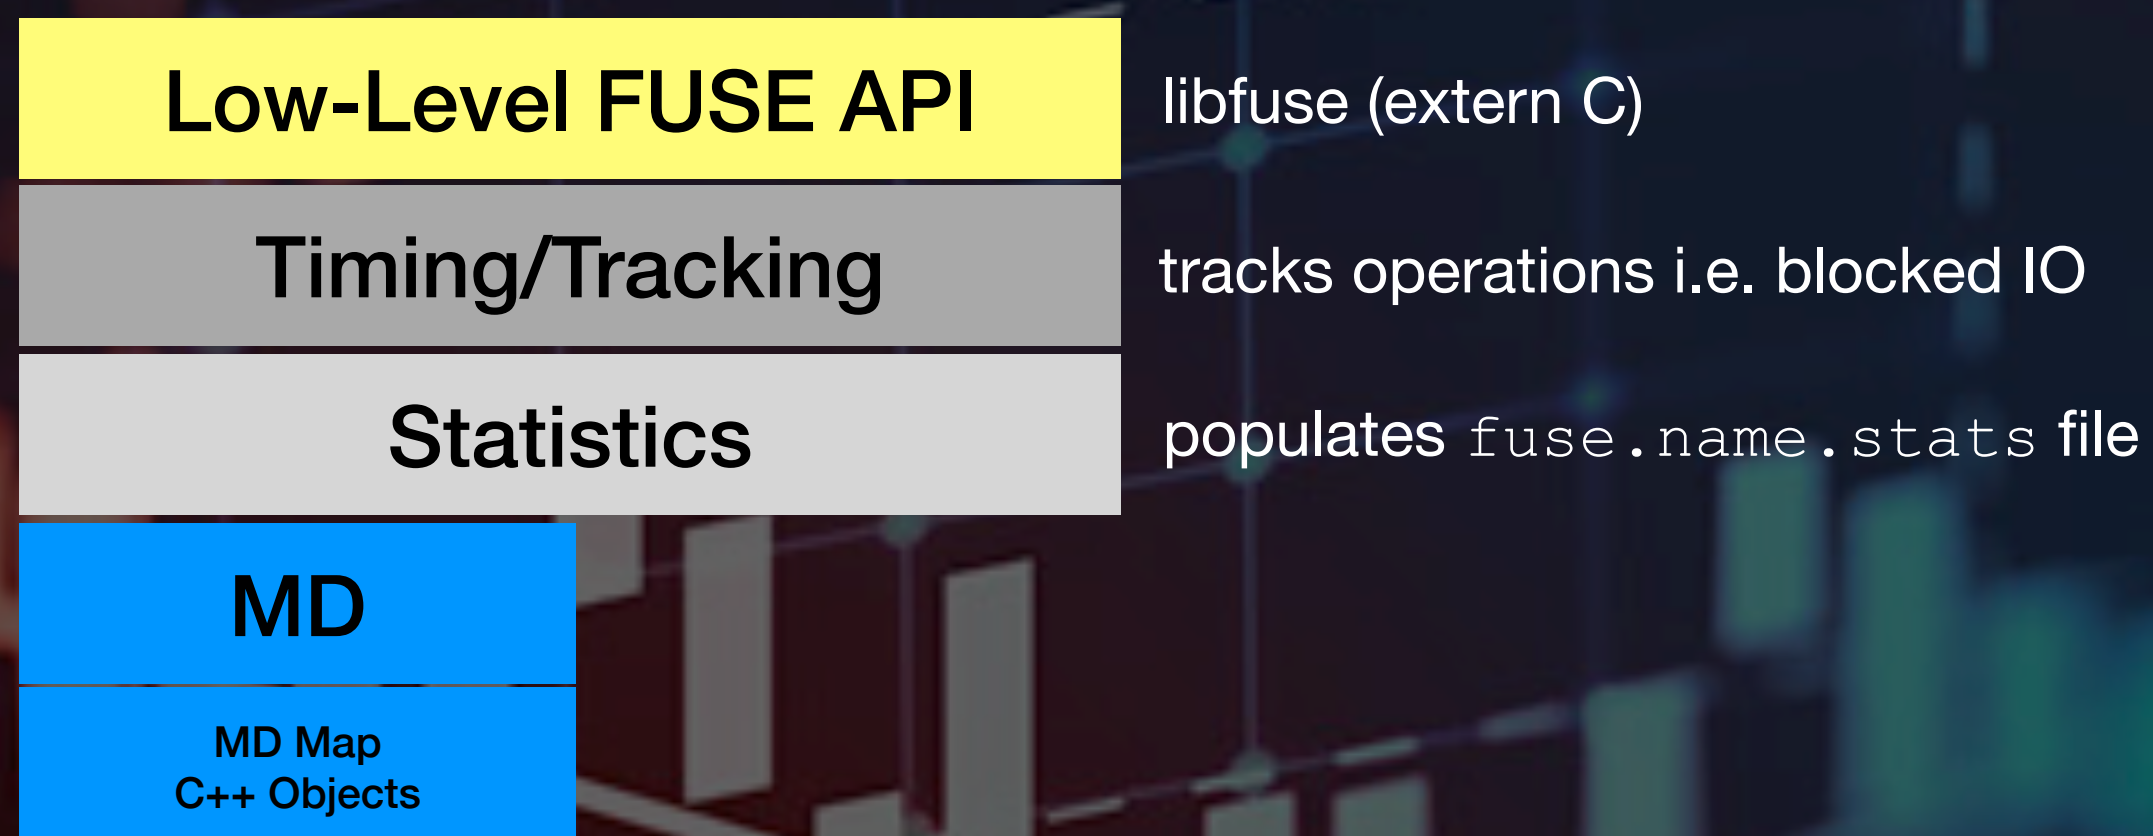
tracks operations i.e. blocked IO

Statistics

populates `fuse.name.stats` file

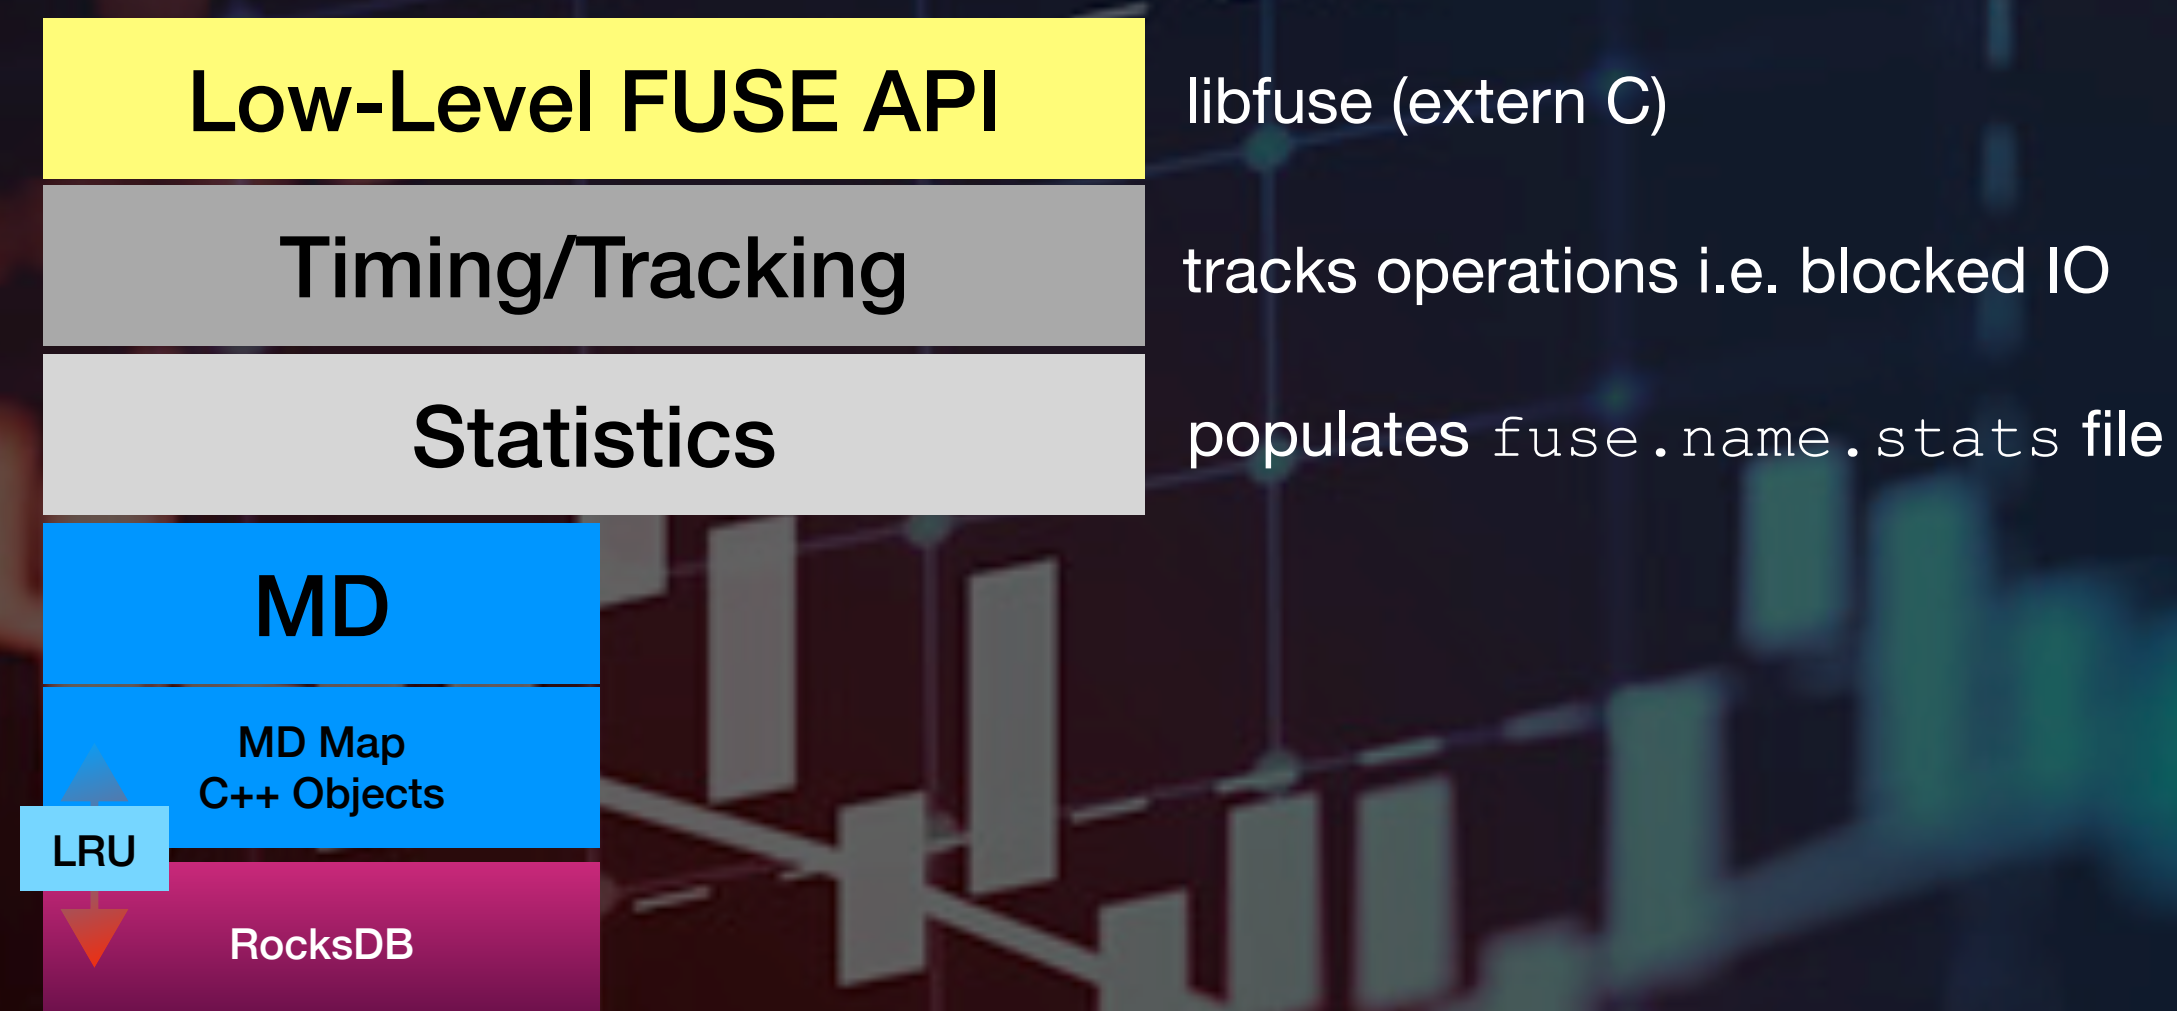
How EOS FUSE works ...

eosxd *architecture*



How EOS FUSE works ...

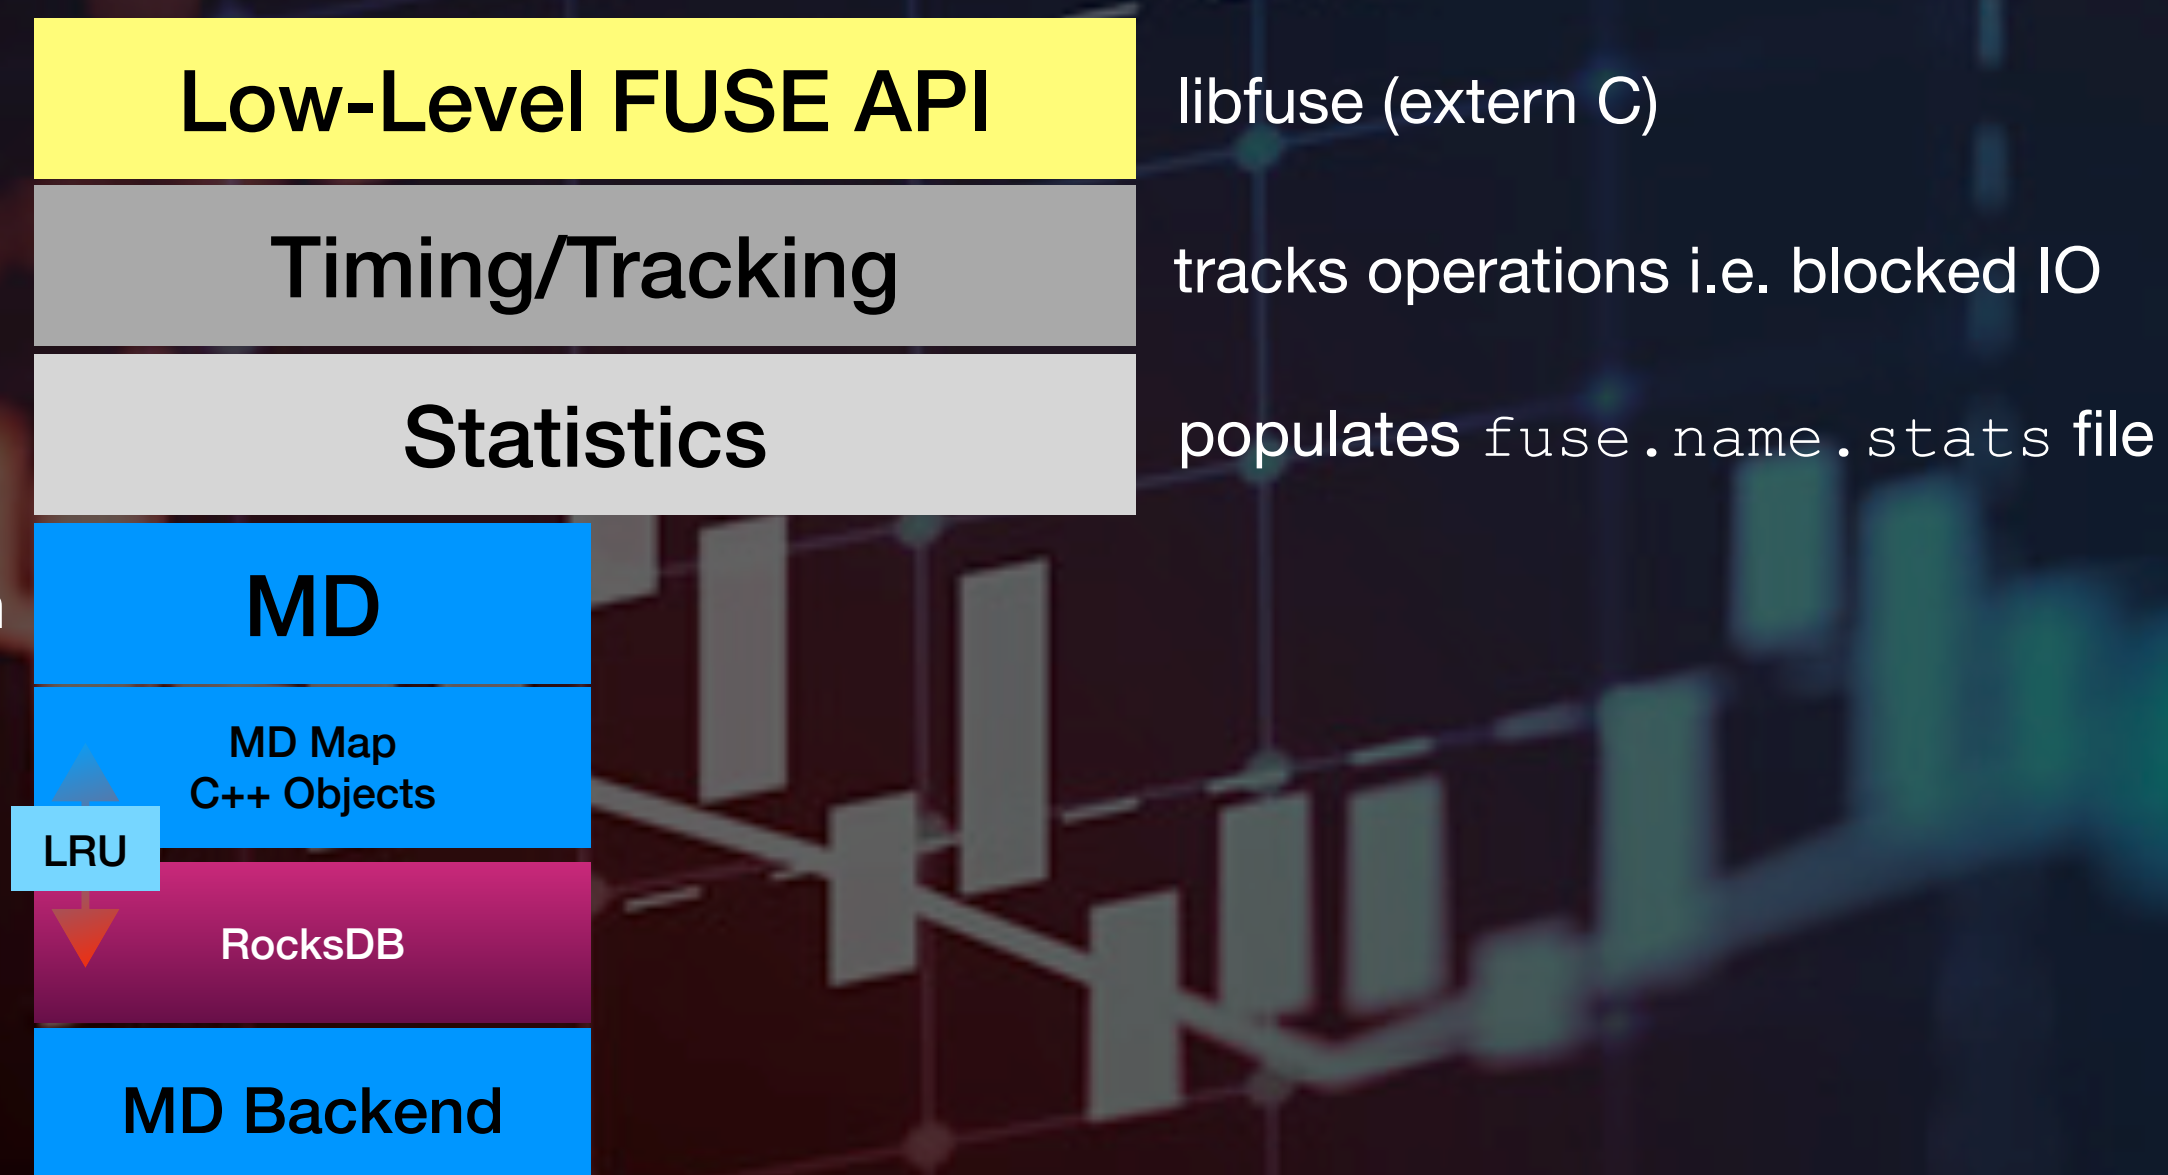
eosxd architecture



How EOS FUSE works ...

eosxd architecture

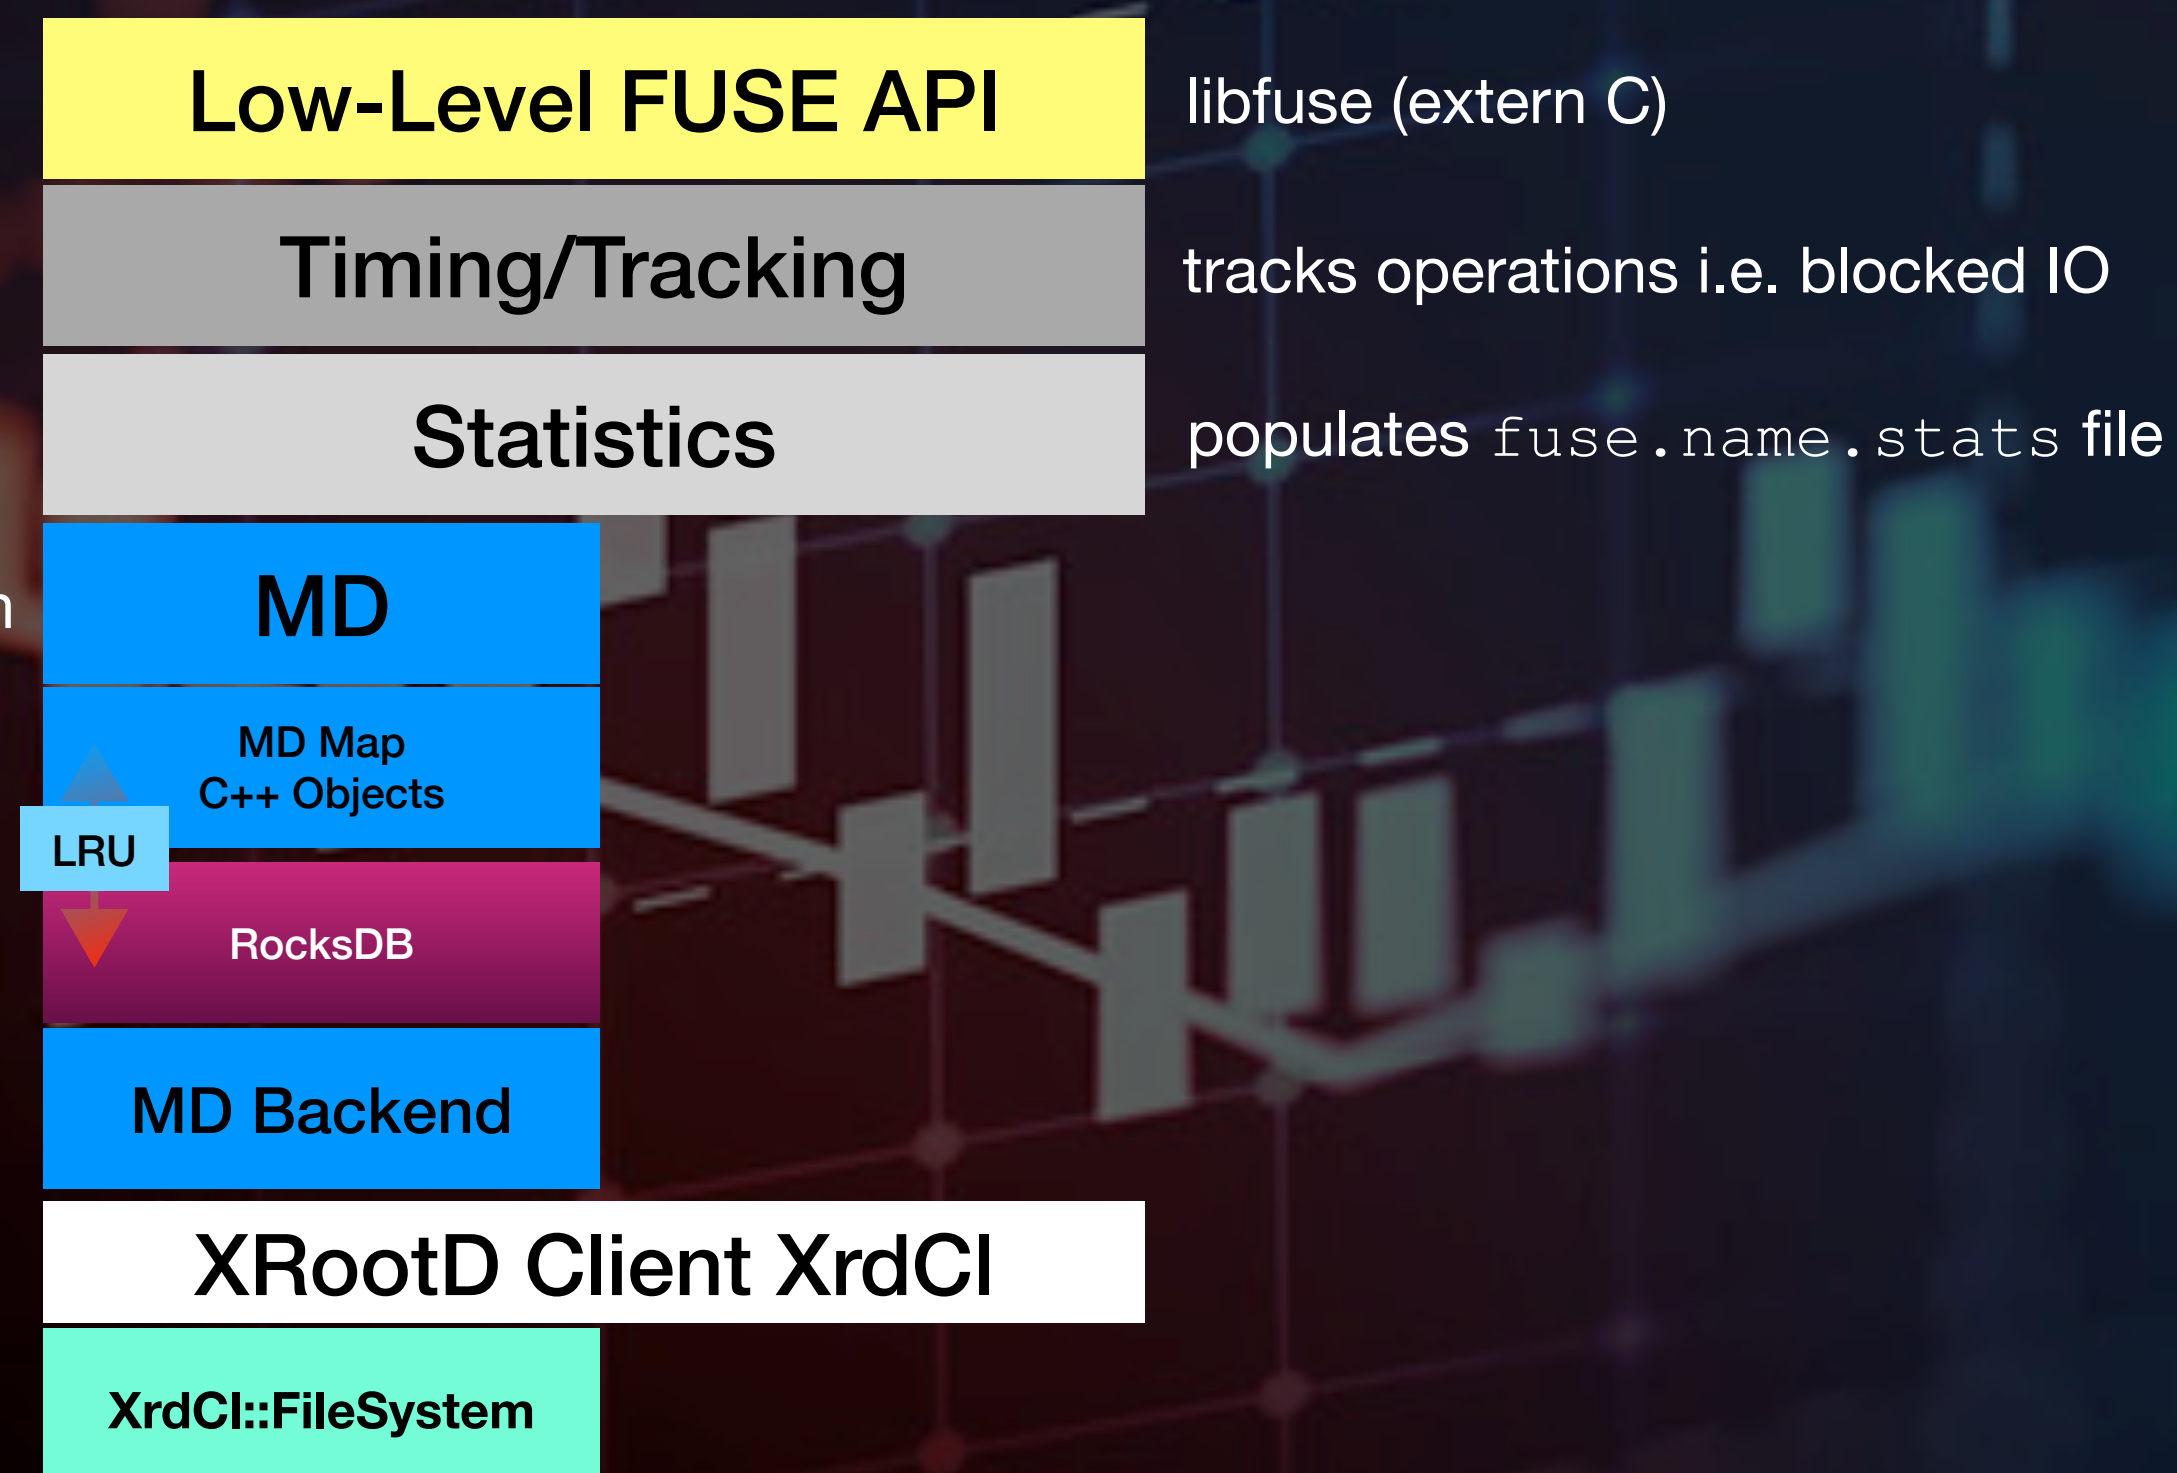
meta data IO path



How EOS FUSE works ...

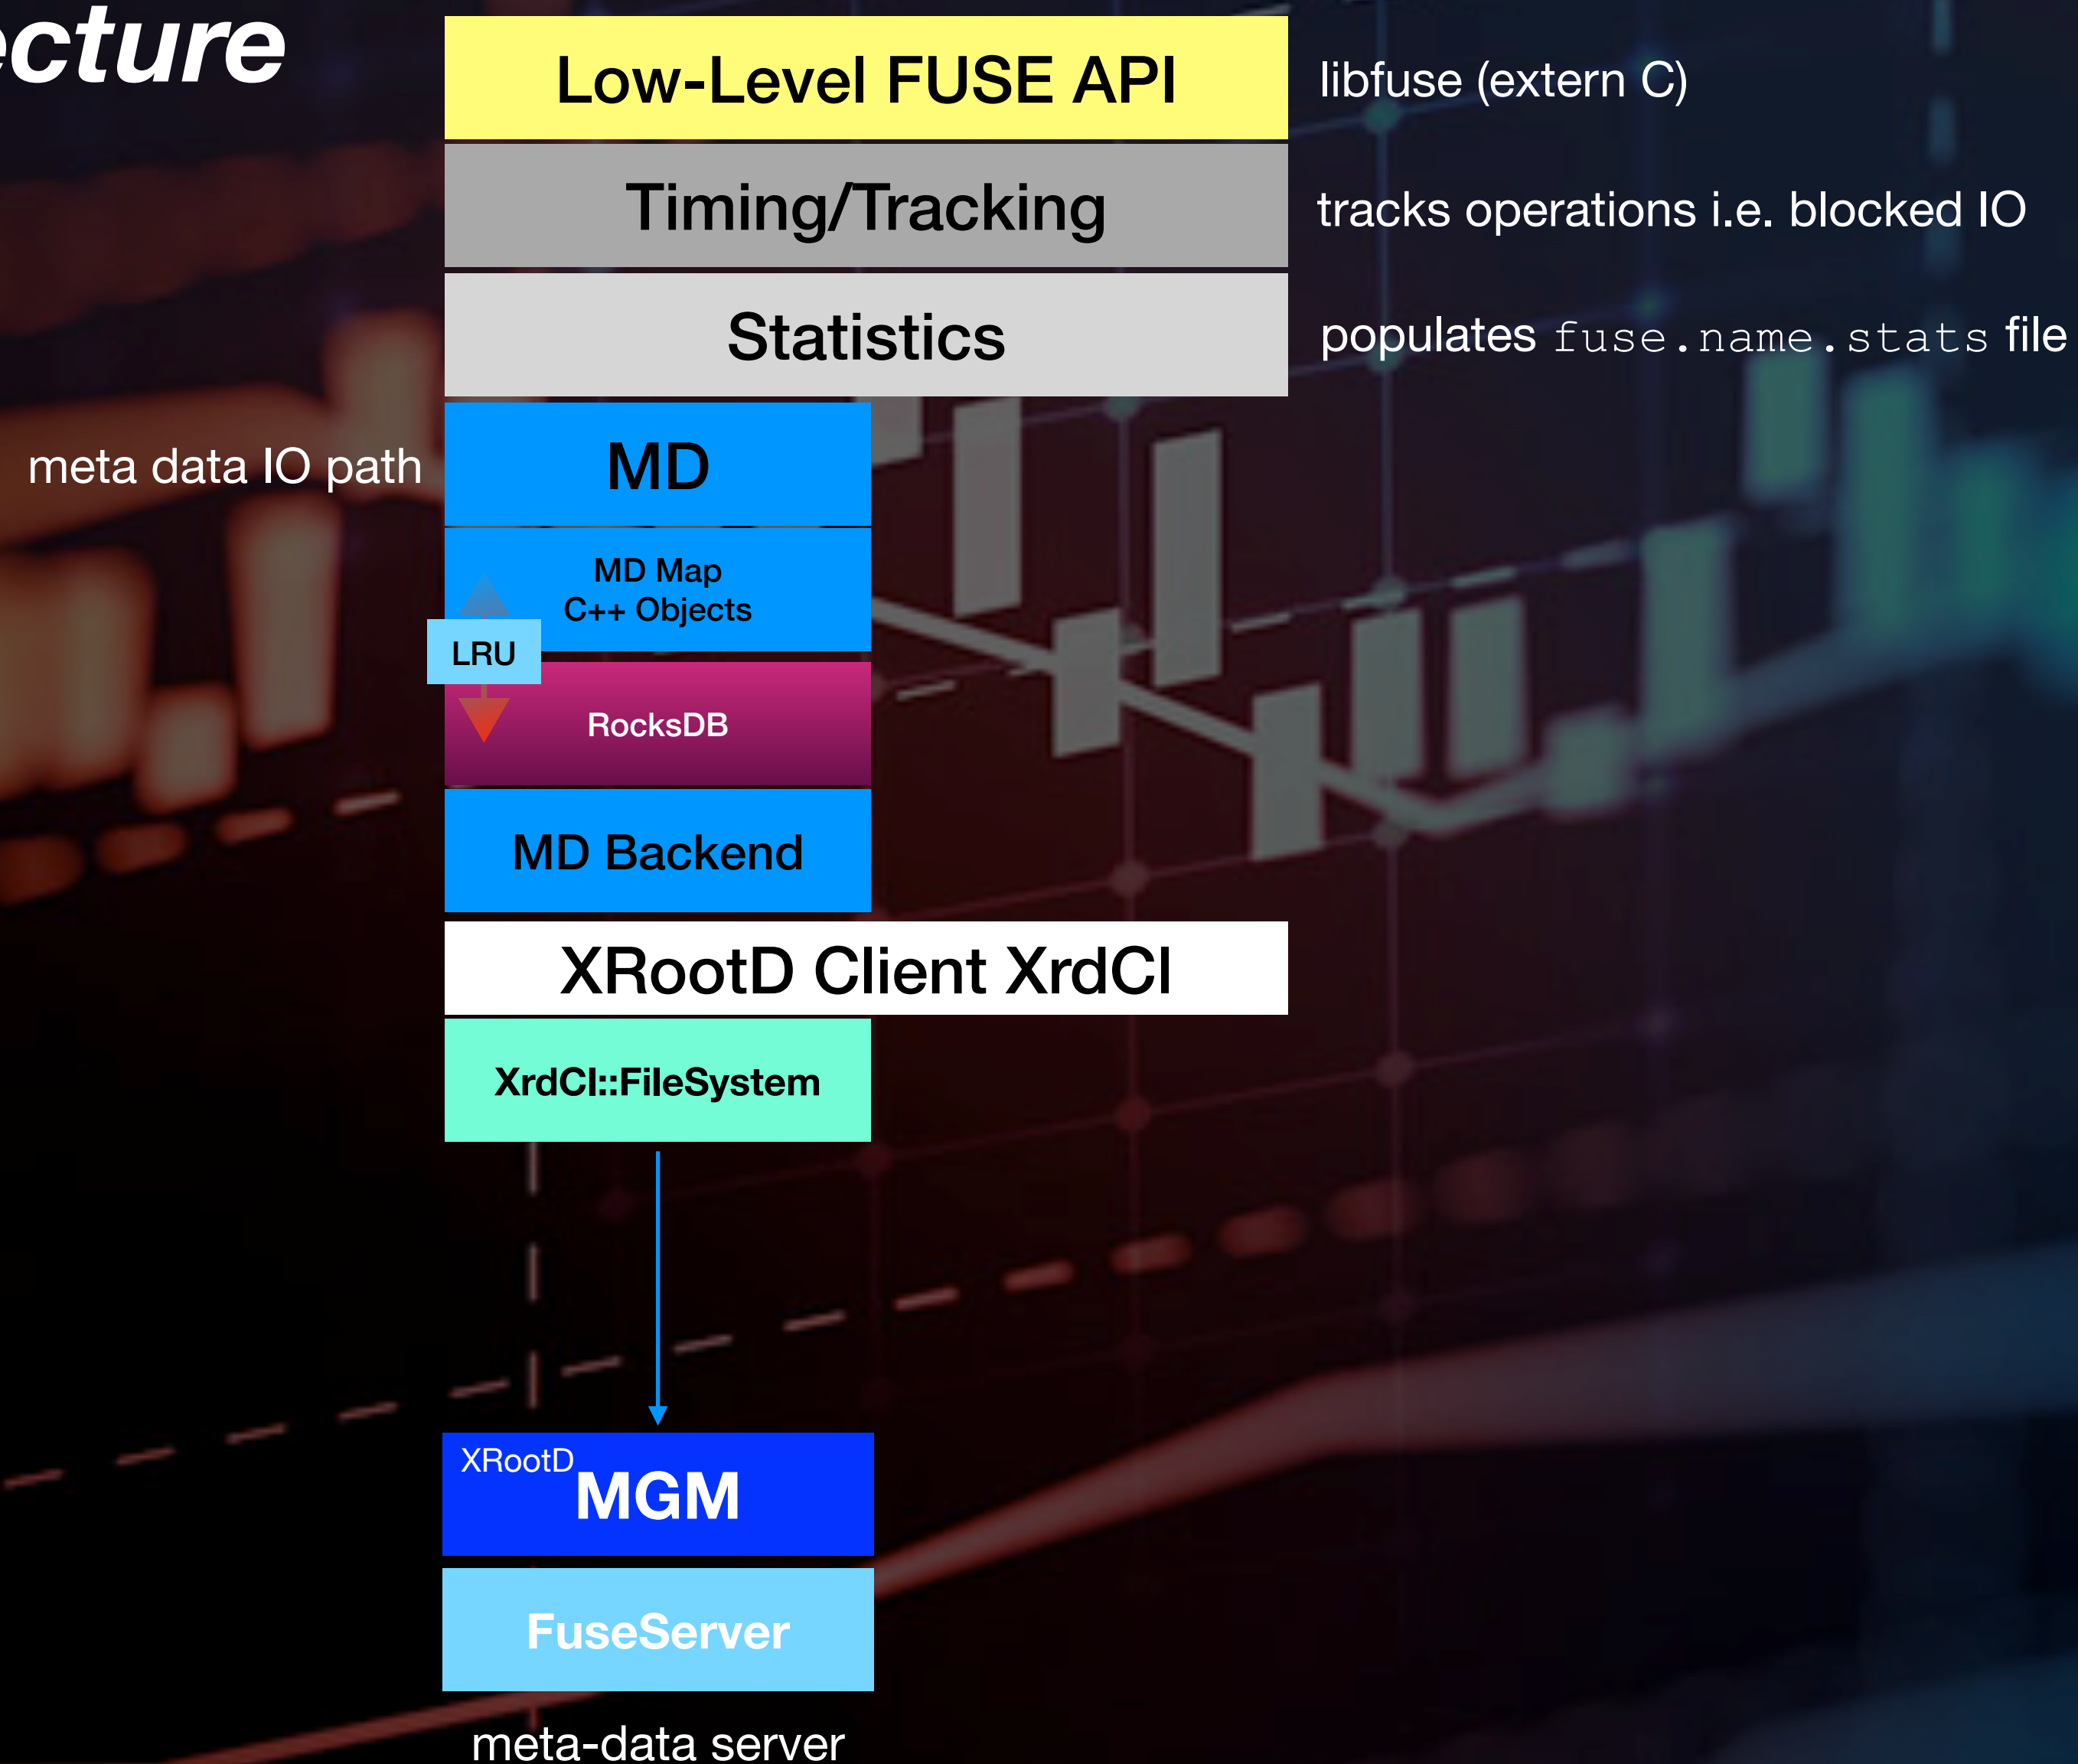
eosxd *architecture*

meta data IO path



How EOS FUSE works ...

eosxd architecture



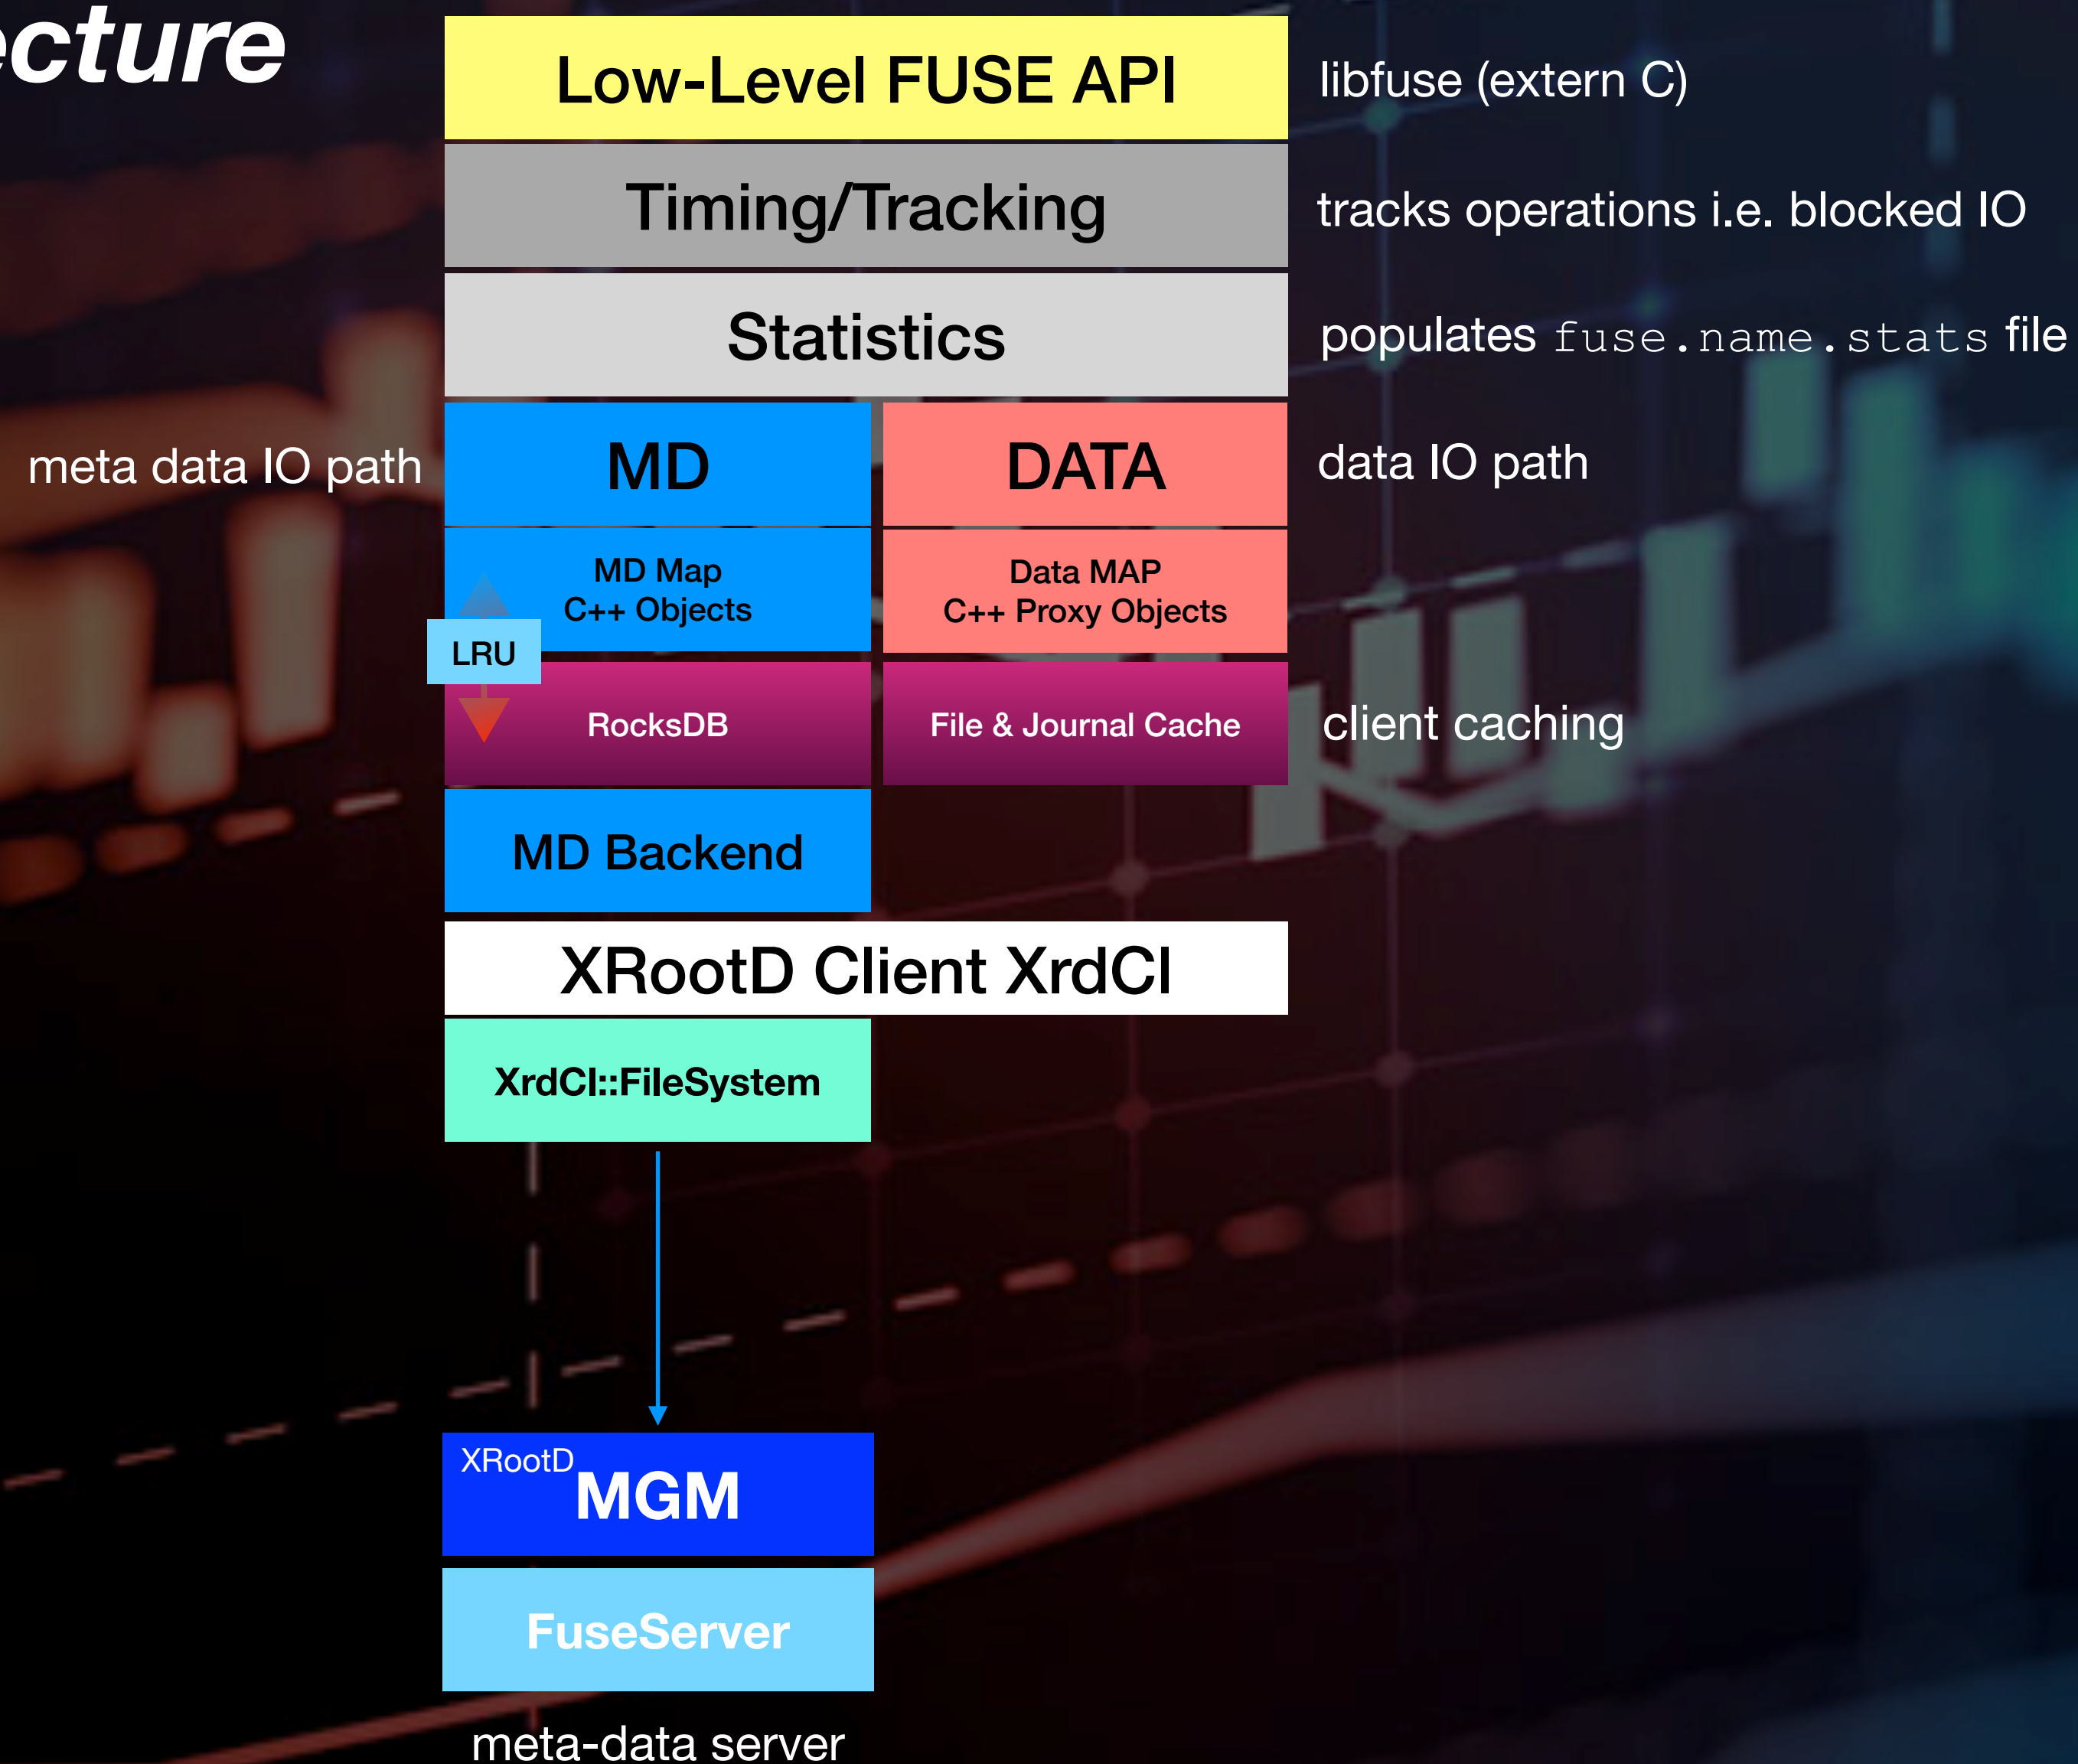
The diagram illustrates the architecture of the XRootD Client XrdCl, showing the flow of data and control between various components. The components are organized into layers:

- Low-Level FUSE API** (Yellow box): The top layer, interfacing with `libfuse` (extern C).
- Timing/Tracking** (Grey box): A layer for tracking operations, specifically blocked IO.
- Statistics** (Light grey box): A layer that populates the `fuse.name.stats` file.
- MD (Metadata) and DATA** (Blue and Red boxes): The core data path. The MD layer includes an **MD Map** (C++ Objects) and an **MD Backend**. The DATA layer includes a **Data MAP** (C++ Proxy Objects).
- LRU (Least Recently Used)** (Blue box): A cache mechanism for the MD Map.
- RocksDB** (Red box): A database component used for storing metadata.
- XRootD Client XrdCl** (White box): The main client interface.
- XrdCl::FileSystem** (Green box): The client's view of the file system.
- XRootD MGM (Meta-Data Manager)** (Blue box): The server-side component that manages metadata.
- FuseServer** (Light blue box): The server-side component that implements the FUSE interface.

The flow of data and control is indicated by arrows: a blue arrow points from the **MD Map** to the **LRU** cache, and a red arrow points from the **LRU** cache to the **RocksDB** database. A green arrow points from the **XrdCl::FileSystem** to the **XRootD MGM**, and a blue arrow points from the **XRootD MGM** to the **FuseServer**.

How EOS FUSE works ...

eosxd architecture



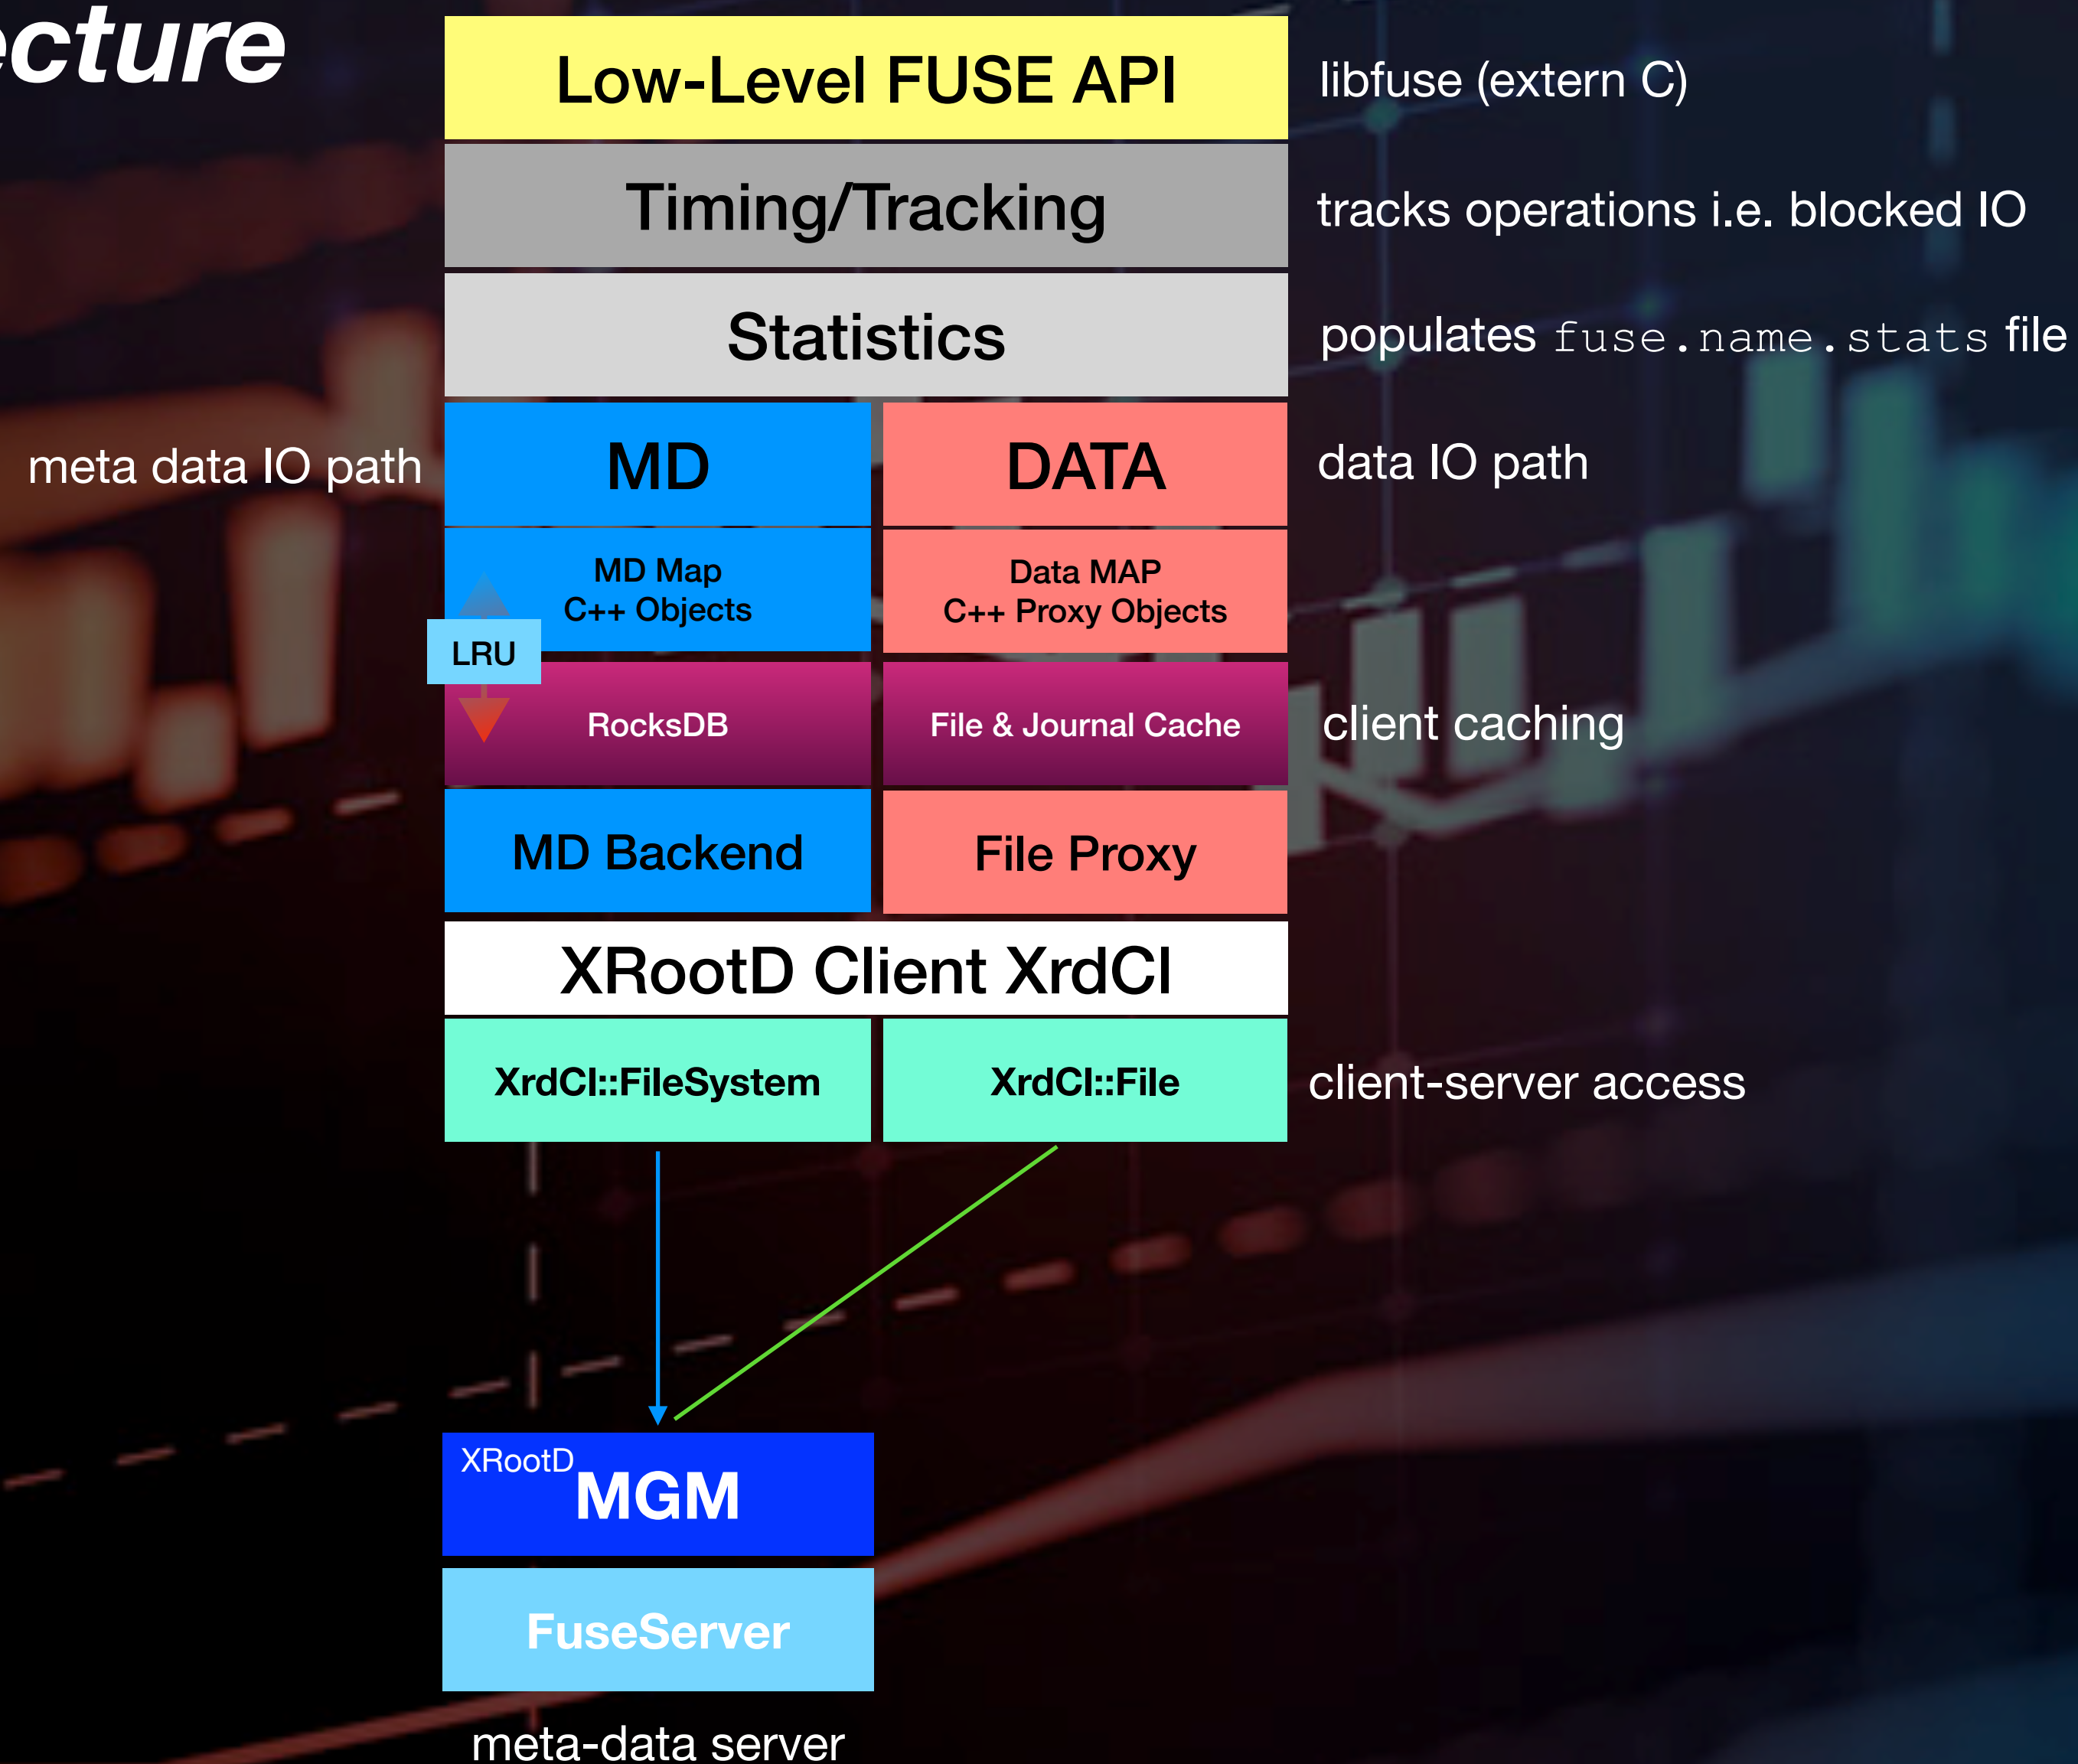
The diagram illustrates the architecture of the XRootD Client XrdCl, organized into several layers and components:

- Low-Level FUSE API** (Yellow box): The top layer, interfacing with `libfuse` (extern C).
- Timing/Tracking** (Grey box): A layer that tracks operations, specifically blocked IO.
- Statistics** (Light Grey box): A layer that populates the `fuse.name.stats` file.
- MD (Metadata) and DATA** (Blue and Red boxes): The core data IO path, divided into two main sections:
 - MD (Left Column):**
 - MD Map C++ Objects** (Blue box): Interacts with `LRU` (Least Recently Used) cache.
 - RocksDB** (Dark Blue box): The underlying storage for the MD Map.
 - MD Backend** (Blue box): The interface for the MD Map.
 - DATA (Right Column):**
 - Data MAP C++ Proxy Objects** (Red box): Interacts with `client caching`.
 - File & Journal Cache** (Dark Red box): The underlying storage for the Data MAP.
 - File Proxy** (Red box): The interface for the Data MAP.
- XRootD Client XrdCl** (White box): The main client interface, which includes the `XrdCl::FileSystem` component.
- XRootD MGM (Meta-Data Manager)** (Blue box): The central component that manages the meta-data, receiving input from the `XrdCl::FileSystem` and the `FuseServer`.
- FuseServer** (Light Blue box): The server component that handles the meta-data requests, acting as the **meta-data server**.

Arrows indicate the flow of data and control between these components, showing a clear separation of concerns and a modular design.

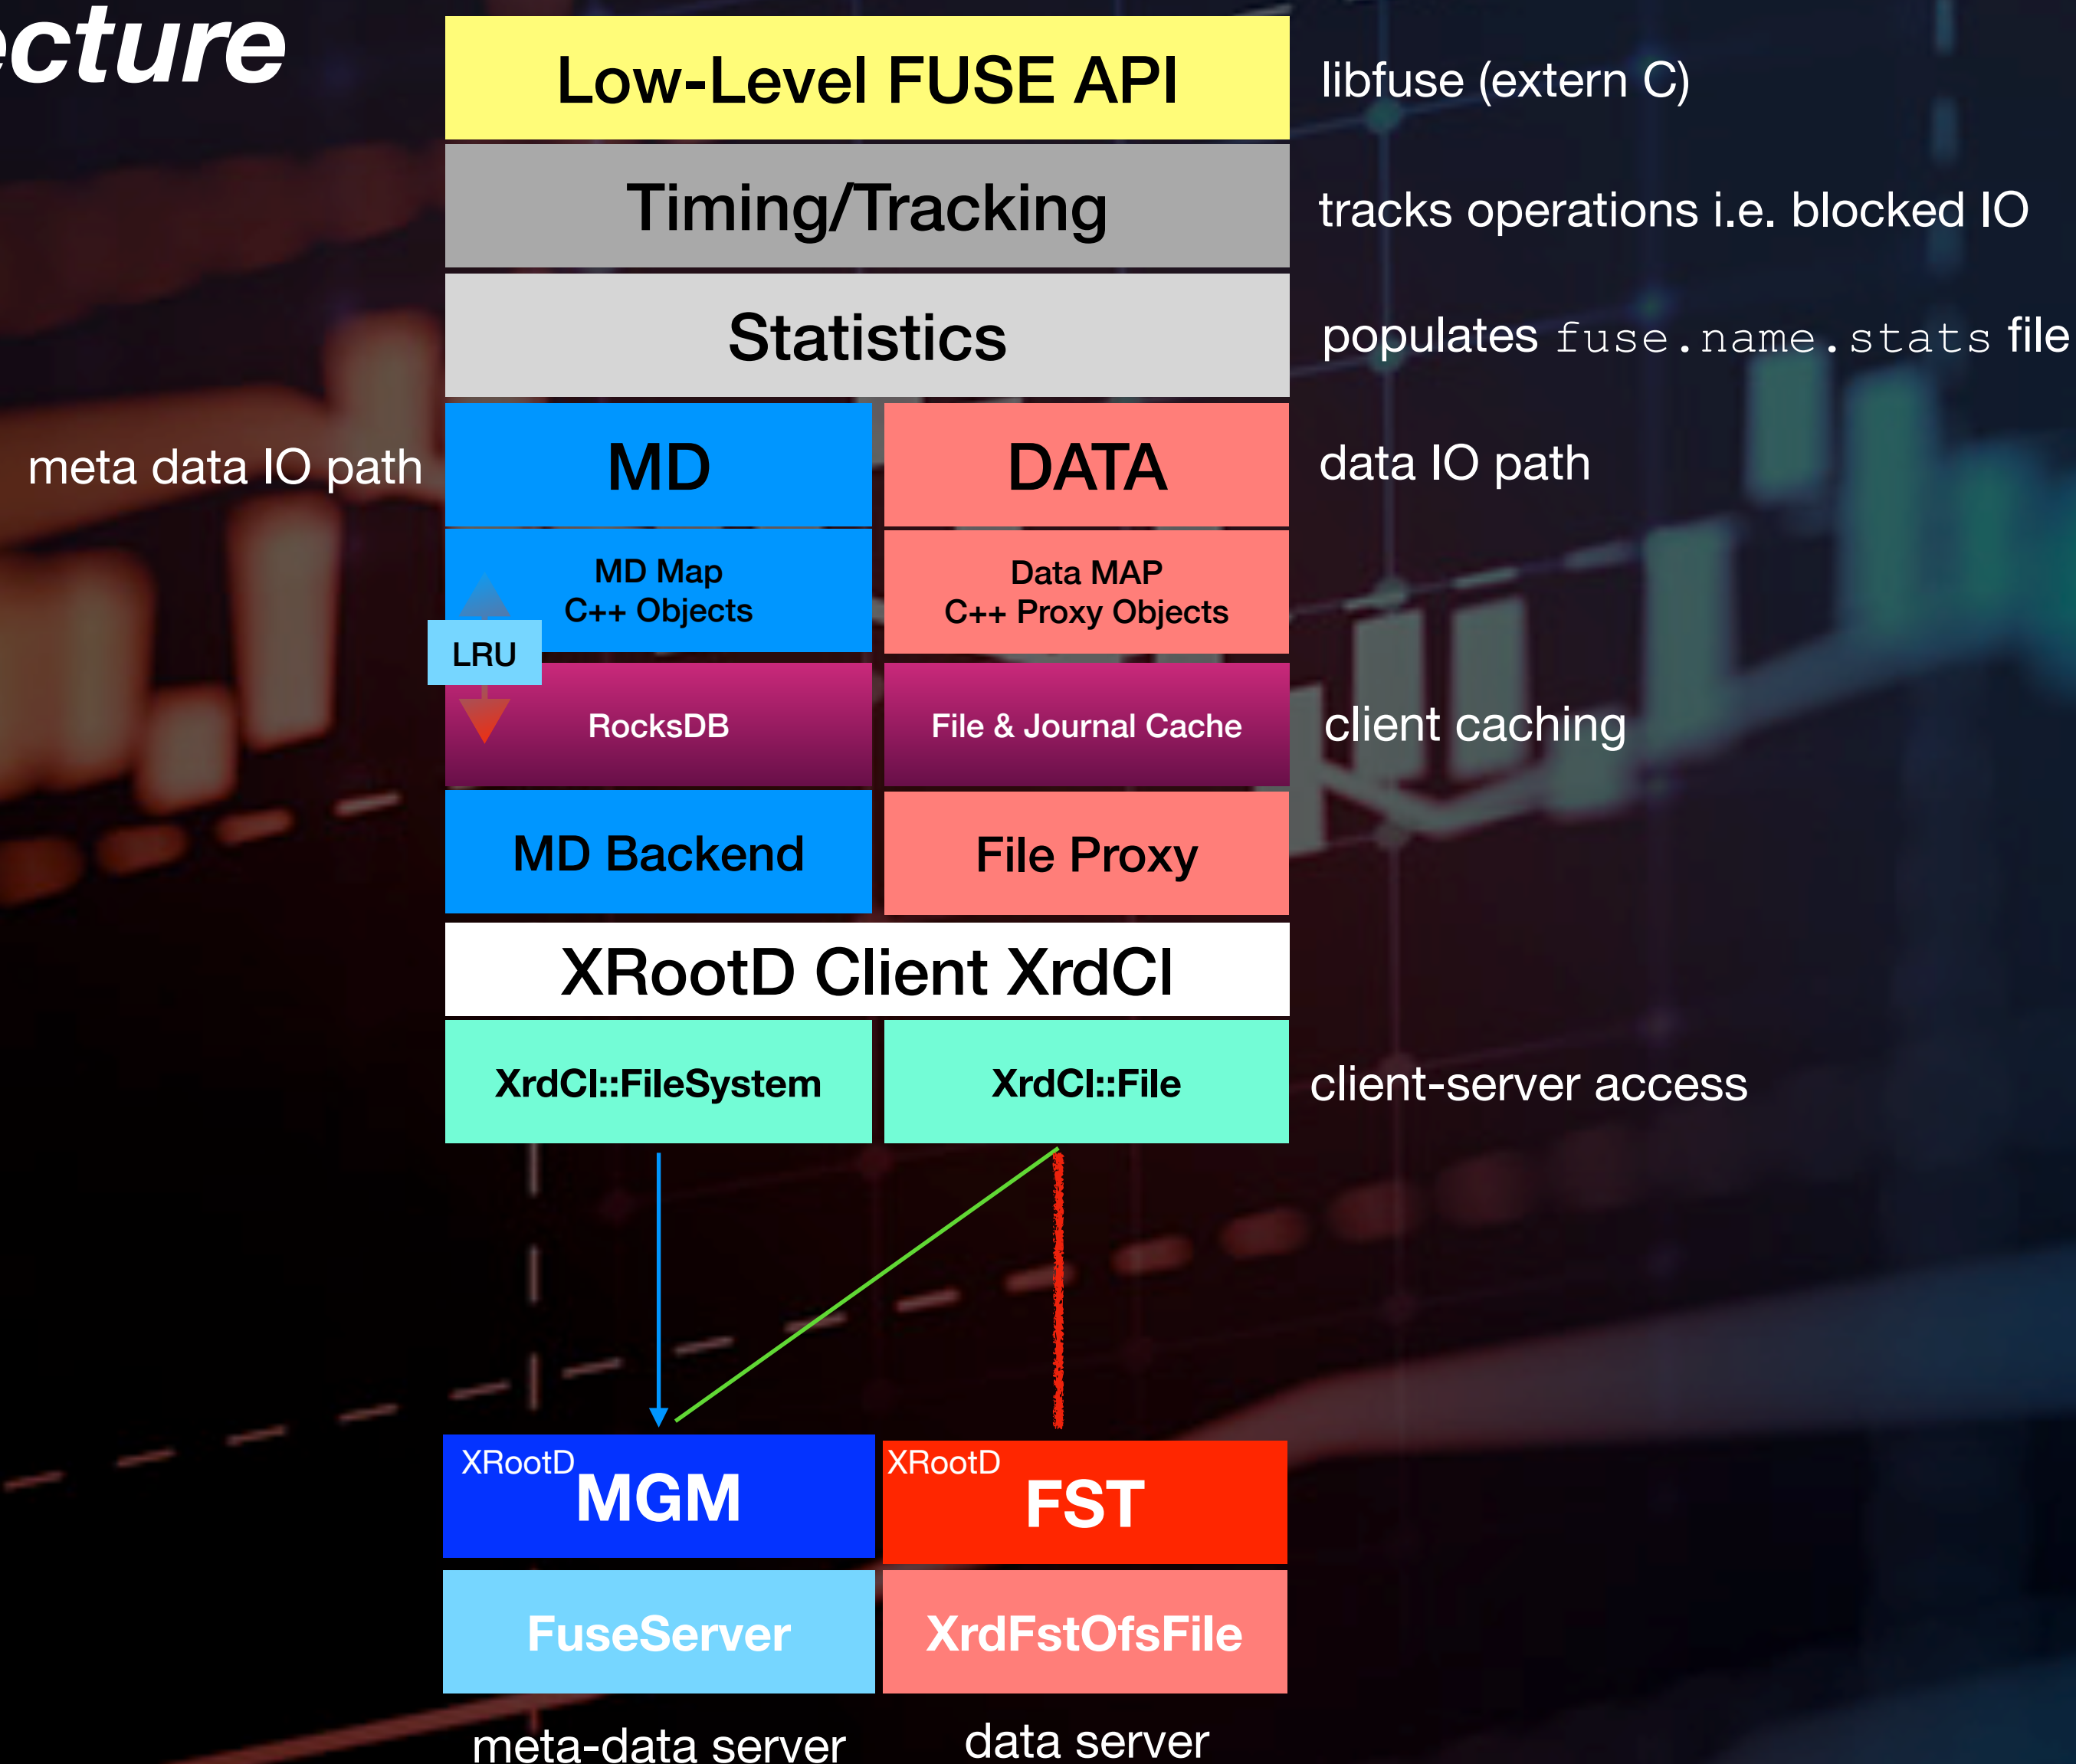
How EOS FUSE works ...

eosxd architecture



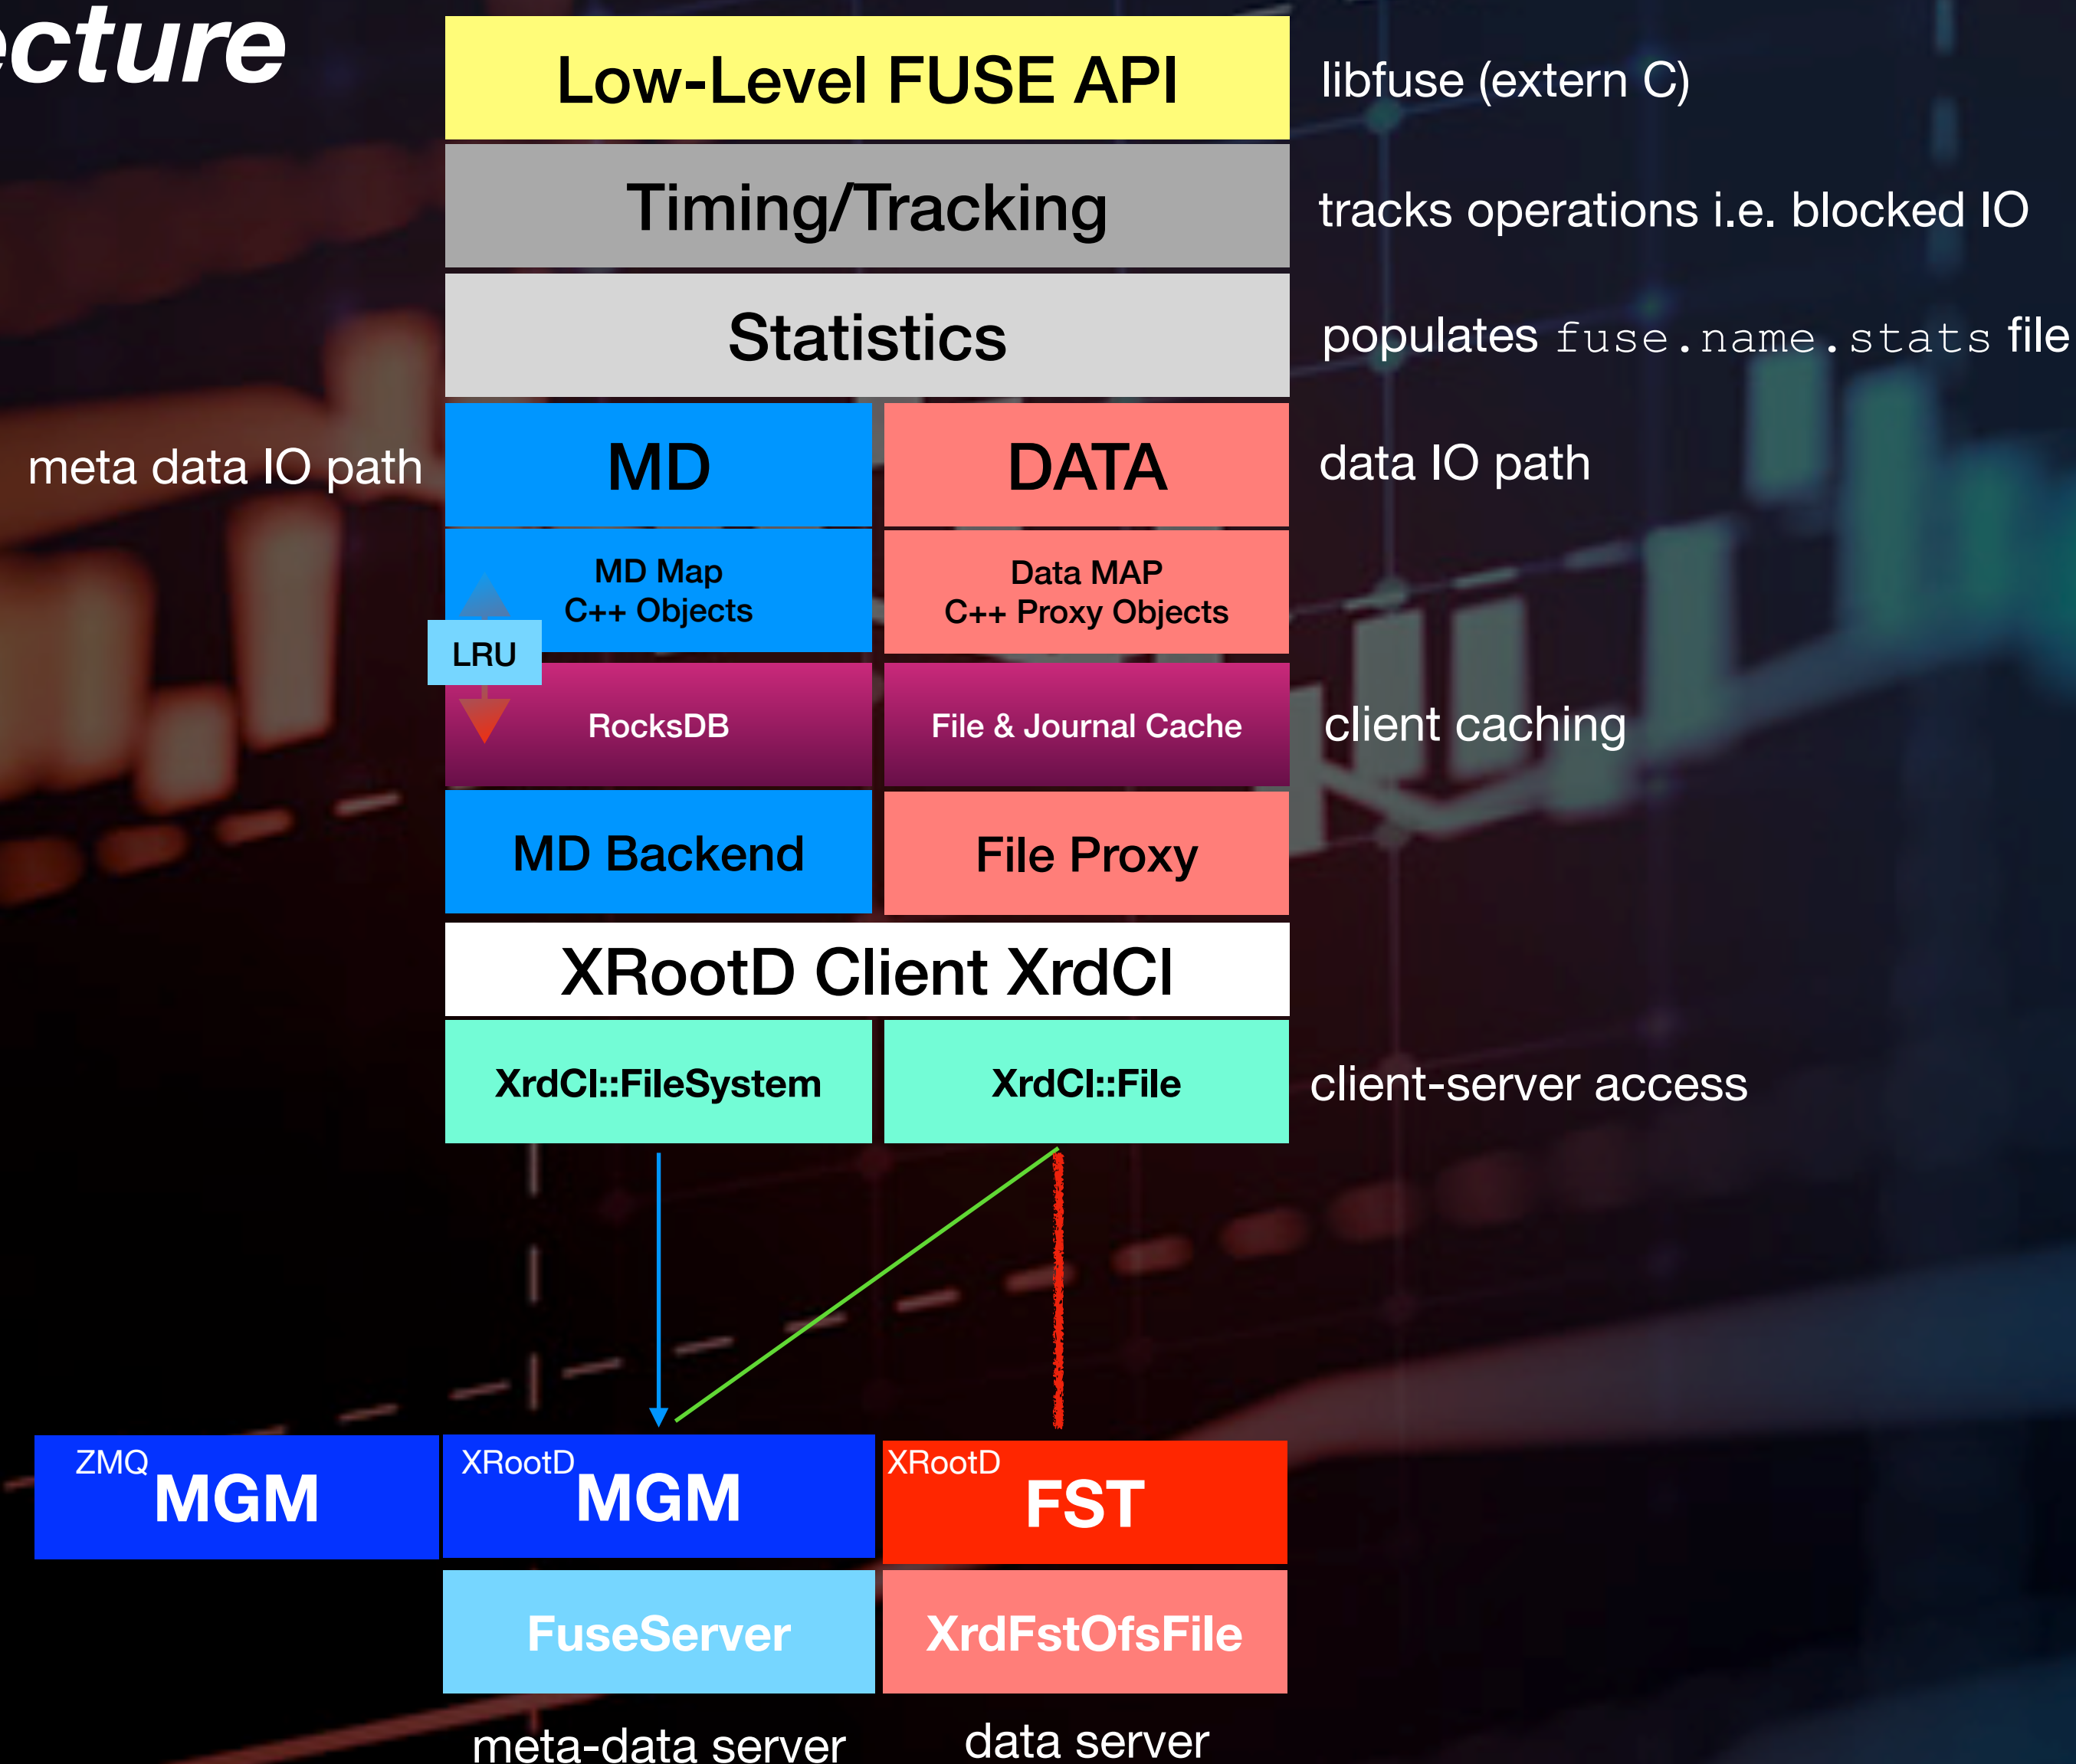
How EOS FUSE works ...

eosxd architecture



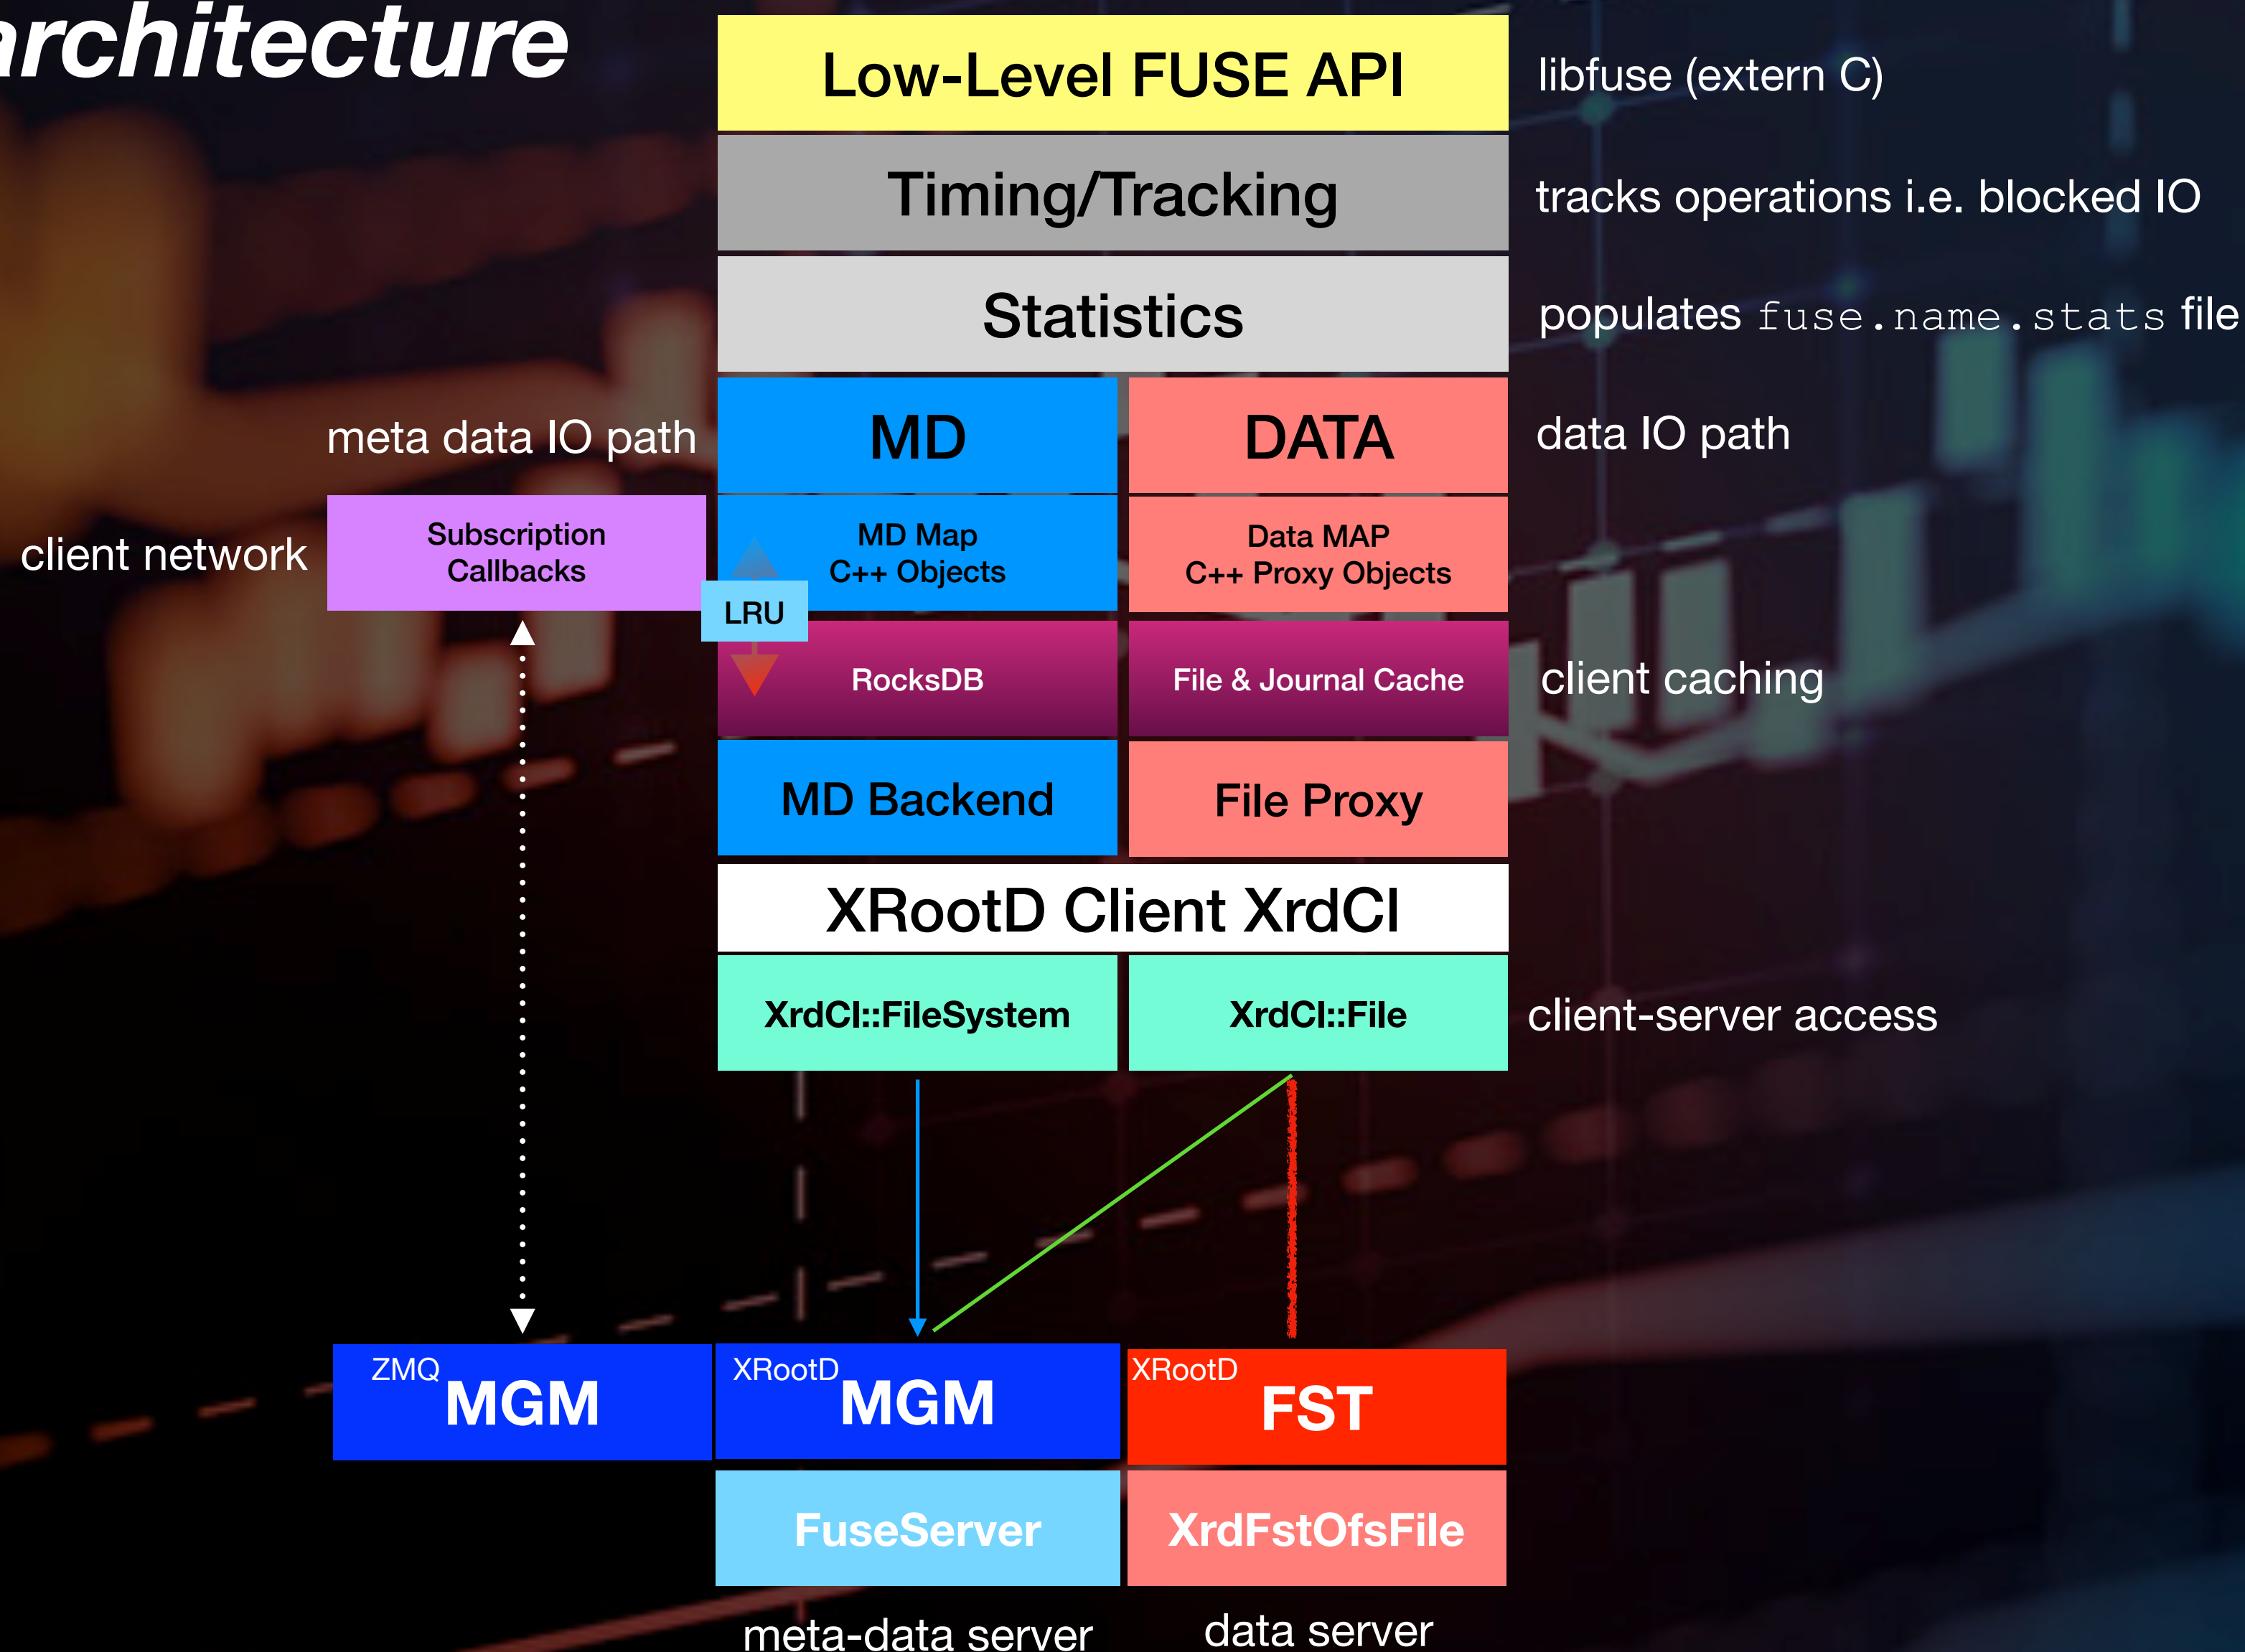
How EOS FUSE works ...

eosxd architecture



How EOS FUSE works ...

eosxd architecture



eosxd

Namespace FuseServer



eosxd

Namespace FuseServer

- The Fuse Server server-side CLI is used to ...

eosxd

Namespace FuseServer

- The Fuse Server server-side CLI is used to ...
 - show connected clients

eosxd

Namespace FuseServer

- The Fuse Server server-side CLI is used to ...
 - show connected clients
 - and most important state information (stuck IO etc.)

eosxd

Namespace FuseServer

- The Fuse Server server-side CLI is used to ...
 - show connected clients
 - and most important state information (stuck IO etc.)
 - evict unwanted clients

eosxd

Namespace FuseServer

- **The Fuse Server server-side CLI is used to ...**
 - **show connected clients**
 - and most important state information (stuck IO etc.)
 - **evict unwanted clients**
 - **show notification subscriptions and permission delegation (capabilities)**

eosxd

Namespace FuseServer

- **The Fuse Server server-side CLI is used to ...**
 - **show connected clients**
 - and most important state information (stuck IO etc.)
 - **evict unwanted clients**
 - **show notification subscriptions and permission delegation (capabilities)**
 - **drop state information**

eosxd

Namespace FuseServer

- **The Fuse Server server-side CLI is used to ...**
 - **show connected clients**
 - and most important state information (stuck IO etc.)
 - **evict unwanted clients**
 - **show notification subscriptions and permission delegation (capabilities)**
 - **drop state information**
 - capabilities (directory subscriptions)

eosxd

Namespace FuseServer

- **The Fuse Server server-side CLI is used to ...**
 - **show connected clients**
 - and most important state information (stuck IO etc.)
 - **evict unwanted clients**
 - **show notification subscriptions and permission delegation (capabilities)**
 - **drop state information**
 - capabilities (directory subscriptions)
 - locks (delegated locks given to a client)

eosxd

Namespace FuseServer

- **The Fuse Server server-side CLI is used to ...**
 - **show connected clients**
 - and most important state information (stuck IO etc.)
 - **evict unwanted clients**
 - **show notification subscriptions and permission delegation (capabilities)**
 - **drop state information**
 - capabilities (directory subscriptions)
 - locks (delegated locks given to a client)
 - **configure heartbeat and broadcast configuration**

eosxd

Namespace FuseServer

- **The Fuse Server server-side CLI is used to ...**
 - **show connected clients**
 - and most important state information (stuck IO etc.)
 - **evict unwanted clients**
 - **show notification subscriptions and permission delegation (capabilities)**
 - **drop state information**
 - capabilities (directory subscriptions)
 - locks (delegated locks given to a client)
 - **configure heartbeat and broadcast configuration**
 - how often each client has to send a heartbeat

eosxd

Namespace FuseServer

- **The Fuse Server server-side CLI is used to ...**
 - **show connected clients**
 - and most important state information (stuck IO etc.)
 - **evict unwanted clients**
 - **show notification subscriptions and permission delegation (capabilities)**
 - **drop state information**
 - capabilities (directory subscriptions)
 - locks (delegated locks given to a client)
 - **configure heartbeat and broadcast configuration**
 - how often each client has to send a heartbeat
 - audience suppression to avoid too high callback rates

eosxd

Broadcasts & Callbacks



eosxd

Broadcasts & Callbacks

external broadcasts

an external Client (not FUSE) triggers a broadcast of file meta-data


extC¹

eosxd

Broadcasts & Callbacks

external broadcasts

an external Client (not FUSE) triggers a broadcast of file meta-data





eosxd

Broadcasts & Callbacks

external broadcasts

an external Client (not FUSE) triggers a broadcast of file meta-data

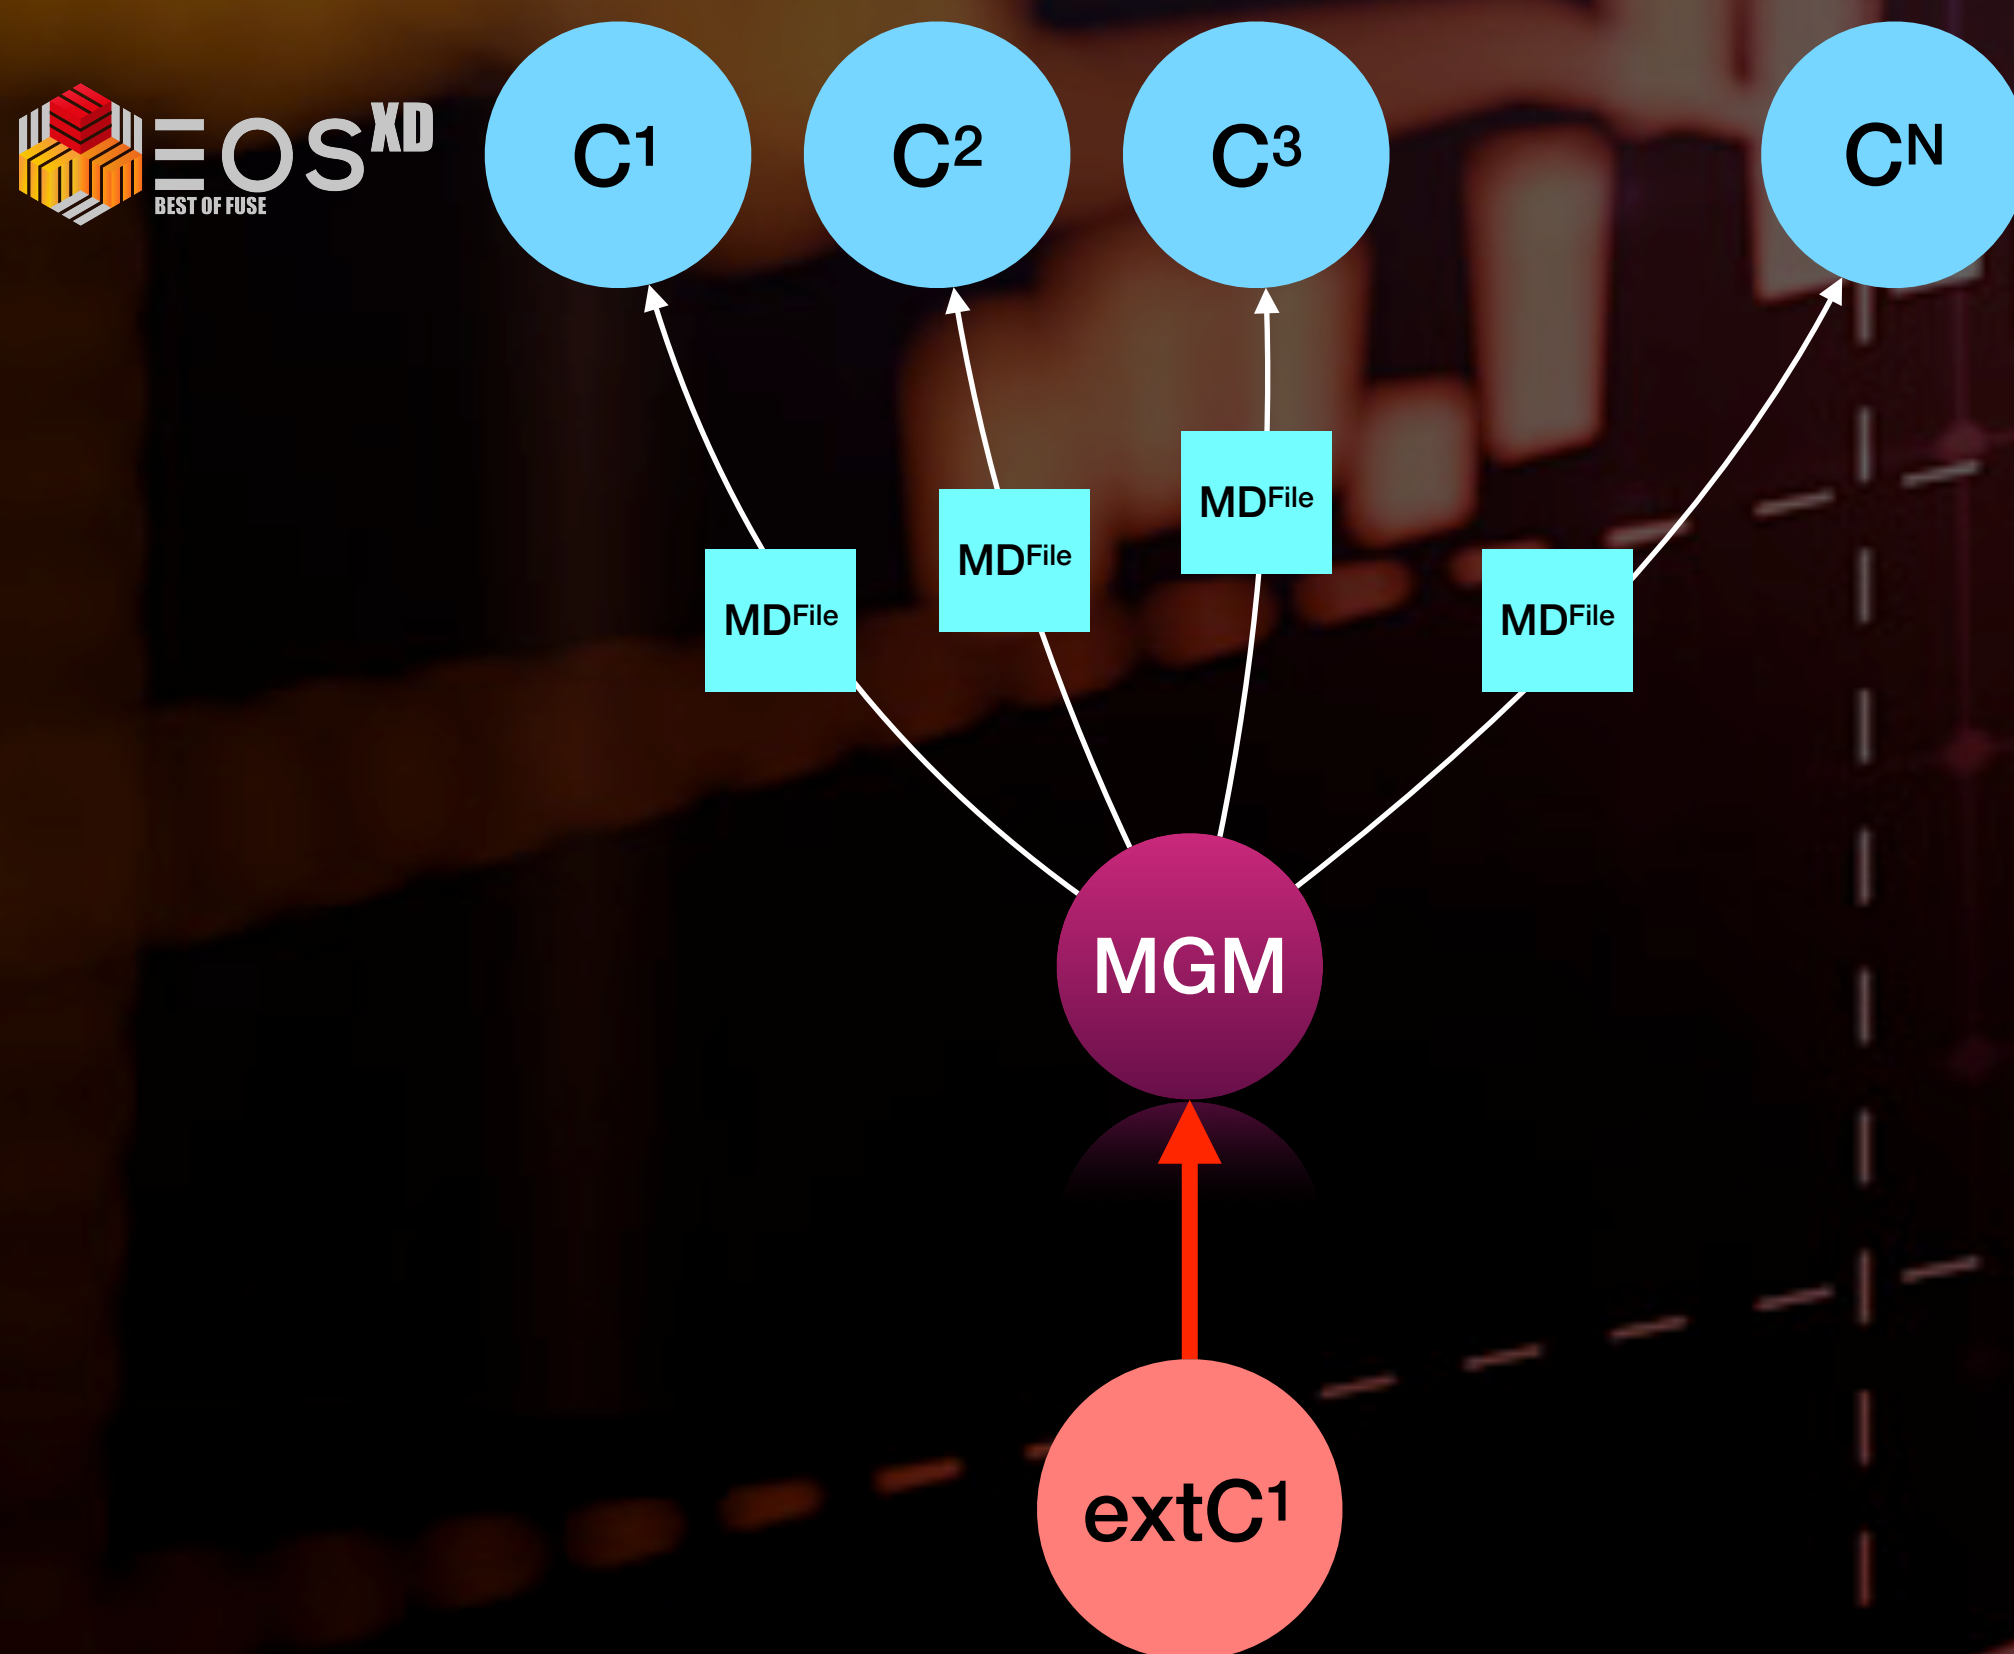


eosxd

Broadcasts & Callbacks

external broadcasts

an external Client (not FUSE) triggers a broadcast of file meta-data

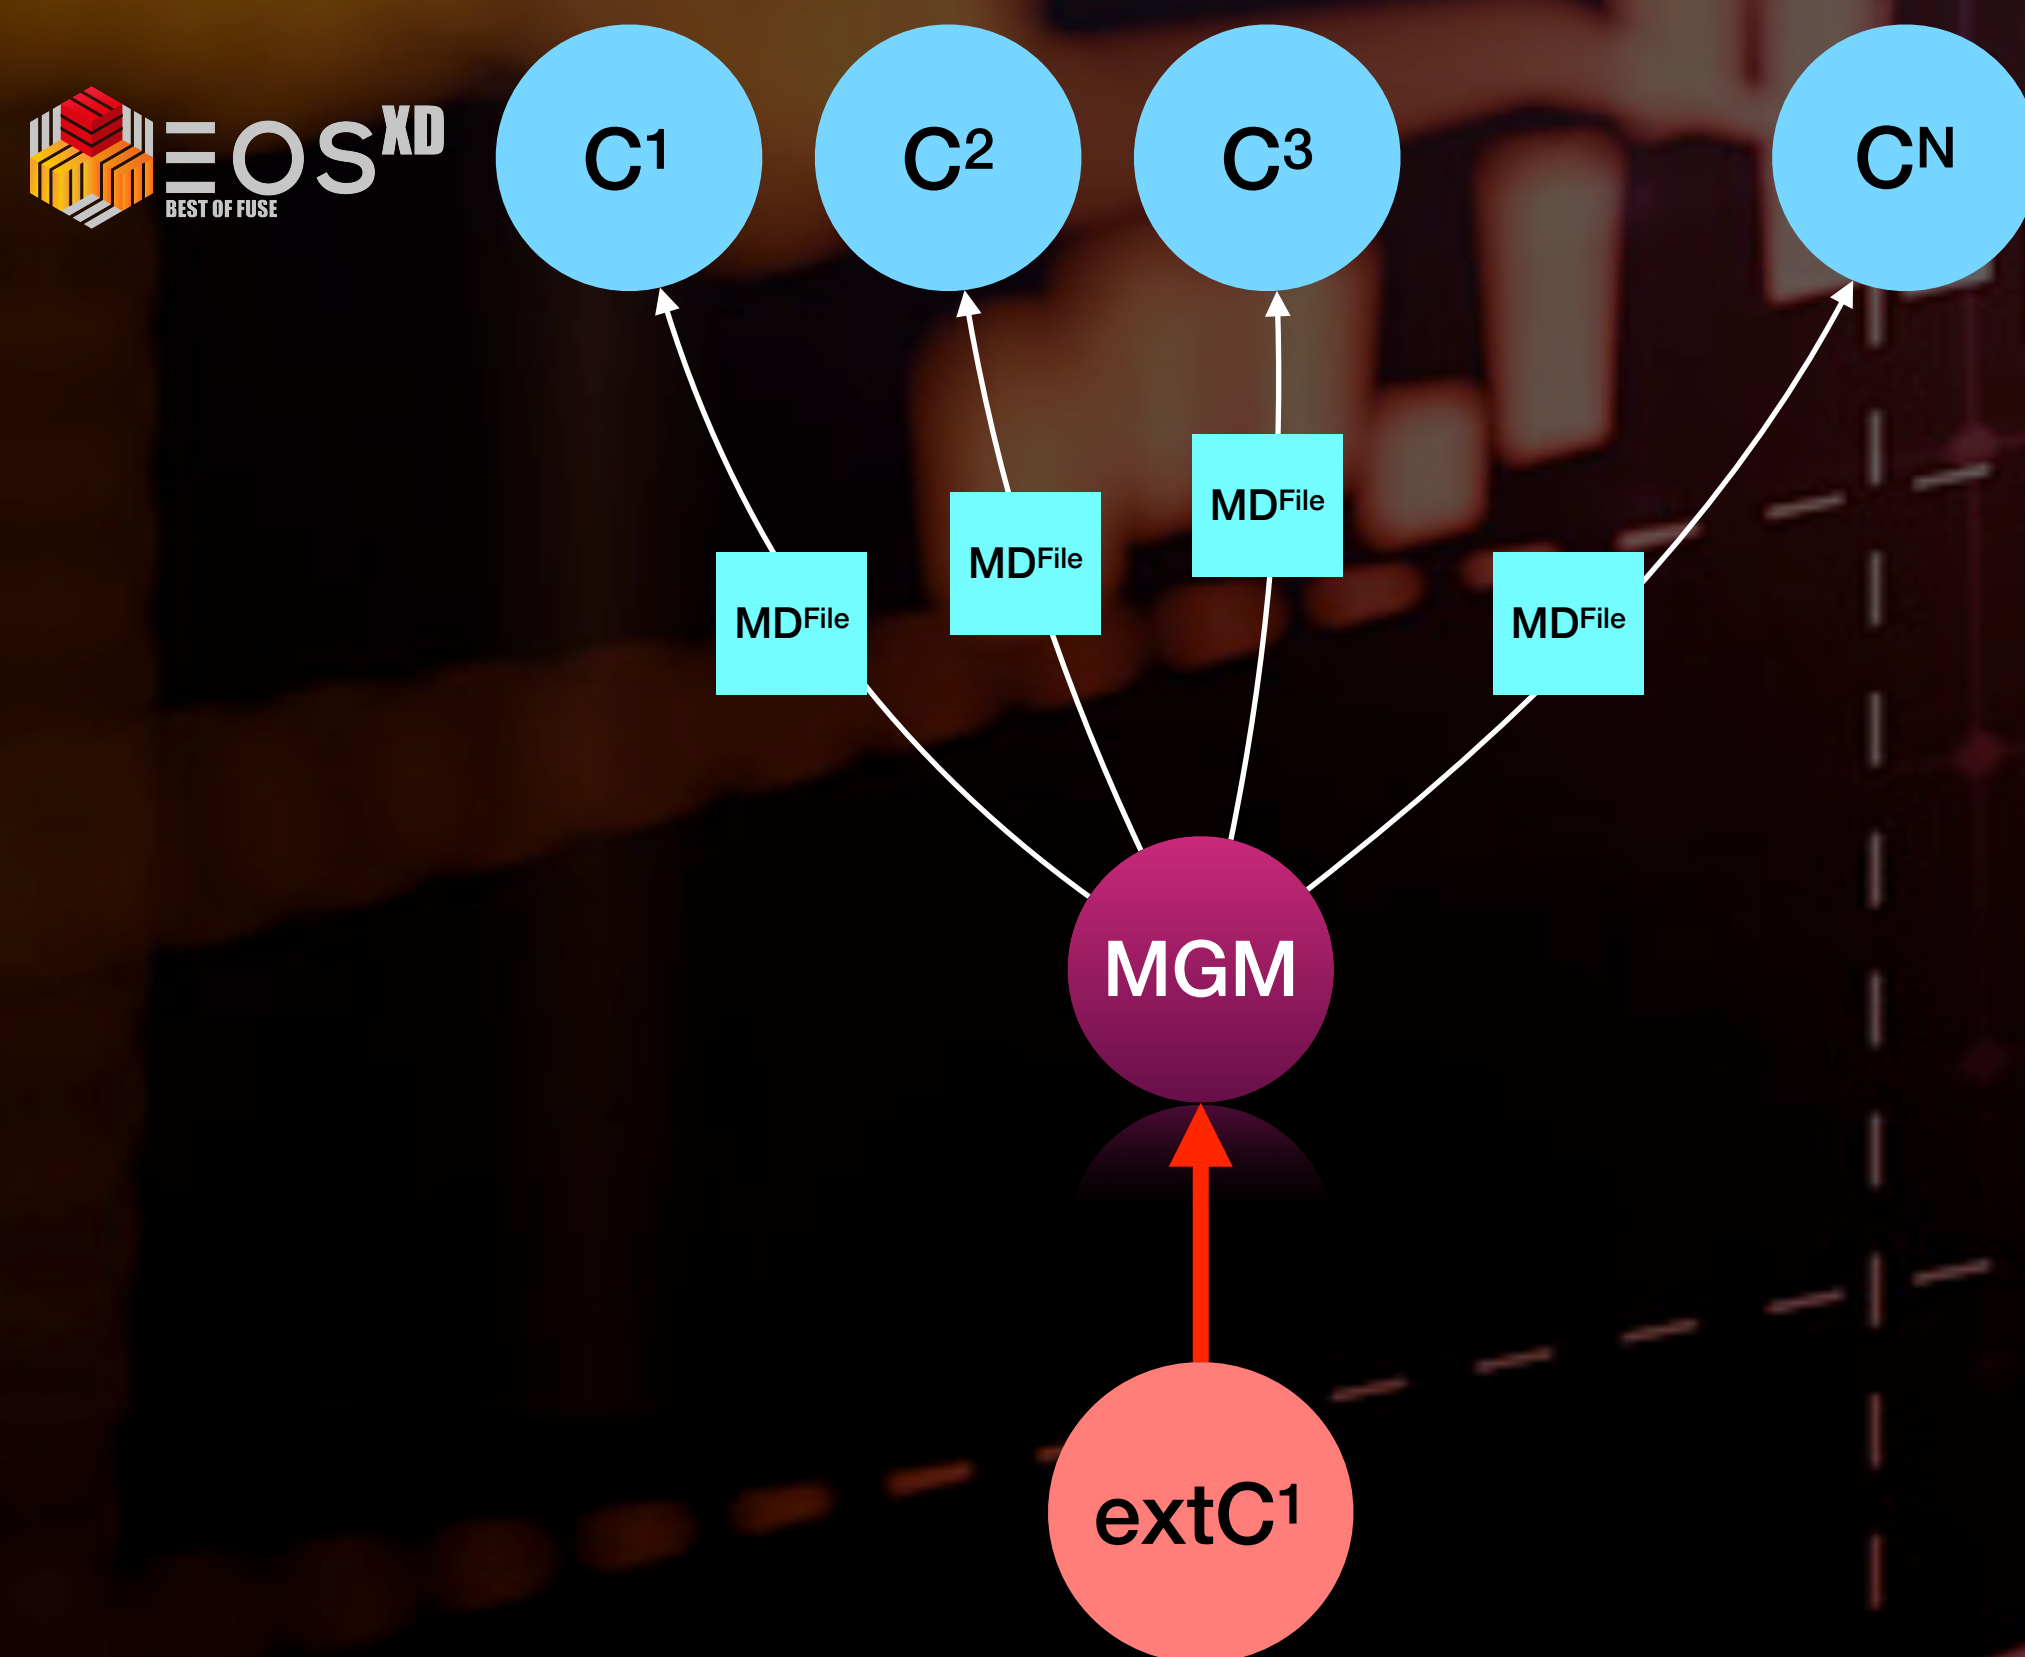


eosxd

Broadcasts & Callbacks


external broadcasts

an external Client (not FUSE) triggers a broadcast of file meta-data



internal broadcasts

an internal client triggers a broadcast of file meta-data

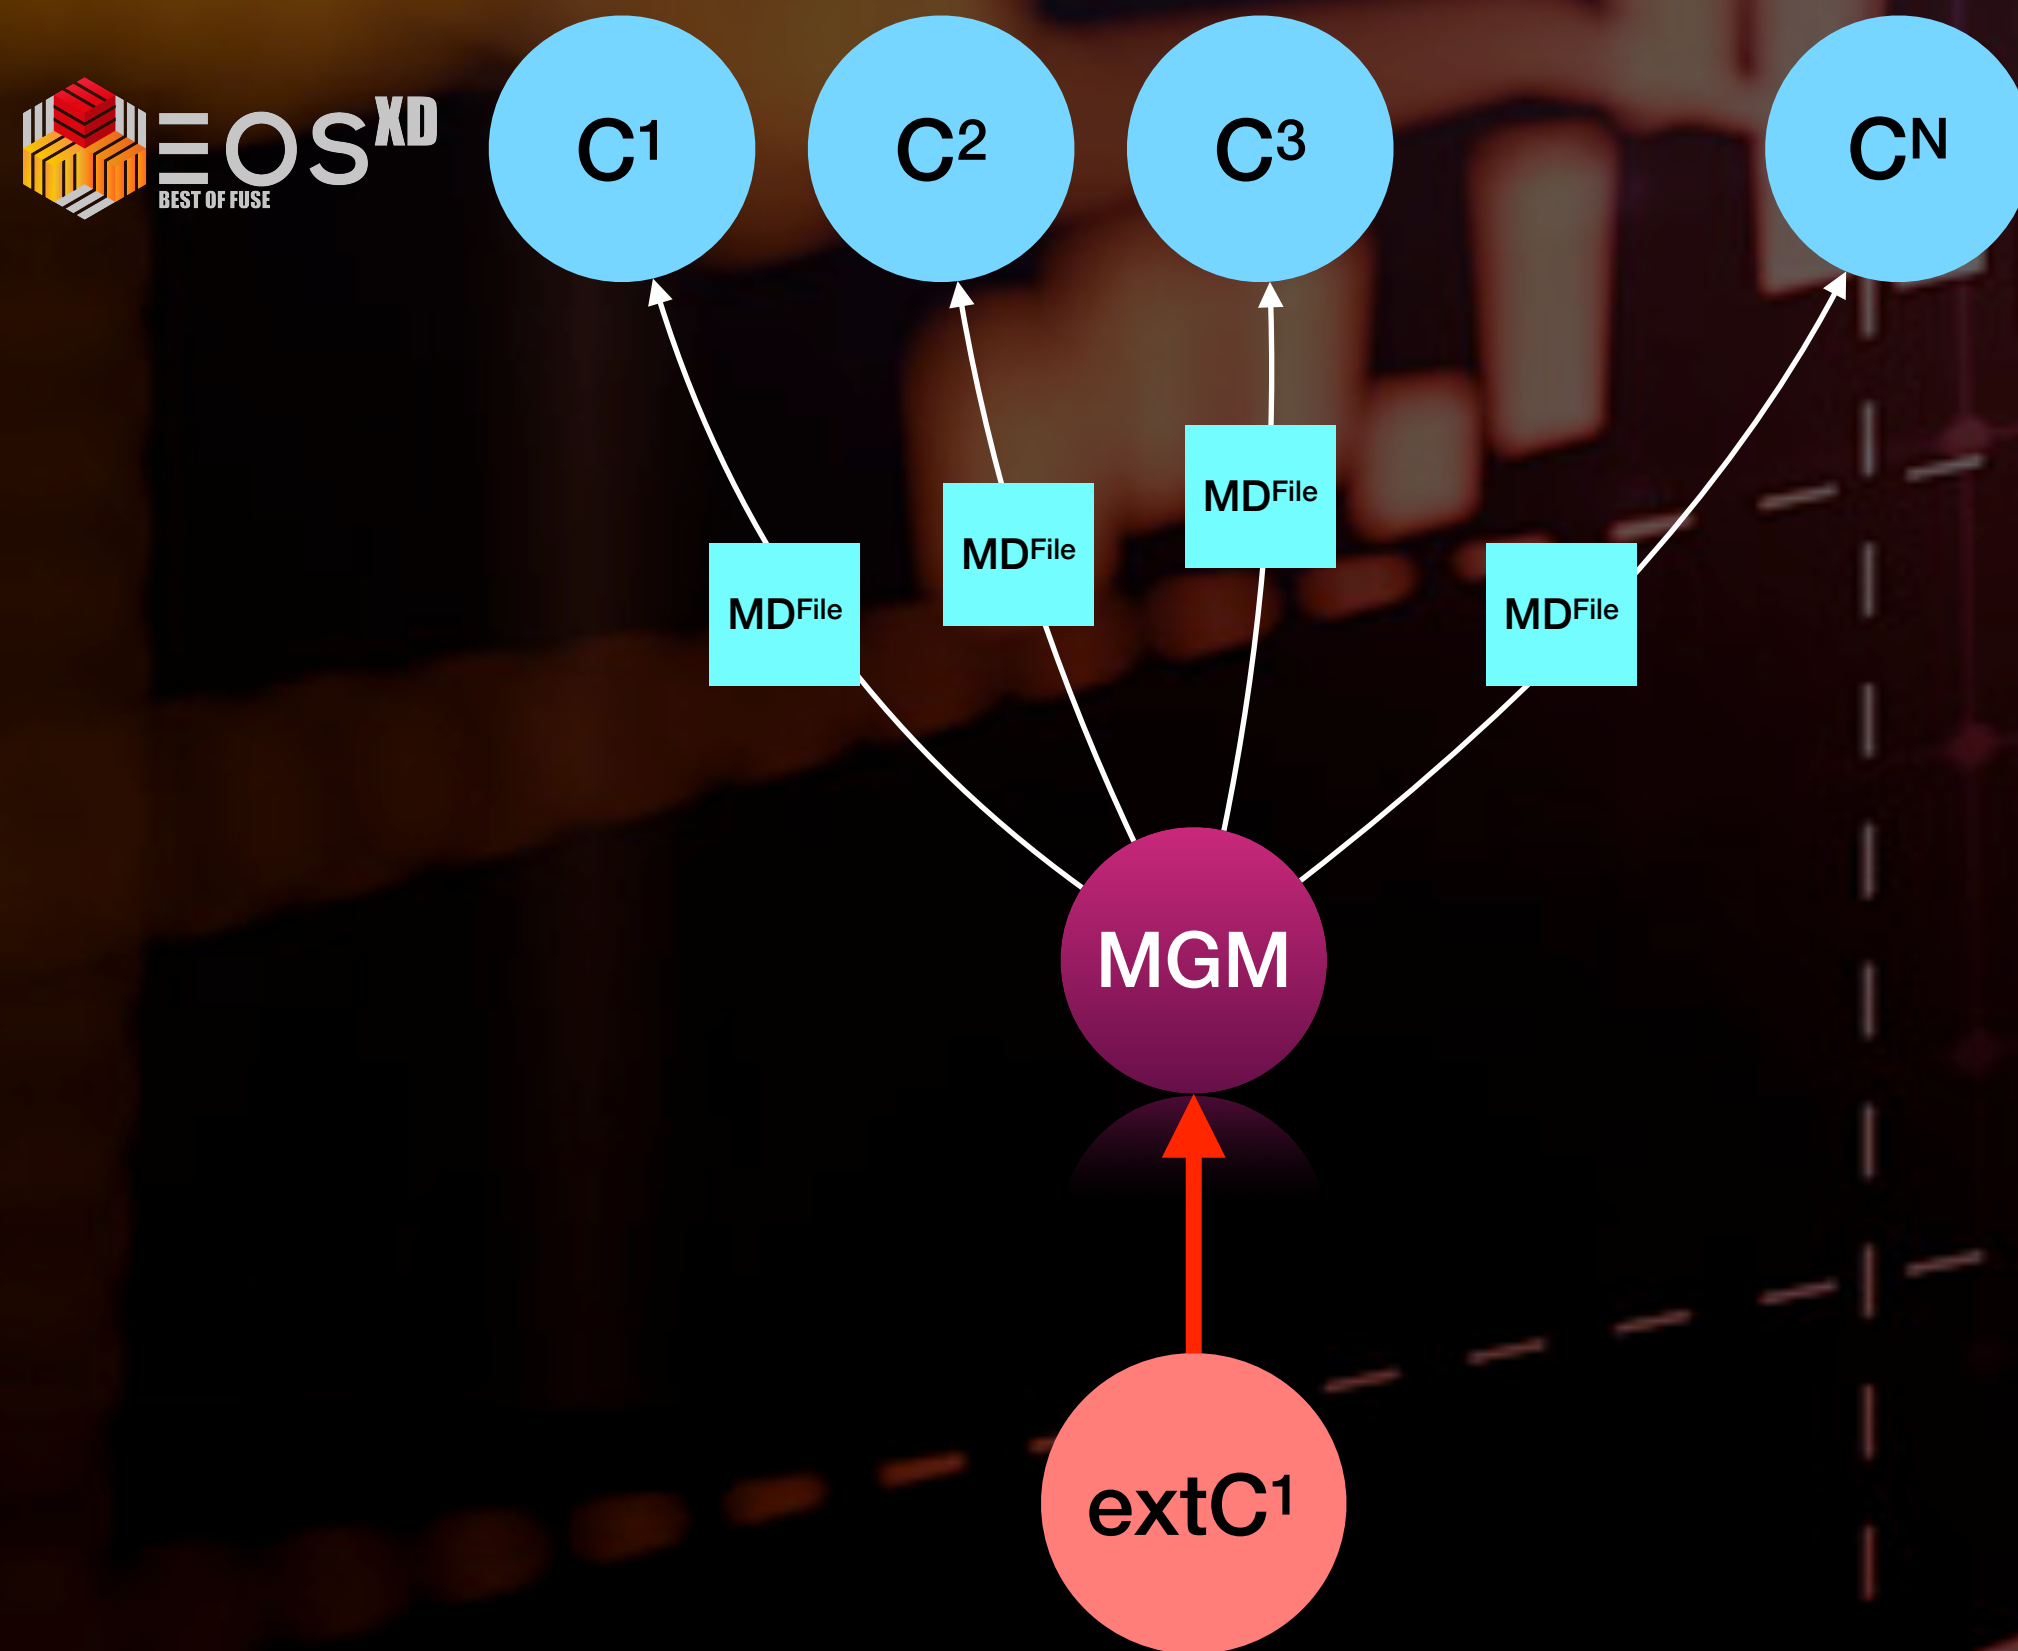


eosxd

Broadcasts & Callbacks

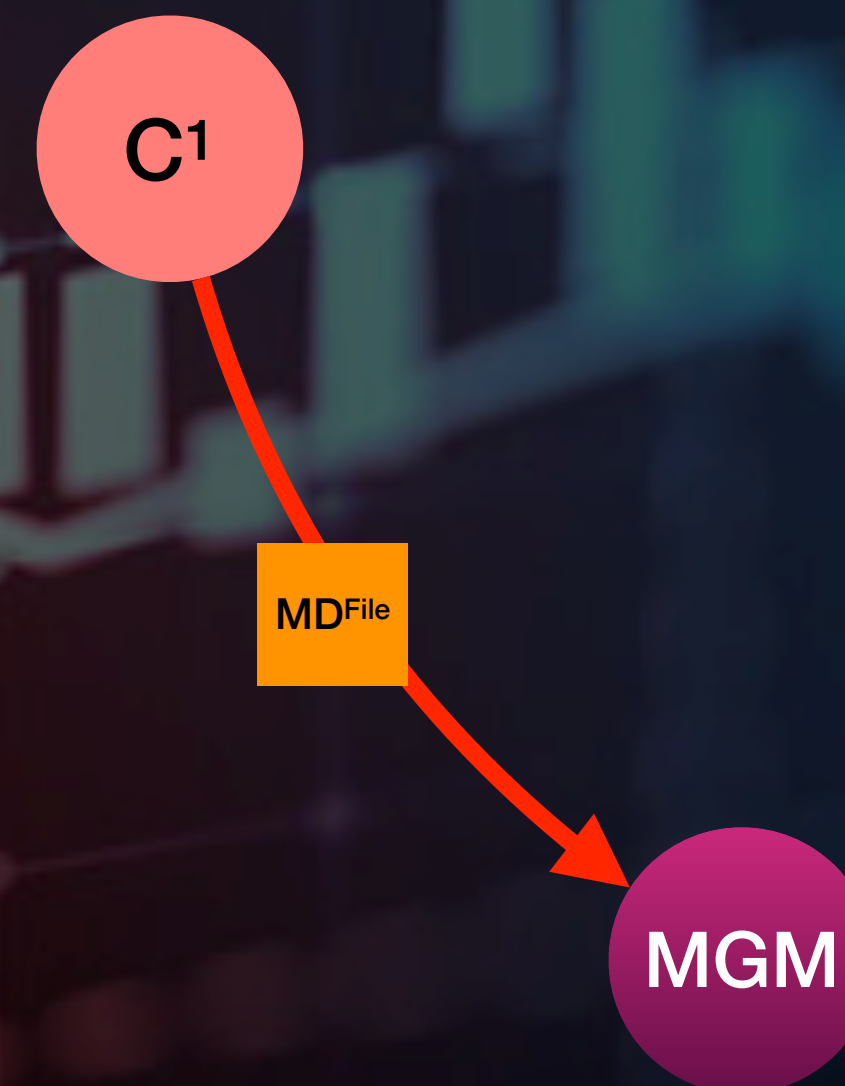
external broadcasts

an external Client (not FUSE) triggers a broadcast of file meta-data



internal broadcasts

an internal client triggers a broadcast of file meta-data

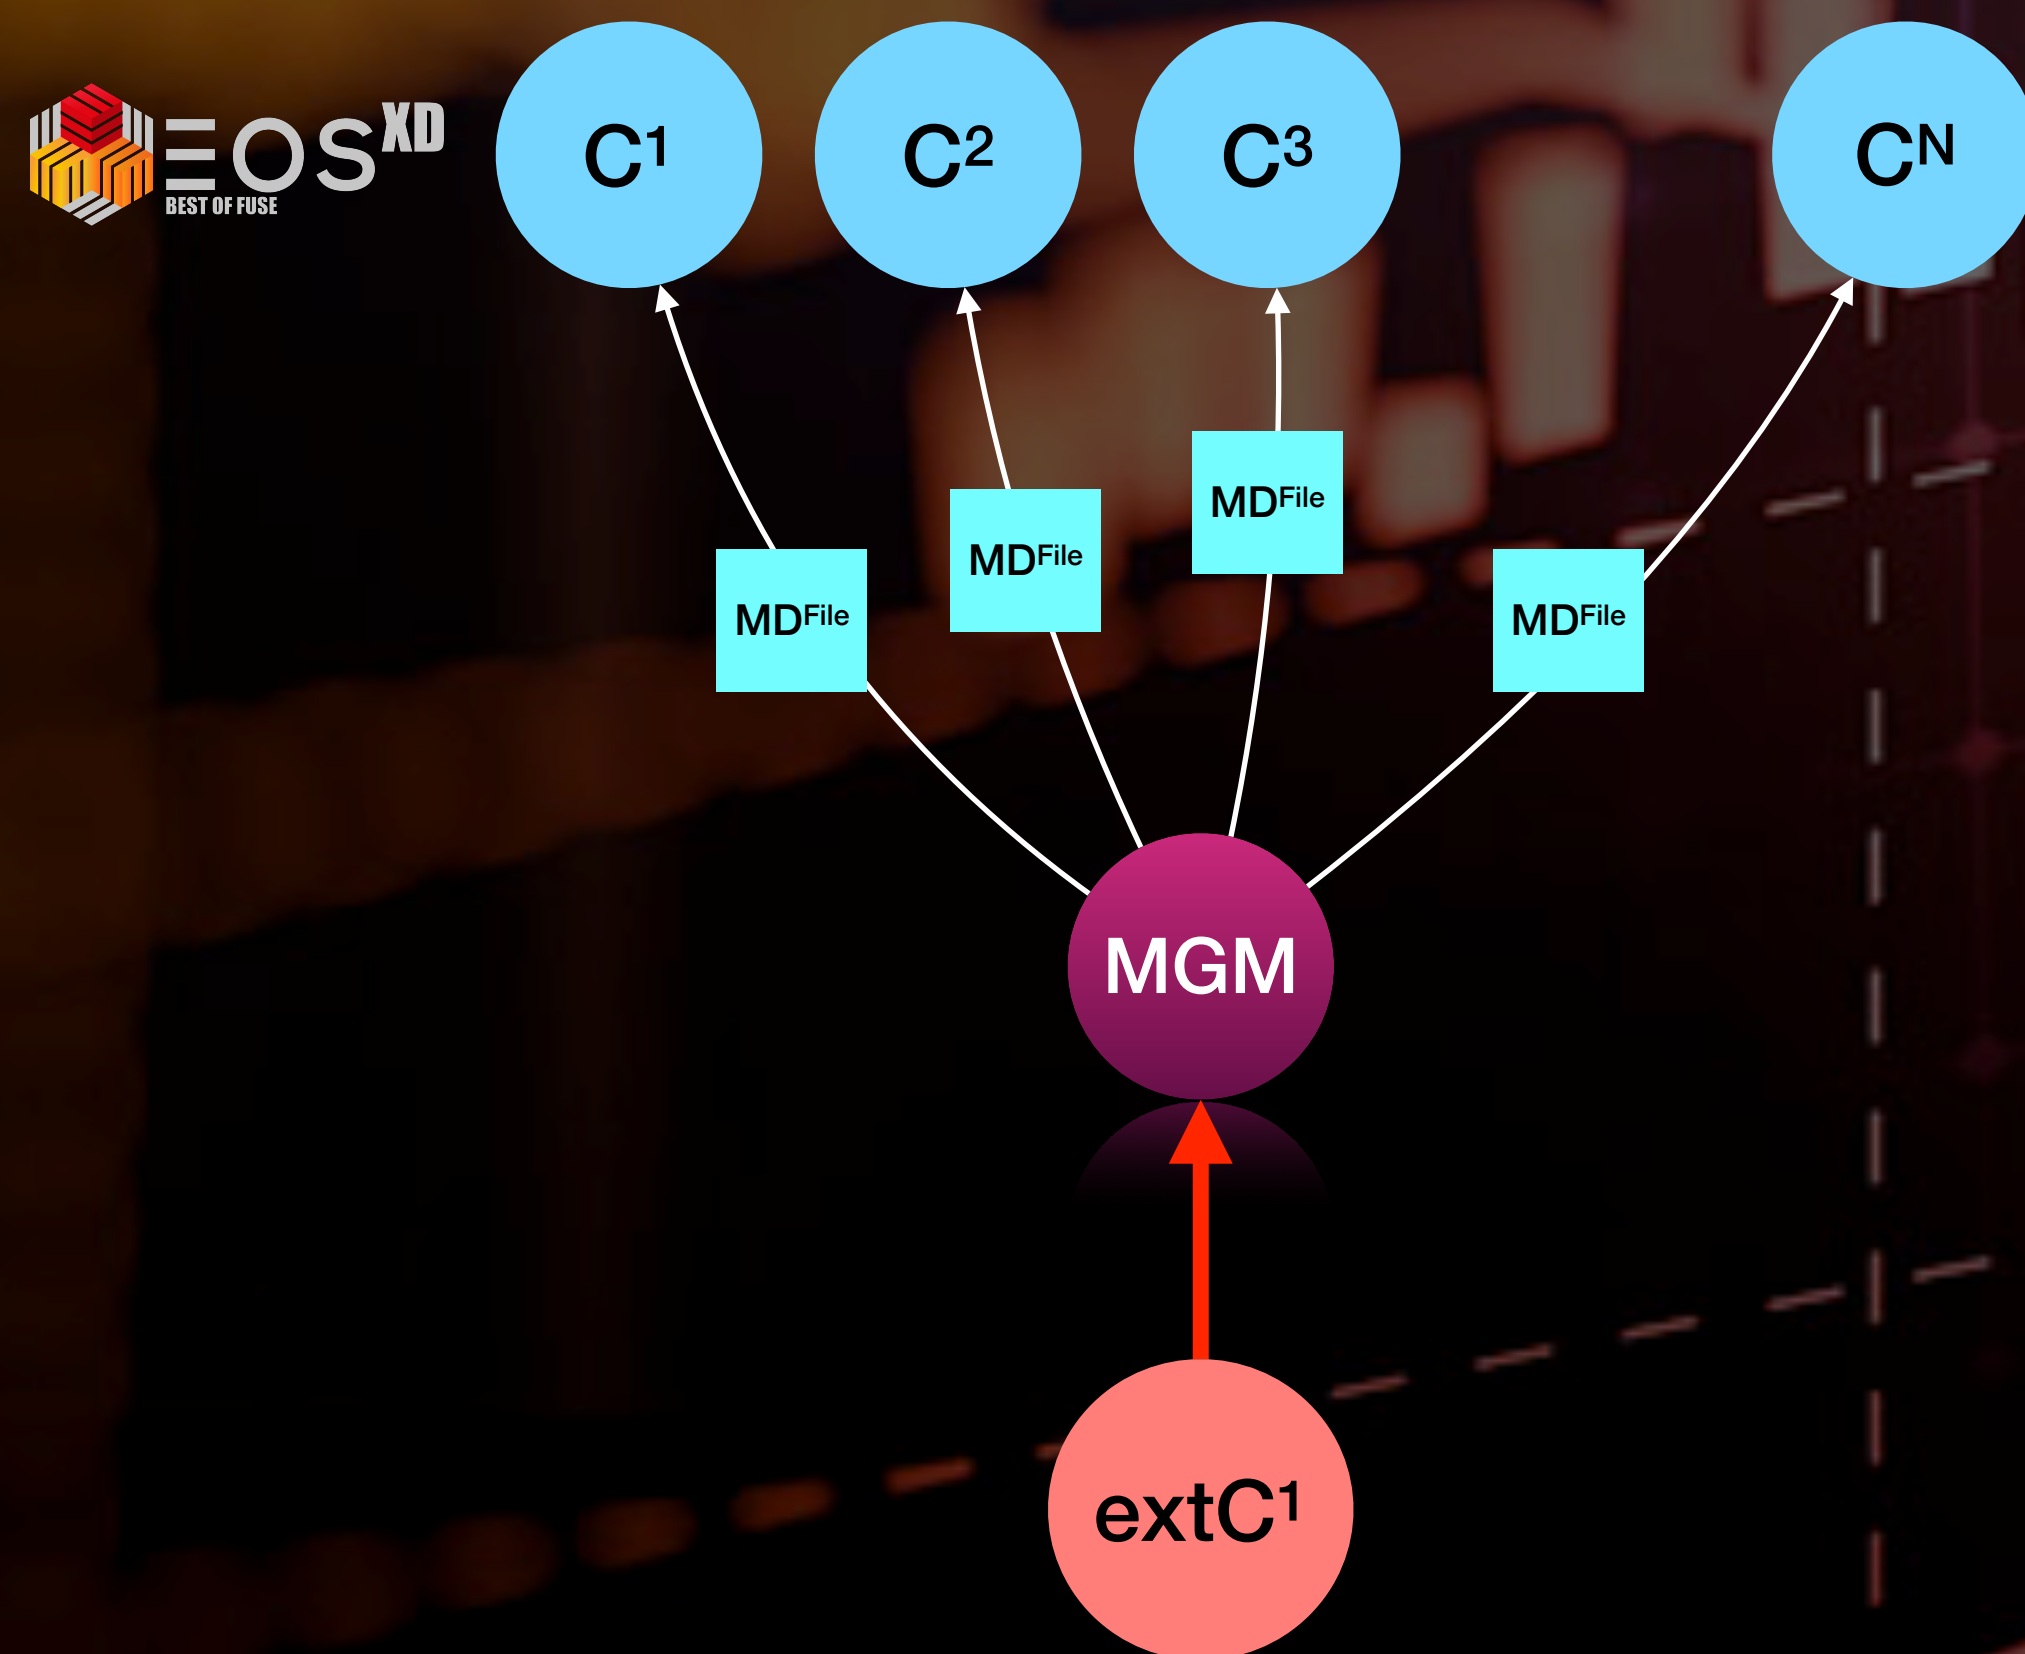


eosxd

Broadcasts & Callbacks

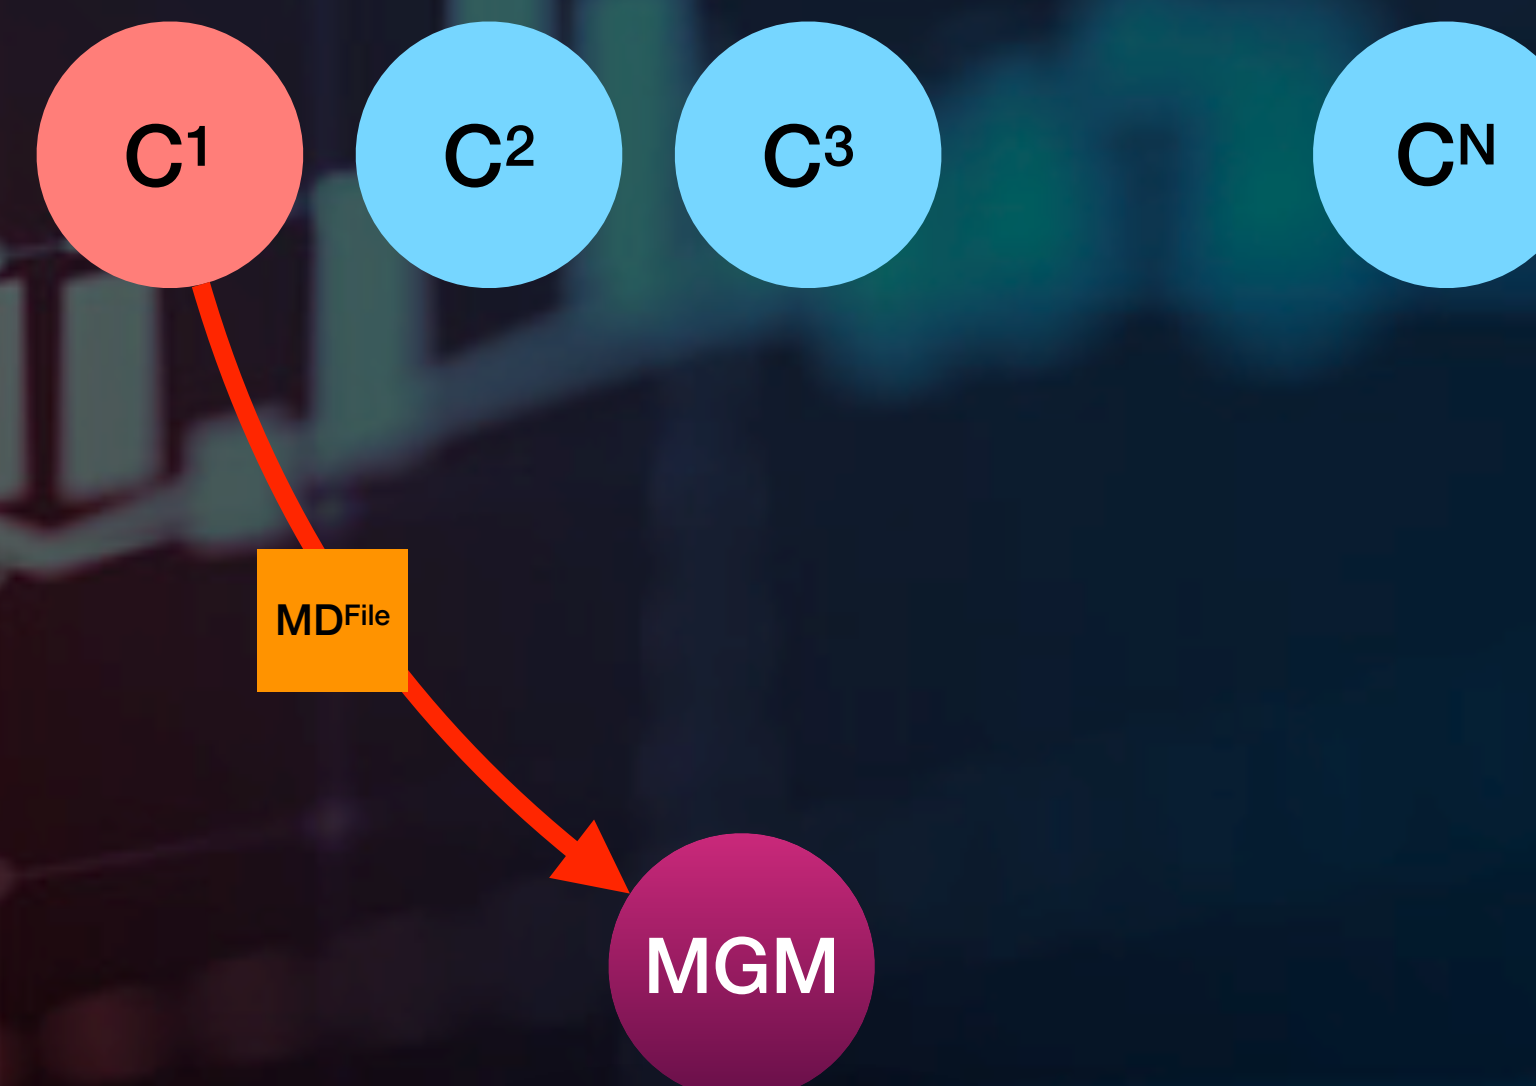
external broadcasts

an external Client (not FUSE) triggers a broadcast of file meta-data



internal broadcasts

an internal client triggers a broadcast of file meta-data




eosxd

Broadcasts & Callbacks

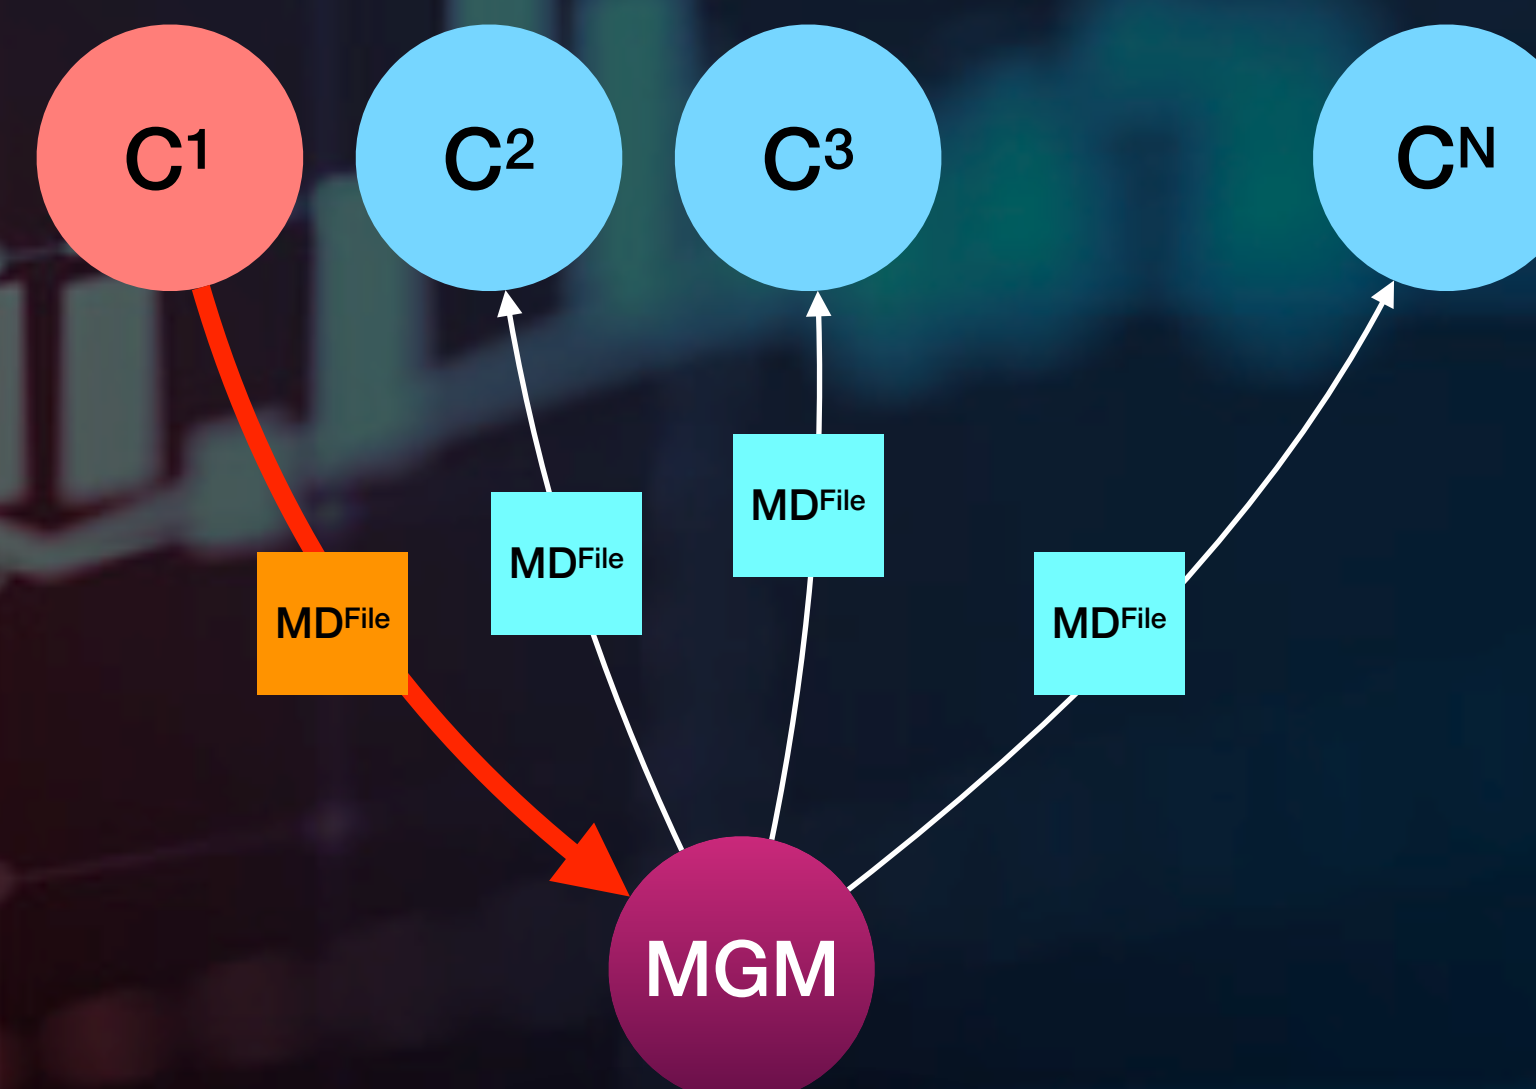
external broadcasts

an external Client (not FUSE) triggers a broadcast of file meta-data



internal broadcasts

an internal client triggers a broadcast of file meta-data



eosxd

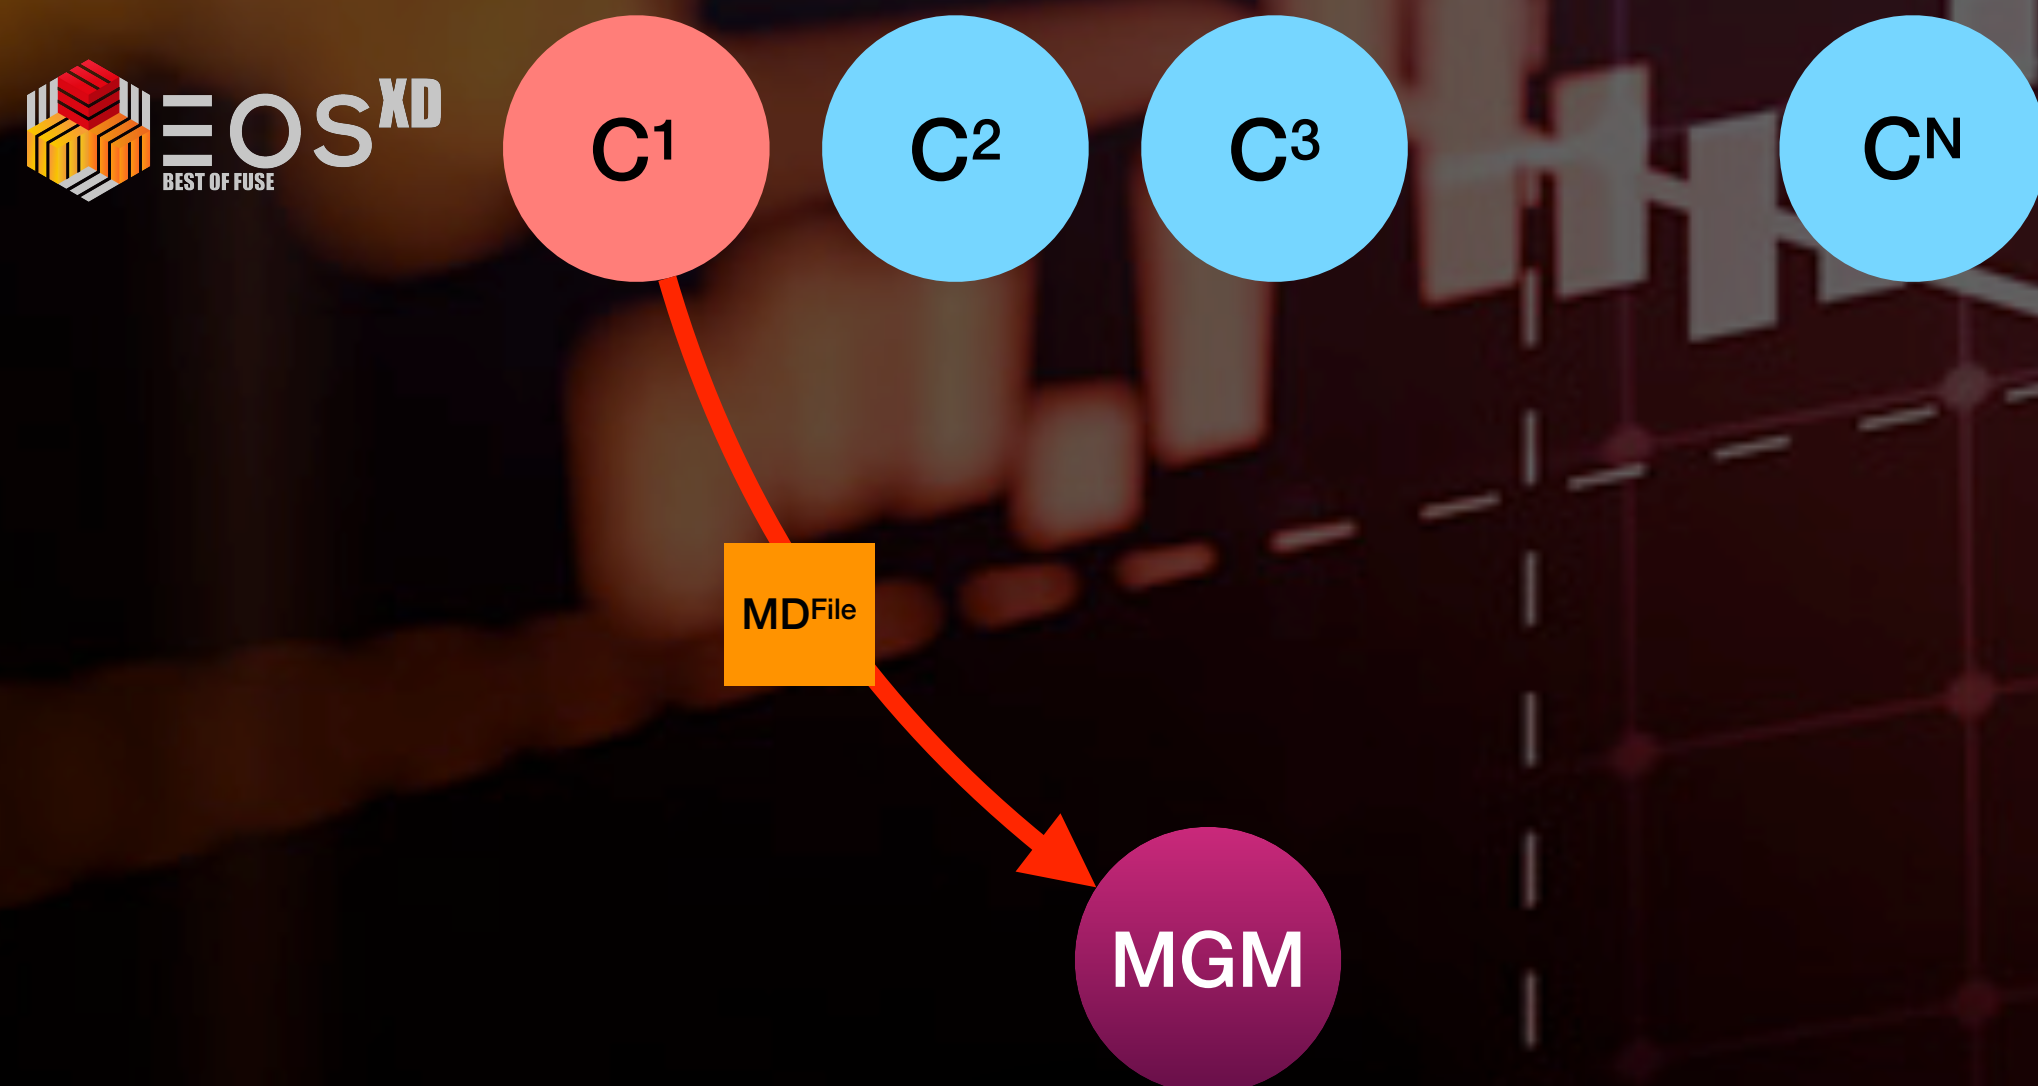
Broadcast Suppression



- an internal client would trigger a **broad cast** of file meta-data
- the receiving **audience is too large** and the broadcast gets suppressed
- the not updating clients keep with an **unsynchronised state** until their subscriptions expired and a directory refresh is done
 - **up to 300s by default**

eosxd

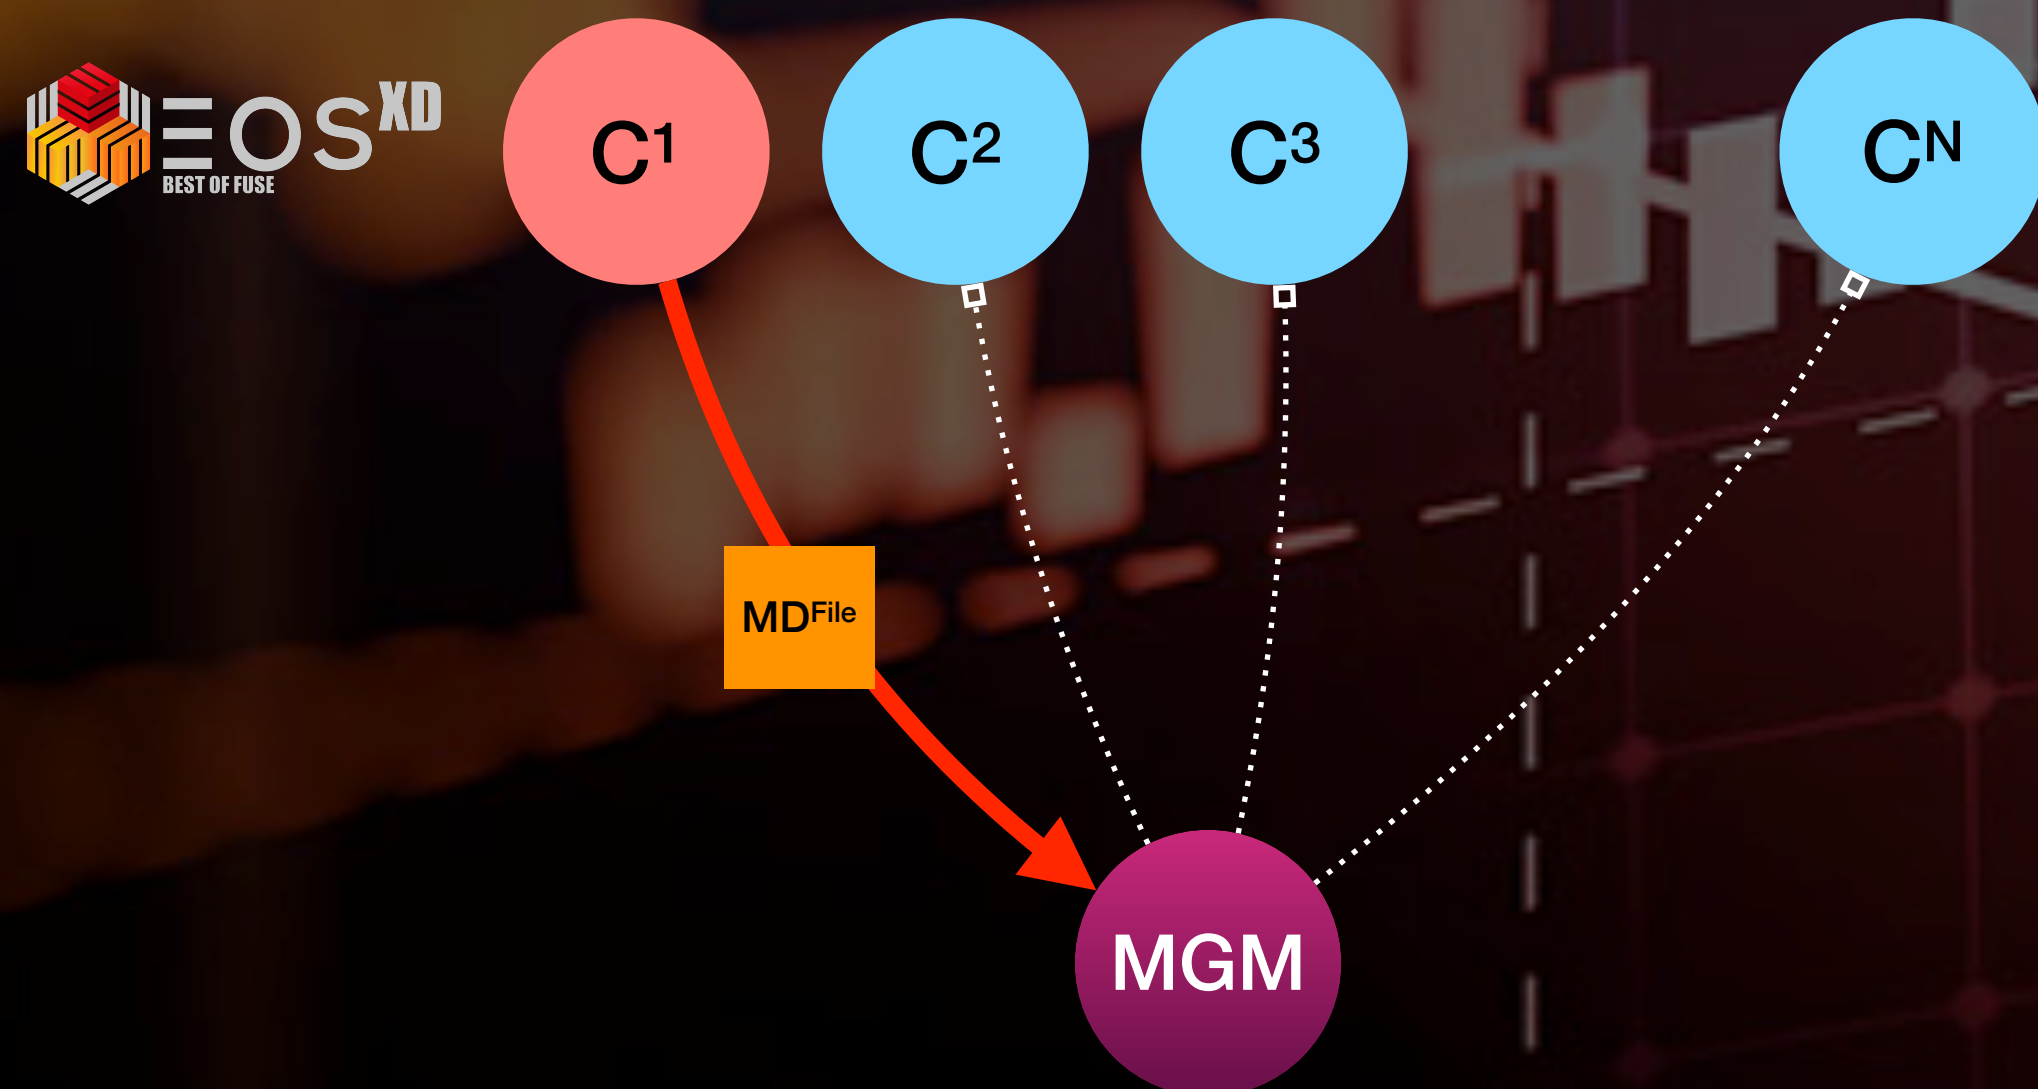
Broadcast Suppression



- an internal client would trigger a **broad cast** of file meta-data
- the receiving **audience is too large** and the broadcast gets suppressed
- the not updating clients keep with an **unsynchronised state** until their subscriptions expired and a directory refresh is done
 - **up to 300s by default**

eosxd

Broadcast Suppression



- an internal client would trigger a **broad cast** of file meta-data
- the receiving **audience is too large** and the broadcast gets suppressed
- the not updating clients keep with an **unsynchronised state** until their subscriptions expired and a directory refresh is done
 - **up to 300s by default**

eosxd

Namespace FuseServer



eos fusex ls

lient : eosxd	b7g11p1215.cern.ch	5.1.14	online	Wed, 19 Apr 2023	11:10:21	GMT	5.75	1.09	c6a0cff8-dea2-11ed-b606-80d4a5b16cc9	p=32919	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
lient : eosxd	b7g11p1215.cern.ch	5.1.14	online	Wed, 19 Apr 2023	10:50:17	GMT	9.58	1.15	f8974f94-de9f-11ed-b8d9-80d4a5b16cc9	p=15866	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
lient : eosxd	b7g11p1233.cern.ch	5.1.14	online	Wed, 19 Apr 2023	01:32:17	GMT	8.29	0.11	052cbb16-de52-11ed-a9b0-f02fa78bd033	p=16475	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
lient : eosxd	b7g11p1233.cern.ch	5.1.14	online	Wed, 19 Apr 2023	09:41:30	GMT	3.47	1.06	5cb5607e-de96-11ed-8c96-f02fa78bd033	p=28162	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
lient : eosxd	b7g11p1233.cern.ch	5.1.14	online	Tue, 18 Apr 2023	20:14:33	GMT	0.72	0.42	a254458a-de25-11ed-9f2b-f02fa78bd033	p=4924	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
lient : eosxd	b7g11p1233.cern.ch	5.1.14	online	Tue, 11 Apr 2023	10:16:06	GMT	4.14	0.82	deaf092c-d851-11ed-9ee9-f02fa78bd033	p=29242	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
lient : eosxd	b7g11p1247.cern.ch	5.1.14	online	Wed, 19 Apr 2023	14:20:41	GMT	2.88	1.14	5d61f9ac-debd-11ed-8591-80d4a5b16ca9	p=48105	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
lient : eosxd	b7g11p1247.cern.ch	5.1.14	online	Wed, 19 Apr 2023	13:48:57	GMT	5.66	0.23	ee644f86-deb8-11ed-a1db-80d4a5b16ca9	p=41863	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
lient : eosxd	b7g11p1277.cern.ch	5.1.14	online	Tue, 18 Apr 2023	23:25:57	GMT	3.66	0.12	5f0047d2-de40-11ed-851a-f02fa78bd06b	p=34036	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
lient : eosxd	b7g11p1277.cern.ch	5.1.14	online	Tue, 18 Apr 2023	05:49:22	GMT	0.12	0.49	c4c82e16-ddac-11ed-b1b1-f02fa78bd06b	p=31930	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
lient : eosxd	b7g11p1277.cern.ch	5.1.14	online	Wed, 19 Apr 2023	15:21:22	GMT	4.39	0.14	d73d1452-dec5-11ed-8776-f02fa78bd06b	p=26841	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
lient : eosxd	b7g11p1277.cern.ch	5.1.14	online	Wed, 19 Apr 2023	09:38:14	GMT	9.33	0.18	e8367846-de95-11ed-813b-f02fa78bd06b	p=40887	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
lient : eosxd	b7g11p1358.cern.ch	5.1.14	online	Wed, 19 Apr 2023	19:42:32	GMT	3.73	0.24	53a4c1b0-deea-11ed-9a8f-80d4a5b16c87	p=40981	caps=4	fds=0	autofs	[vacant]	act mount=/eos/home



Client : eosxd	b7g11p1215.cern.ch	5.1.14	online	Wed, 19 Apr 2023	11:10:21 GMT	5.75	1.09	c6a0cff8-dea2-11ed-b606-80d4a5b16cc9	p=32919	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
Client : eosxd	b7g11p1215.cern.ch	5.1.14	online	Wed, 19 Apr 2023	10:50:17 GMT	9.58	1.15	f8974f94-de9f-11ed-b8d9-80d4a5b16cc9	p=15866	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
Client : eosxd	b7g11p1233.cern.ch	5.1.14	online	Wed, 19 Apr 2023	01:32:17 GMT	8.29	0.11	052cbb16-de52-11ed-a9b0-f02fa78bd033	p=16475	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
Client : eosxd	b7g11p1233.cern.ch	5.1.14	online	Wed, 19 Apr 2023	09:41:30 GMT	3.47	1.06	5cb5607e-de96-11ed-8c96-f02fa78bd033	p=28162	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
Client : eosxd	b7g11p1233.cern.ch	5.1.14	online	Tue, 18 Apr 2023	20:14:33 GMT	0.72	0.42	a254458a-de25-11ed-9f2b-f02fa78bd033	p=4924	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
Client : eosxd	b7g11p1233.cern.ch	5.1.14	online	Tue, 11 Apr 2023	10:16:06 GMT	4.14	0.82	deaf092c-d851-11ed-9ee9-f02fa78bd033	p=29242	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
Client : eosxd	b7g11p1247.cern.ch	5.1.14	online	Wed, 19 Apr 2023	14:20:41 GMT	2.88	1.14	5d61f9ac-debd-11ed-8591-80d4a5b16ca9	p=48105	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
Client : eosxd	b7g11p1247.cern.ch	5.1.14	online	Wed, 19 Apr 2023	13:48:57 GMT	5.66	0.23	ee644f86-deb8-11ed-a1db-80d4a5b16ca9	p=41863	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
Client : eosxd	b7g11p1277.cern.ch	5.1.14	online	Tue, 18 Apr 2023	23:25:57 GMT	3.66	0.12	5f0047d2-de40-11ed-851a-f02fa78bd06b	p=34036	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
Client : eosxd	b7g11p1277.cern.ch	5.1.14	online	Tue, 18 Apr 2023	05:49:22 GMT	0.12	0.49	c4c82e16-ddac-11ed-b1b1-f02fa78bd06b	p=31930	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
Client : eosxd	b7g11p1277.cern.ch	5.1.14	online	Wed, 19 Apr 2023	15:21:22 GMT	4.39	0.14	d73d1452-dec5-11ed-8776-f02fa78bd06b	p=26841	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
Client : eosxd	b7g11p1277.cern.ch	5.1.14	online	Wed, 19 Apr 2023	09:38:14 GMT	9.33	0.18	e8367846-de95-11ed-813b-f02fa78bd06b	p=40887	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
Client : eosxd	b7g11p1358.cern.ch	5.1.14	online	Wed, 19 Apr 2023	19:42:32 GMT	3.73	0.24	53a4c1b0-deea-11ed-9a8f-80d4a5b16c87	p=40981	caps=4	fds=0	autofs	[vacant]	act mount=/eos/home



eos fusex ls

```
list client nodes
```

client	ip	hostname	version	status	date	time	tz	lat	lon	id	key	value	caps	fds	autofs	act	mount
client : eosxd	b7g11p1215.cern.ch	5.1.14	online	Wed, 19 Apr 2023	11:10:21	GMT	5.75	1.09	c6a0cff8-dea2-11ed-b606-80d4a5b16cc9	p=32919	caps=0	fds=0	autofs	[vacant]	act	mount=/eos/home	
client : eosxd	b7g11p1215.cern.ch	5.1.14	online	Wed, 19 Apr 2023	10:50:17	GMT	9.58	1.15	f8974f94-de9f-11ed-b8d9-80d4a5b16cc9	p=15866	caps=0	fds=0	autofs	[vacant]	act	mount=/eos/home	
client : eosxd	b7g11p1233.cern.ch	5.1.14	online	Wed, 19 Apr 2023	01:32:17	GMT	8.29	0.11	052cbb16-de52-11ed-a9b0-f02fa78bd033	p=16475	caps=0	fds=0	autofs	[vacant]	act	mount=/eos/home	
client : eosxd	b7g11p1233.cern.ch	5.1.14	online	Wed, 19 Apr 2023	01:30	GMT	3.47	1.06	5cb5607e-de96-11ed-8c96-f02fa78bd033	p=28162	caps=0	fds=0	autofs	[vacant]	act	mount=/eos/home	
client : eosxd	b7g11p1233.cern.ch	5.1.14	online	Tue, 18 Apr 2023	20:14:33	GMT	0.72	0.42	a254458a-de25-11ed-9f2b-f02fa78bd033	p=4924	caps=0	fds=0	autofs	[vacant]	act	mount=/eos/home	
client : eosxd	b7g11p1233.cern.ch	5.1.14	online	Tue, 11 Apr 2023	10:16:06	GMT	4.14	0.82	deaf092c-d851-11ed-9ee9-f02fa78bd033	p=29242	caps=0	fds=0	autofs	[vacant]	act	mount=/eos/home	
client : eosxd	b7g11p1247.cern.ch	5.1.14	online	Wed, 19 Apr 2023	14:20:41	GMT	2.88	1.14	5d61f9ac-debd-11ed-8591-80d4a5b16ca9	p=48105	caps=0	fds=0	autofs	[vacant]	act	mount=/eos/home	
client : eosxd	b7g11p1247.cern.ch	5.1.14	online	Wed, 19 Apr 2023	13:48:57	GMT	5.66	0.23	ee644f86-deb8-11ed-a1db-80d4a5b16ca9	p=41863	caps=0	fds=0	autofs	[vacant]	act	mount=/eos/home	
client : eosxd	b7g11p1277.cern.ch	5.1.14	online	Tue, 18 Apr 2023	23:25:57	GMT	3.66	0.12	5f0047d2-de40-11ed-851a-f02fa78bd06b	p=34036	caps=0	fds=0	autofs	[vacant]	act	mount=/eos/home	
client : eosxd	b7g11p1277.cern.ch	5.1.14	online	Tue, 18 Apr 2023	05:49:22	GMT	0.12	0.49	c4c82e16-ddac-11ed-b1b1-f02fa78bd06b	p=31930	caps=0	fds=0	autofs	[vacant]	act	mount=/eos/home	
client : eosxd	b7g11p1277.cern.ch	5.1.14	online	Wed, 19 Apr 2023	15:21:22	GMT	4.39	0.14	d73d1452-dec5-11ed-8776-f02fa78bd06b	p=26841	caps=0	fds=0	autofs	[vacant]	act	mount=/eos/home	
client : eosxd	b7g11p1277.cern.ch	5.1.14	online	Wed, 19 Apr 2023	09:38:14	GMT	9.33	0.18	e8367846-de95-11ed-813b-f02fa78bd06b	p=40887	caps=0	fds=0	autofs	[vacant]	act	mount=/eos/home	
client : eosxd	b7g11p1358.cern.ch	5.1.14	online	Wed, 19 Apr 2023	19:42:32	GMT	3.73	0.24	53a4c1b0-deea-11ed-9a8f-80d4a5b16c87	p=40981	caps=4	fds=0	autofs	[vacant]	act	mount=/eos/home	



eos fusex ls

```
list client nodes
```

eos fusex ls -l

list client nodes and some statistics per client

Client	IP	Version	Status	Time	Timezone	Lat	Long	Client ID	Client Name	Client Type	Client State	Client Mount
client : eosxd	b7g11p1215.cern.ch	5.1.14	online	Wed, 19 Apr 2023 11:10:21	GMT	5.75	1.09	c6a0cff8-dea2-11ed-b606-80d4a5b16cc9	p=32919	caps=0	fds=0	autofs [vacant] act mount=/eos/home
client : eosxd	b7g11p1215.cern.ch	5.1.14	online	Wed, 19 Apr 2023 10:50:17	GMT	9.58	1.15	f8974f94-de9f-11ed-b8d9-80d4a5b16cc9	p=15866	caps=0	fds=0	autofs [vacant] act mount=/eos/home
client : eosxd	b7g11p1233.cern.ch	5.1.14	online	Wed, 19 Apr 2023 01:32:17	GMT	8.29	0.11	052cbb16-de52-11ed-a9b0-f02fa78bd033	p=16475	caps=0	fds=0	autofs [vacant] act mount=/eos/home
client : eosxd	b7g11p1233.cern.ch	5.1.14	online	Wed, 19 Apr 2023 01:11:30	GMT	3.47	1.06	5cb5607e-de96-11ed-8c96-f02fa78bd033	p=28162	caps=0	fds=0	autofs [vacant] act mount=/eos/home
client : eosxd	b7g11p1233.cern.ch	5.1.14	online	Tue, 18 Apr 2023 20:14:33	GMT	0.72	0.42	a254458a-de25-11ed-9f2b-f02fa78bd033	p=4924	caps=0	fds=0	autofs [vacant] act mount=/eos/home
client : eosxd	b7g11p1233.cern.ch	5.1.14	online	Tue, 11 Apr 2023 10:16:06	GMT	4.14	0.82	deaf092c-d851-11ed-9ee9-f02fa78bd033	p=29242	caps=0	fds=0	autofs [vacant] act mount=/eos/home
client : eosxd	b7g11p1247.cern.ch	5.1.14	online	Wed, 19 Apr 2023 14:20:41	GMT	2.88	1.14	5d61f9ac-debd-11ed-8591-80d4a5b16ca9	p=48105	caps=0	fds=0	autofs [vacant] act mount=/eos/home
client : eosxd	b7g11p1247.cern.ch	5.1.14	online	Wed, 19 Apr 2023 11:08:08	GMT	5.66	0.23	ee644f86-deb8-11ed-a1db-80d4a5b16ca9	p=41863	caps=0	fds=0	autofs [vacant] act mount=/eos/home
client : eosxd	b7g11p1277.cern.ch	5.1.14	online	Tue, 18 Apr 2023 23:25:57	GMT	3.66	0.12	5f0047d2-de40-11ed-851a-f02fa78bd06b	p=34036	caps=0	fds=0	autofs [vacant] act mount=/eos/home
client : eosxd	b7g11p1277.cern.ch	5.1.14	online	Tue, 18 Apr 2023 15:21:22	GMT	0.12	0.49	c4c82e16-ddac-11ed-b1b1-f02fa78bd06b	p=31930	caps=0	fds=0	autofs [vacant] act mount=/eos/home
client : eosxd	b7g11p1277.cern.ch	5.1.14	online	Wed, 19 Apr 2023 15:21:22	GMT	4.39	0.14	d73d1452-dec5-11ed-8776-f02fa78bd06b	p=26841	caps=0	fds=0	autofs [vacant] act mount=/eos/home
client : eosxd	b7g11p1277.cern.ch	5.1.14	online	Wed, 19 Apr 2023 09:38:14	GMT	9.33	0.18	e8367846-de95-11ed-813b-f02fa78bd06b	p=40887	caps=0	fds=0	autofs [vacant] act mount=/eos/home
client : eosxd	b7g11p1358.cern.ch	5.1.14	online	Wed, 19 Apr 2023 19:42:32	GMT	3.73	0.24	53a4c1b0-deea-11ed-9a8f-80d4a5b16c87	p=40981	caps=4	fds=0	autofs [vacant] act mount=/eos/home

eosxd

Namespace FuseServer

eos fusex ls

client : eosxd	b7g11p1215.cern.ch 5.1.14	online	Wed, 19 Apr 2023 11:10:21 GMT 5.75	1.09	c6a0cff8-dea2-11ed-b606-80d4a5b16cc9	p=32919	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
client : eosxd	b7g11p1215.cern.ch 5.1.14	online	Wed, 19 Apr 2023 10:50:17 GMT 9.58	1.15	f8974f94-de9f-11ed-b8d9-80d4a5b16cc9	p=15866	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
client : eosxd	b7g11p1233.cern.ch 5.1.14	online	Wed, 19 Apr 2023 01:32:17 GMT 8.29	0.11	052cbb16-de52-11ed-a9b0-f02fa78bd033	p=16475	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
client : eosxd	b7g11p1233.cern.ch 5.1.14	online	Wed, 19 Apr 2023 01:41:30 GMT 3.47	1.06	5cb5607e-de96-11ed-8c96-f02fa78bd033	p=28162	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
client : eosxd	b7g11p1233.cern.ch 5.1.14	online	Tue, 18 Apr 2023 20:14:33 GMT 0.72	0.42	a254458a-de25-11ed-9f2b-f02fa78bd033	p=4924	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
client : eosxd	b7g11p1233.cern.ch 5.1.14	online	Tue, 11 Apr 2023 10:16:06 GMT 4.14	0.82	deaf092c-d851-11ed-9ee9-f02fa78bd033	p=29242	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
client : eosxd	b7g11p1247.cern.ch 5.1.14	online	Wed, 19 Apr 2023 14:20:41 GMT 2.88	1.14	5d61f9ac-debd-11ed-8591-80d4a5b16ca9	p=48105	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
client : eosxd	b7g11p1247.cern.ch 5.1.14	online	Wed, 19 Apr 2023 13:08:13 GMT 5.66	0.23	ee644f86-deb8-11ed-a1db-80d4a5b16ca9	p=41863	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
client : eosxd	b7g11p1277.cern.ch 5.1.14	online	Tue, 18 Apr 2023 23:25:57 GMT 3.66	0.12	5f0047d2-de40-11ed-851a-f02fa78bd06b	p=34036	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
client : eosxd	b7g11p1277.cern.ch 5.1.14	online	Tue, 18 Apr 2023 03:45:12 GMT 0.12	0.49	c4c82e16-ddac-11ed-b1b1-f02fa78bd06b	p=31930	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
client : eosxd	b7g11p1277.cern.ch 5.1.14	online	Wed, 19 Apr 2023 15:21:22 GMT 4.39	0.14	d73d1452-dec5-11ed-8776-f02fa78bd06b	p=26841	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
client : eosxd	b7g11p1277.cern.ch 5.1.14	online	Wed, 19 Apr 2023 09:38:14 GMT 9.33	0.18	e8367846-de95-11ed-813b-f02fa78bd06b	p=40887	caps=0	fds=0	autofs	[vacant]	act mount=/eos/home
client : eosxd	b7g11p1358.cern.ch 5.1.14	online	Wed, 19 Apr 2023 19:42:32 GMT 3.73	0.24	53a4c1b0-deea-11ed-9a8f-80d4a5b16c87	p=40981	caps=4	fds=0	autofs	[vacant]	act mount=/eos/home

eos fusex ls

list client nodes

eos fusex ls -l

list client nodes and some
statistics per client

eos fusex evict

force unmount, trigger a stack trace,
ask for logfile, truncate logfiles

eosxd

Namespace FuseServer

eos fusex ls

client : eosxd b7g11p1215.cern.ch 5.1.14 online Wed, 19 Apr 2023 11:10:21 GMT 5.75 1.09 c6a0cff8-dea2-11ed-b606-80d4a5b16cc9 p=32919 caps=0 fds=0 autofs [vacant] act mount=/eos/home	
client : eosxd b7g11p1215.cern.ch 5.1.14 online Wed, 19 Apr 2023 10:50:17 GMT 9.58 1.15 f8974f94-de9f-11ed-b8d9-80d4a5b16cc9 p=15866 caps=0 fds=0 autofs [vacant] act mount=/eos/home	
client : eosxd b7g11p1233.cern.ch 5.1.14 online Wed, 19 Apr 2023 01:32:17 GMT 8.29 0.11 052cbb16-de52-11ed-a9b0-f02fa78bd033 p=16475 caps=0 fds=0 autofs [vacant] act mount=/eos/home	
client : eosxd b7g11p1233.cern.ch 5.1.14 online Wed, 19 Apr 2023 01:41:30 GMT 3.47 1.06 5cb5607e-de96-11ed-8c96-f02fa78bd033 p=28162 caps=0 fds=0 autofs [vacant] act mount=/eos/home	
client : eosxd b7g11p1233.cern.ch 5.1.14 online Tue, 18 Apr 2023 20:14:33 GMT 0.72 0.42 a254458a-de25-11ed-9f2b-f02fa78bd033 p=4924 caps=0 fds=0 autofs [vacant] act mount=/eos/home	
client : eosxd b7g11p1233.cern.ch 5.1.14 online Tue, 11 Apr 2023 10:16:06 GMT 4.14 0.82 deaf092c-d851-11ed-9ee9-f02fa78bd033 p=29242 caps=0 fds=0 autofs [vacant] act mount=/eos/home	
client : eosxd b7g11p1247.cern.ch 5.1.14 online Wed, 19 Apr 2023 14:20:41 GMT 2.88 1.14 5d61f9ac-debd-11ed-8591-80d4a5b16ca9 p=48105 caps=0 fds=0 autofs [vacant] act mount=/eos/home	
client : eosxd b7g11p1247.cern.ch 5.1.14 online Wed, 19 Apr 2023 13:08:13 GMT 5.66 0.23 ee644f86-deb8-11ed-a1db-80d4a5b16ca9 p=41863 caps=0 fds=0 autofs [vacant] act mount=/eos/home	
client : eosxd b7g11p1277.cern.ch 5.1.14 online Tue, 18 Apr 2023 23:25:57 GMT 3.66 0.12 5f0047d2-de40-11ed-851a-f02fa78bd06b p=34036 caps=0 fds=0 autofs [vacant] act mount=/eos/home	
client : eosxd b7g11p1277.cern.ch 5.1.14 online Tue, 18 Apr 2023 09:49:12 GMT 0.12 0.49 c4c82e16-ddac-11ed-b1b1-f02fa78bd06b p=31930 caps=0 fds=0 autofs [vacant] act mount=/eos/home	
client : eosxd b7g11p1277.cern.ch 5.1.14 online Wed, 19 Apr 2023 15:21:22 GMT 4.39 0.14 d73d1452-dec5-11ed-8776-f02fa78bd06b p=26841 caps=0 fds=0 autofs [vacant] act mount=/eos/home	
client : eosxd b7g11p1277.cern.ch 5.1.14 online Wed, 19 Apr 2023 09:38:14 GMT 9.33 0.18 e8367846-de95-11ed-813b-f02fa78bd06b p=40887 caps=0 fds=0 autofs [vacant] act mount=/eos/home	
client : eosxd b7g11p1358.cern.ch 5.1.14 online Wed, 19 Apr 2023 19:42:32 GMT 3.73 0.24 53a4c1b0-deea-11ed-9a8f-80d4a5b16c87 p=40981 caps=4 fds=0 autofs [vacant] act mount=/eos/home	
eos fusex ls	list client nodes
eos fusex ls -l	list client nodes and some statistics per client
eos fusex evict	force unmount, trigger a stack trace, ask for logfile, truncate logfiles
eos fusex conf	configure heartbeat interval configure audience suppression

eosxd

Namespace FuseServer

eos fusex ls		eos fusex ls
eos fusex ls -l		list client nodes and some statistics per client
eos fusex evict		force unmount, trigger a stack trace, ask for logfile, truncate logfiles
eos fusex conf		configure heartbeat interval configure audience suppression
eos fusex caps		show directory subscriptions

eosxd

Namespace FuseServer

	eos fusex ls	list client nodes
	eos fusex ls -l	list client nodes and some statistics per client
	eos fusex evict	force unmount, trigger a stack trace, ask for logfile, truncate logfiles
	eos fusex conf	configure heartbeat interval configure audience suppression
	eos fusex caps	show directory subscriptions
	eos fusex dropcap(s)	drop one or all subscriptions from a client

```
eos fusex ls
```

```
client : eosxd b7g11p1215.cern.ch 5.1.14 online Wed, 19 Apr 2023 11:10:21 GMT 5.75 1.09 c6a0cff8-dea2-11ed-b606-80d4a5b16cc9 p=32919 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1215.cern.ch 5.1.14 online Wed, 19 Apr 2023 10:50:17 GMT 9.58 1.15 f8974f94-de9f-11ed-b8d9-80d4a5b16cc9 p=15866 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1233.cern.ch 5.1.14 online Wed, 19 Apr 2023 01:32:17 GMT 8.29 0.11 052cbb16-de52-11ed-a9b0-f02fa78bd033 p=16475 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1233.cern.ch 5.1.14 online Wed, 19 Apr 2023 01:11:30 GMT 3.47 1.06 5cb5607e-de96-11ed-8c96-f02fa78bd033 p=28162 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1233.cern.ch 5.1.14 online Tue, 18 Apr 2023 20:14:33 GMT 0.72 0.42 a254458a-de25-11ed-9f2b-f02fa78bd033 p=4924 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1233.cern.ch 5.1.14 online Tue, 11 Apr 2023 10:16:06 GMT 4.14 0.82 deaf092c-d851-11ed-9ee9-f02fa78bd033 p=29242 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1247.cern.ch 5.1.14 online Wed, 19 Apr 2023 14:20:41 GMT 2.88 1.14 5d61f9ac-debd-11ed-8591-80d4a5b16ca9 p=48105 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1247.cern.ch 5.1.14 online Wed, 19 Apr 2023 13:13:38 GMT 5.66 0.23 ee644f86-deb8-11ed-a1db-80d4a5b16ca9 p=41863 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1277.cern.ch 5.1.14 online Tue, 18 Apr 2023 23:25:57 GMT 3.66 0.12 5f0047d2-de40-11ed-851a-f02fa78bd06b p=34036 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1277.cern.ch 5.1.14 online Tue, 18 Apr 2023 03:45:12 GMT 0.12 0.49 c4c82e16-ddac-11ed-b1b1-f02fa78bd06b p=31930 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1277.cern.ch 5.1.14 online Wed, 19 Apr 2023 15:21:22 GMT 4.39 0.14 d73d1452-dec5-11ed-8776-f02fa78bd06b p=26841 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1277.cern.ch 5.1.14 online Wed, 19 Apr 2023 09:38:14 GMT 9.33 0.18 e8367846-de95-11ed-813b-f02fa78bd06b p=40887 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1358.cern.ch 5.1.14 online Wed, 19 Apr 2023 19:42:32 GMT 3.73 0.24 53a4c1b0-deea-11ed-9a8f-80d4a5b16c87 p=40981 caps=4 fds=0 autofs [vacant] act mount=/eos/home
```


eosxd

Namespace FuseServer

	eos fusex ls	list client nodes
	eos fusex ls -l	list client nodes and some statistics per client
	eos fusex evict	force unmount, trigger a stack trace, ask for logfile, truncate logfiles
	eos fusex conf	configure heartbeat interval configure audience suppression
	eos fusex caps	show directory subscriptions
	eos fusex dropcap(s)	drop one or all subscriptions from a client
	eos fusex droplocks	drop all locks for given inode and process

```
eos fusex ls
```

```
client : eosxd b7g11p1215.cern.ch 5.1.14 online Wed, 19 Apr 2023 11:10:21 GMT 5.75 1.09 c6a0cff8-dea2-11ed-b606-80d4a5b16cc9 p=32919 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1215.cern.ch 5.1.14 online Wed, 19 Apr 2023 10:50:17 GMT 9.58 1.15 f8974f94-de9f-11ed-b8d9-80d4a5b16cc9 p=15866 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1233.cern.ch 5.1.14 online Wed, 19 Apr 2023 01:32:17 GMT 8.29 0.11 052cbb16-de52-11ed-a9b0-f02fa78bd033 p=16475 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1233.cern.ch 5.1.14 online Wed, 19 Apr 2023 01:11:30 GMT 3.47 1.06 5cb5607e-de96-11ed-8c96-f02fa78bd033 p=28162 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1233.cern.ch 5.1.14 online Tue, 18 Apr 2023 20:14:33 GMT 0.72 0.42 a254458a-de25-11ed-9f2b-f02fa78bd033 p=4924 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1233.cern.ch 5.1.14 online Tue, 11 Apr 2023 10:16:06 GMT 4.14 0.82 deaf092c-d851-11ed-9ee9-f02fa78bd033 p=29242 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1247.cern.ch 5.1.14 online Wed, 19 Apr 2023 14:20:41 GMT 2.88 1.14 5d61f9ac-debd-11ed-8591-80d4a5b16ca9 p=48105 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1247.cern.ch 5.1.14 online Wed, 19 Apr 2023 13:13:38 GMT 5.66 0.23 ee644f86-deb8-11ed-a1db-80d4a5b16ca9 p=41863 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1277.cern.ch 5.1.14 online Tue, 18 Apr 2023 23:25:57 GMT 3.66 0.12 5f0047d2-de40-11ed-851a-f02fa78bd06b p=34036 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1277.cern.ch 5.1.14 online Tue, 18 Apr 2023 03:45:12 GMT 0.12 0.49 c4c82e16-ddac-11ed-b1b1-f02fa78bd06b p=31930 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1277.cern.ch 5.1.14 online Wed, 19 Apr 2023 15:21:22 GMT 4.39 0.14 d73d1452-dec5-11ed-8776-f02fa78bd06b p=26841 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1277.cern.ch 5.1.14 online Wed, 19 Apr 2023 09:38:14 GMT 9.33 0.18 e8367846-de95-11ed-813b-f02fa78bd06b p=40887 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1358.cern.ch 5.1.14 online Wed, 19 Apr 2023 19:42:32 GMT 3.73 0.24 53a4c1b0-deea-11ed-9a8f-80d4a5b16c87 p=40981 caps=4 fds=0 autofs [vacant] act mount=/eos/home
```


eosxd

Namespace FuseServer

eos fusex ls

```
client : eosxd b7g11p1215.cern.ch 5.1.14 online Wed, 19 Apr 2023 11:10:21 GMT 5.75 1.09 c6a0cff8-dea2-11ed-b606-80d4a5b16cc9 p=32919 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1215.cern.ch 5.1.14 online Wed, 19 Apr 2023 10:50:17 GMT 9.58 1.15 f8974f94-de9f-11ed-b8d9-80d4a5b16cc9 p=15866 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1233.cern.ch 5.1.14 online Wed, 19 Apr 2023 01:32:17 GMT 8.29 0.11 052cbb16-de52-11ed-a9b0-f02fa78bd033 p=16475 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1233.cern.ch 5.1.14 online Wed, 19 Apr 2023 01:11:30 GMT 3.47 1.06 5cb5607e-de96-11ed-8c96-f02fa78bd033 p=28162 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1233.cern.ch 5.1.14 online Tue, 18 Apr 2023 20:14:33 GMT 0.72 0.42 a254458a-de25-11ed-9f2b-f02fa78bd033 p=4924 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1233.cern.ch 5.1.14 online Tue, 11 Apr 2023 10:16:06 GMT 4.14 0.82 deaf092c-d851-11ed-9ee9-f02fa78bd033 p=29242 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1247.cern.ch 5.1.14 online Wed, 19 Apr 2023 14:20:41 GMT 2.88 1.14 5d61f9ac-debd-11ed-8591-80d4a5b16ca9 p=48105 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1247.cern.ch 5.1.14 online Wed, 19 Apr 2023 13:13:38 GMT 5.66 0.23 ee644f86-deb8-11ed-a1db-80d4a5b16ca9 p=41863 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1277.cern.ch 5.1.14 online Tue, 18 Apr 2023 23:25:57 GMT 3.66 0.12 5f0047d2-de40-11ed-851a-f02fa78bd06b p=34036 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1277.cern.ch 5.1.14 online Tue, 18 Apr 2023 09:49:12 GMT 0.12 0.49 c4c82e16-ddac-11ed-b1b1-f02fa78bd06b p=31930 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1277.cern.ch 5.1.14 online Wed, 19 Apr 2023 15:21:22 GMT 4.39 0.14 d73d1452-dec5-11ed-8776-f02fa78bd06b p=26841 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1277.cern.ch 5.1.14 online Wed, 19 Apr 2023 09:38:14 GMT 9.33 0.18 e8367846-de95-11ed-813b-f02fa78bd06b p=40887 caps=0 fds=0 autofs [vacant] act mount=/eos/home
client : eosxd b7g11p1358.cern.ch 5.1.14 online Wed, 19 Apr 2023 19:42:32 GMT 3.73 0.24 53a4c1b0-deea-11ed-9a8f-80d4a5b16c87 p=40981 caps=4 fds=0 autofs [vacant] act mount=/eos/home
```

eos fusex ls

list client nodes

eos fusex ls -l

list client nodes and some statistics per client

eos fusex evict

force unmount, trigger a stack trace, ask for logfile, truncate logfiles

eos fusex conf

configure heartbeat interval
configure audience suppression

eos fusex caps

show directory subscriptions

eos fusex dropcap(s)

drop one or all subscriptions from a client

eos fusex droplocks

drop all locks for given inode and process

eos fusex ls -l

```
client : eosxd lxplus700.cern.ch 5.1.14 online
..... ino : 40868
..... ino-to-del : 0
..... ino-backlog : 0
..... ino-ever : 511893
..... ino-ever-del : 18426
..... threads : 60
..... total-ram : 29.988 GB
..... free-ram : 0.280 GB
..... vsize : 2.016 GB
..... rsize : 0.193 GB
..... wr-buf-mb : 0 MB
..... ra-buf-mb : 24 MB
..... load1 : 19.33
..... leasetime : 300 s
..... open-files : 22
..... logfile-size : 18576
..... rbytes : 20397308261
..... wbytes : 24307570949
..... n-op : 4656820
..... rd60 : 0.00 MB/s
..... wr60 : 0.33 MB/s
..... iops60 : 303.83
..... xoff : 830
..... ra-xoff : 0
..... ra-nobuf : 9489
..... wr-nobuf : 0
..... idle : 24
..... blockedms : 0.00 [ ]
```


eosxd

Configuration

a named mount like

eosxd -ofsname=*name* **/eos/**

eosxd

Configuration

a named mount like

eosxd -ofsname=*name* **/eos/**

uses this configuration file:

/etc/eos/fuse.*name***.conf**

eosxd

Configuration

a named mount like

eosxd -ofsname=*name* **/eos/**

uses this configuration file:

/etc/eos/fuse.*name***.conf**

configless

eosxd -ofsname=*instance.domain:/eos/* **/eos/**

eosxd

Client side Disk Caching

File & Journal Cache

```
"cache" : {  
  "type" : "disk",  
  "size-mb" : 512,  
  "size-ino" : 65536,  
  "file-cache-max-kb", 256,  
  "file-journal-max-kb", 131072  
  "journal-mb" : 2048,  
  "journal-ino" : 65536,  
  "clean-threshold" : 85.0,  
  "location" : "/var/cache/eos/fusex/cache/",  
  "journal" : "/var/cache/eos/fusex/journal/",  
  "read-ahead-strategy" : "static",  
  "read-ahead-bytes-nominal" : 262144,  
  "read-ahead-bytes-max" : 2097152,  
  "read-ahead-blocks-max" : 16,  
  "max-read-ahead-buffer" : 134217728,  
  "max-write-buffer" : 134217728  
}
```


eosxd

Client side Disk Caching

File & Journal Cache

```
"cache" : {  
  "type" : "disk",  
  "size-mb" : 512,  
  "size-ino" : 65536,  
  "file-cache-max-kb", 256,  
  "file-journal-max-kb", 131072  
  "journal-mb" : 2048,  
  "journal-ino" : 65536,  
  "clean-threshold" : 85.0,  
  "location" : "/var/cache/eos/fusex/cache/",  
  "journal" : "/var/cache/eos/fusex/journal/",  
  "read-ahead-strategy" : "static",  
  "read-ahead-bytes-nominal" : 262144,  
  "read-ahead-bytes-max" : 2097152,  
  "read-ahead-blocks-max" : 16,  
  "max-read-ahead-buffer" : 134217728,  
  "max-write-buffer" : 134217728  
}
```

- There are **two caches** which are persisted while a mount is still up

eosxd

Client side Disk Caching

File & Journal Cache

```
"cache" : {  
  "type" : "disk",  
  "size-mb" : 512,  
  "size-ino" : 65536,  
  "file-cache-max-kb", 256,  
  "file-journal-max-kb", 131072  
  "journal-mb" : 2048,  
  "journal-ino" : 65536,  
  "clean-threshold" : 85.0,  
  "location" : "/var/cache/eos/fusex/cache/",  
  "journal" : "/var/cache/eos/fusex/journal/",  
  "read-ahead-strategy" : "static",  
  "read-ahead-bytes-nominal" : 262144,  
  "read-ahead-bytes-max" : 2097152,  
  "read-ahead-blocks-max" : 16,  
  "max-read-ahead-buffer" : 134217728,  
  "max-write-buffer" : 134217728  
}
```

- There are **two caches** which are persisted while a mount is still up
 - **journal cache**

eosxd

Client side Disk Caching

File & Journal Cache

```
"cache" : {  
  "type" : "disk",  
  "size-mb" : 512,  
  "size-ino" : 65536,  
  "file-cache-max-kb", 256,  
  "file-journal-max-kb", 131072  
  "journal-mb" : 2048,  
  "journal-ino" : 65536,  
  "clean-threshold" : 85.0,  
  "location" : "/var/cache/eos/fusex/cache/",  
  "journal" : "/var/cache/eos/fusex/journal/",  
  "read-ahead-strategy" : "static",  
  "read-ahead-bytes-nominal" : 262144,  
  "read-ahead-bytes-max" : 2097152,  
  "read-ahead-blocks-max" : 16,  
  "max-read-ahead-buffer" : 134217728,  
  "max-write-buffer" : 134217728  
}
```

- There are **two caches** which are persisted while a mount is still up
 - **journal cache**
 - journaling writes
(and also reads if a file has `attr:sys.file.cache=1`)
with 128MB max. size per file

eosxd

Client side Disk Caching

File & Journal Cache

```
"cache" : {  
  "type" : "disk",  
  "size-mb" : 512,  
  "size-ino" : 65536,  
  "file-cache-max-kb", 256,  
  "file-journal-max-kb", 131072  
  "journal-mb" : 2048,  
  "journal-ino" : 65536,  
  "clean-threshold" : 85.0,  
  "location" : "/var/cache/eos/fusex/cache/",  
  "journal" : "/var/cache/eos/fusex/journal/",  
  "read-ahead-strategy" : "static",  
  "read-ahead-bytes-nominal" : 262144,  
  "read-ahead-bytes-max" : 2097152,  
  "read-ahead-blocks-max" : 16,  
  "max-read-ahead-buffer" : 134217728,  
  "max-write-buffer" : 134217728  
}
```

- There are **two caches** which are persisted while a mount is still up
 - **journal cache**
 - journalling writes
(and also reads if a file has `attr:sys.file.cache=1`)
with 128MB max. size per file
 - is used to recover writes

eosxd

Client side Disk Caching

File & Journal Cache

```
"cache" : {  
  "type" : "disk",  
  "size-mb" : 512,  
  "size-ino" : 65536,  
  "file-cache-max-kb", 256,  
  "file-journal-max-kb", 131072  
  "journal-mb" : 2048,  
  "journal-ino" : 65536,  
  "clean-threshold" : 85.0,  
  "location" : "/var/cache/eos/fusex/cache/",  
  "journal" : "/var/cache/eos/fusex/journal/",  
  "read-ahead-strategy" : "static",  
  "read-ahead-bytes-nominal" : 262144,  
  "read-ahead-bytes-max" : 2097152,  
  "read-ahead-blocks-max" : 16,  
  "max-read-ahead-buffer" : 134217728,  
  "max-write-buffer" : 134217728  
}
```

- There are **two caches** which are persisted while a mount is still up
 - **journal cache**
 - journalling writes
(and also reads if a file has `attr:sys.file.cache=1`)
with 128MB max. size per file
 - is used to recover writes
 - **file (start) cache**

eosxd

Client side Disk Caching

File & Journal Cache

```
"cache" : {  
  "type" : "disk",  
  "size-mb" : 512,  
  "size-ino" : 65536,  
  "file-cache-max-kb", 256,  
  "file-journal-max-kb", 131072  
  "journal-mb" : 2048,  
  "journal-ino" : 65536,  
  "clean-threshold" : 85.0,  
  "location" : "/var/cache/eos/fusex/cache/",  
  "journal" : "/var/cache/eos/fusex/journal/",  
  "read-ahead-strategy" : "static",  
  "read-ahead-bytes-nominal" : 262144,  
  "read-ahead-bytes-max" : 2097152,  
  "read-ahead-blocks-max" : 16,  
  "max-read-ahead-buffer" : 134217728,  
  "max-write-buffer" : 134217728  
}
```

- There are **two caches** which are persisted while a mount is still up
 - **journal cache**
 - journalling writes
(and also reads if a file has `attr:sys.file.cache=1`)
with 128MB max. size per file
 - is used to recover writes
 - **file (start) cache**
 - caches the first 256kb of each file

eosxd

Client side Disk Caching

File & Journal Cache

```
"cache" : {  
  "type" : "disk",  
  "size-mb" : 512,  
  "size-ino" : 65536,  
  "file-cache-max-kb", 256,  
  "file-journal-max-kb", 131072  
  "journal-mb" : 2048,  
  "journal-ino" : 65536,  
  "clean-threshold" : 85.0,  
  "location" : "/var/cache/eos/fusex/cache/",  
  "journal" : "/var/cache/eos/fusex/journal/",  
  "read-ahead-strategy" : "static",  
  "read-ahead-bytes-nominal" : 262144,  
  "read-ahead-bytes-max" : 2097152,  
  "read-ahead-blocks-max" : 16,  
  "max-read-ahead-buffer" : 134217728,  
  "max-write-buffer" : 134217728  
}
```

- There are **two caches** which are persisted while a mount is still up
 - **journal cache**
 - journalling writes
(and also reads if a file has `attr:sys.file.cache=1`)
with 128MB max. size per file
 - is used to recover writes
 - **file (start) cache**
 - caches the first 256kb of each file
 - triggered when a file is opened

eosxd

Virtual/Functional Attributes



Information/Getter

Information/Getter



Information/Getter

[illegible]



Information/Getter

[illegible]

Virtual/Functional Attributes

Information/Getter

[illegible]

Information/Getter

[illegible]

Information/Getter

[illegible]

Information/Getter

[illegible]

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	number of files here
eos.name	EOS instance name

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	humber of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	number of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname
eos.mgmurl	EOS instance URL root://...

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	humber of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname
eos.mgmurl	EOS instance URL root://...
eos.url.xroot	EOS XRootD TURL of the given file

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	number of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname
eos.mgmurl	EOS instance URL root://...
eos.url.xroot	EOS XRootD TURL of the given file
eos.stats	get statistics output file

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	number of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname
eos.mgmurl	EOS instance URL root://...
eos.url.xroot	EOS XRootD TURL of the given file
eos.stats	get statistics output file
eos.quota	get quota output for the calling user

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	number of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname
eos.mgmurl	EOS instance URL root://...
eos.url.xroot	EOS XRootD TURL of the given file
eos.stats	get statistics output file
eos.quota	get quota output for the calling user

Authentication IF/Getter

Virtual Ext. Attribute	Meaning

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	number of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname
eos.mgmurl	EOS instance URL root://...
eos.url.xroot	EOS XRootD TURL of the given file
eos.stats	get statistics output file
eos.quota	get quota output for the calling user

Authentication IF/Getter

Virtual Ext. Attribute	Meaning
eos.identity	get identity of the calling client

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	humber of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname
eos.mgmurl	EOS instance URL root://...
eos.url.xroot	EOS XRootD TURL of the given file
eos.stats	get statistics output file
eos.quota	get quota output for the calling user

Authentication IF/Getter

Virtual Ext. Attribute	Meaning
eos.identity	get identity of the calling client
eos.identityparent	get identity of the calling client for the parent directory

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	humber of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname
eos.mgmurl	EOS instance URL root://...
eos.url.xroot	EOS XRootD TURL of the given file
eos.stats	get statistics output file
eos.quota	get quota output for the calling user

Authentication IF/Getter

Virtual Ext. Attribute	Meaning
eos.identity	get identity of the calling client
eos.identityparent	get identity of the calling client for the parent directory
eos.reconnect	force to drop current connection and create a new one

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	number of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname
eos.mgmurl	EOS instance URL root://...
eos.url.xroot	EOS XRootD TURL of the given file
eos.stats	get statistics output file
eos.quota	get quota output for the calling user

Authentication IF/Getter

Virtual Ext. Attribute	Meaning
eos.identity	get identity of the calling client
eos.identityparent	get identity of the calling client for the parent directory
eos.reconnect	force to drop current connection and create a new one
eos.reconnectparent	as above but for the parent



Information/Getter

Authentication IF/Getter

Modifications / Setter

[illegible]

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	number of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname
eos.mgmurl	EOS instance URL root://...
eos.url.xroot	EOS XRootD TURL of the given file
eos.stats	get statistics output file
eos.quota	get quota output for the calling user

Authentication IF/Getter

Virtual Ext. Attribute	Meaning
eos.identity	get identity of the calling client
eos.identityparent	get identity of the calling client for the parent directory
eos.reconnect	force to drop current connection and create a new one
eos.reconnectparent	as above but for the parent

Modifications / Setter

Virtual Ext. Attribute	ARG	Meaning
system.eos.debug	<level>	Set debug level

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	number of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname
eos.mgmurl	EOS instance URL root://...
eos.url.xroot	EOS XRootD TURL of the given file
eos.stats	get statistics output file
eos.quota	get quota output for the calling user

Authentication IF/Getter

Virtual Ext. Attribute	Meaning
eos.identity	get identity of the calling client
eos.identityparent	get identity of the calling client for the parent directory
eos.reconnect	force to drop current connection and create a new one
eos.reconnectparent	as above but for the parent

Modifications / Setter

Virtual Ext. Attribute	ARG	Meaning
system.eos.debug	<level>	Set debug level
system.eos.dropcap	-	Drop subscription here

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	number of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname
eos.mgmurl	EOS instance URL root://...
eos.url.xroot	EOS XRootD TURL of the given file
eos.stats	get statistics output file
eos.quota	get quota output for the calling user

Authentication IF/Getter

Virtual Ext. Attribute	Meaning
eos.identity	get identity of the calling client
eos.identityparent	get identity of the calling client for the parent directory
eos.reconnect	force to drop current connection and create a new one
eos.reconnectparent	as above but for the parent

Modifications / Setter

Virtual Ext. Attribute	ARG	Meaning
system.eos.debug	<level>	Set debug level
system.eos.dropcap	-	Drop subscription here
system.eos.dropcaps	-	Drop all subscriptions

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	number of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname
eos.mgmurl	EOS instance URL root://...
eos.url.xroot	EOS XRootD TURL of the given file
eos.stats	get statistics output file
eos.quota	get quota output for the calling user

Authentication IF/Getter

Virtual Ext. Attribute	Meaning
eos.identity	get identity of the calling client
eos.identityparent	get identity of the calling client for the parent directory
eos.reconnect	force to drop current connection and create a new one
eos.reconnectparent	as above but for the parent

Modifications / Setter

Virtual Ext. Attribute	ARG	Meaning
system.eos.debug	<level>	Set debug level
system.eos.dropcap	-	Drop subscription here
system.eos.dropcaps	-	Drop all subscriptions
system.eos.resetstat	-	Reset Stat Counter

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	number of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname
eos.mgmurl	EOS instance URL root://...
eos.url.xroot	EOS XRootD TURL of the given file
eos.stats	get statistics output file
eos.quota	get quota output for the calling user

Authentication IF/Getter

Virtual Ext. Attribute	Meaning
eos.identity	get identity of the calling client
eos.identityparent	get identity of the calling client for the parent directory
eos.reconnect	force to drop current connection and create a new one
eos.reconnectparent	as above but for the parent

Modifications / Setter

Virtual Ext. Attribute	ARG	Meaning
system.eos.debug	<level>	Set debug level
system.eos.dropcap	-	Drop subscription here
system.eos.dropcaps	-	Drop all subscriptions
system.eos.resetstat	-	Reset Stat Counter
system.eos.log	private public	Change Lofile Permissions

eosxd

Virtual/Functional Attributes

Information/Getter

Virtual Ext. Attribute	Meaning
eos.btime	file birth time
eos.ttime	latest modification time in subtree
eos.tsize	total size of directory subtree here
eos.dsize	total size of files in this directory
eos.dcount	number of subdirectories here
eos.fcount	number of files here
eos.name	EOS instance name
eos.hostport	EOS hostname:portname
eos.mgmurl	EOS instance URL root://...
eos.url.xroot	EOS XRootD TURL of the given file
eos.stats	get statistics output file
eos.quota	get quota output for the calling user

Authentication IF/Getter

Virtual Ext. Attribute	Meaning
eos.identity	get identity of the calling client
eos.identityparent	get identity of the calling client for the parent directory
eos.reconnect	force to drop current connection and create a new one
eos.reconnectparent	as above but for the parent

Modifications / Setter

Virtual Ext. Attribute	ARG	Meaning
system.eos.debug	<level>	Set debug level
system.eos.dropcap	-	Drop subscription here
system.eos.dropcaps	-	Drop all subscriptions
system.eos.resetstat	-	Reset Stat Counter
system.eos.log	private public	Change Lofile Permissions
system.eos.fuzz	all/config	Toggle Fuzzing mode

eosxd vs eosxd3 changes

eosxd vs eosxd3 changes

- readdirplus

eosxd vs eosxd3 changes

- **readdirplus**

- the function returning a directory listing attaches the stat information for each listed entry

eosxd vs eosxd3 changes

- **readdirplus**

- the function returning a directory listing attaches the stat information for each listed entry
 - avoids additional **getattr** calls for each listed child!

eosxd vs eosxd3 changes

- **readdirplus**

- the function returning a directory listing attaches the stat information for each listed entry
 - avoids additional **getattr** calls for each listed child!

- **forgetmulti**

eosxd vs eosxd3 changes

- **readdirplus**

- the function returning a directory listing attaches the stat information for each listed entry
 - avoids additional **getattr** calls for each listed child!

- **forgetmulti**

- when the kernel wants to forget inodes, instead of recalling one by one a bulk request reduces the number of **forget** calls substantially

eosxd vs eosxd3 changes

- **readdirplus**

- the function returning a directory listing attaches the stat information for each listed entry
 - avoids additional **getattr** calls for each listed child!

- **forgetmulti**

- when the kernel wants to forget inodes, instead of recalling one by one a bulk request reduces the number of **forget** calls substantially

- **clone_fd** option

eosxd vs eosxd3 changes

- **readdirplus**

- the function returning a directory listing attaches the stat information for each listed entry
 - avoids additional **getattr** calls for each listed child!

- **forgetmulti**

- when the kernel wants to forget inodes, instead of recalling one by one a bulk request reduces the number of **forget** calls substantially

- **clone_fd** option

- each FUSE thread can have its own connection to the kernel module via /dev/fuse

eosxd vs eosxd3 changes

- **readdirplus**

- the function returning a directory listing attaches the stat information for each listed entry
 - avoids additional **getattr** calls for each listed child!

- **forgetmulti**

- when the kernel wants to forget inodes, instead of recalling one by one a bulk request reduces the number of **forget** calls substantially

- **clone_fd** option

- each FUSE thread can have its own connection to the kernel module via /dev/fuse
 - in libfuse2 this is a single shared connection

eosxd vs eosxd3 changes

- **readdirplus**

- the function returning a directory listing attaches the stat information for each listed entry
 - avoids additional **getattr** calls for each listed child!

- **forgetmulti**

- when the kernel wants to forget inodes, instead of recalling one by one a bulk request reduces the number of **forget** calls substantially

- **clone_fd** option

- each FUSE thread can have its own connection to the kernel module via /dev/fuse
 - in libfuse2 this is a single shared connection

- **write_back_cache** option

eosxd vs eosxd3 changes

- **readdirplus**

- the function returning a directory listing attaches the stat information for each listed entry
 - avoids additional **getattr** calls for each listed child!

- **forgetmulti**

- when the kernel wants to forget inodes, instead of recalling one by one a bulk request reduces the number of **forget** calls substantially

- **clone_fd** option

- each FUSE thread can have its own connection to the kernel module via /dev/fuse
 - in libfuse2 this is a single shared connection

- **write_back_cache** option

- small writes are not hitting directly eosxd, but they are written into the buffer cache and then flushed as larger operations

eosxd vs eosxd3 changes

- **readdirplus**

- the function returning a directory listing attaches the stat information for each listed entry
 - avoids additional **getattr** calls for each listed child!

- **forgetmulti**

- when the kernel wants to forget inodes, instead of recalling one by one a bulk request reduces the number of **forget** calls substantially

- **clone_fd** option

- each FUSE thread can have its own connection to the kernel module via /dev/fuse
 - in libfuse2 this is a single shared connection

- **write_back_cache** option

- small writes are not hitting directly eosxd, but they are written into the buffer cache and then flushed as larger operations
- increases single byte write performance from 16kHz to 0.5MHz

eosxd vs eosxd3 changes

- **readdirplus**

- the function returning a directory listing attaches the stat information for each listed entry
 - avoids additional **getattr** calls for each listed child!

- **forgetmulti**

- when the kernel wants to forget inodes, instead of recalling one by one a bulk request reduces the number of **forget** calls substantially

- **clone_fd** option

- each FUSE thread can have its own connection to the kernel module via /dev/fuse
 - in libfuse2 this is a single shared connection

- **write_back_cache** option

- small writes are not hitting directly eosxd, but they are written into the buffer cache and then flushed as larger operations
- increases single byte write performance from 16kHz to 0.5MHz

- **FUSE_SET_ATTR_CTIME**

eosxd vs eosxd3 changes

- **readdirplus**

- the function returning a directory listing attaches the stat information for each listed entry
 - avoids additional **getattr** calls for each listed child!

- **forgetmulti**

- when the kernel wants to forget inodes, instead of recalling one by one a bulk request reduces the number of **forget** calls substantially

- **clone_fd** option

- each FUSE thread can have its own connection to the kernel module via /dev/fuse
 - in libfuse2 this is a single shared connection

- **write_back_cache** option

- small writes are not hitting directly eosxd, but they are written into the buffer cache and then flushed as larger operations
- increases single byte write performance from 16kHz to 0.5MHz

- **FUSE_SET_ATTR_CTIME**

- allows to set the change time with the **utimes** function

eosxd vs eosxd3 starting



eosxd vs eosxd3 starting



eosxd starts the libfuse2 version

eosxd vs eosxd3 starting



eosxd starts the libfuse2 version
eosxd3 starts the libfuse3 version

eosxd vs eosxd3 starting

eosxd starts the **libfuse2** version

eosxd3 starts the **libfuse3** version

- If you define in `/etc/sysconfig/eosxd`: `EOSXD="/usr/sbin/eosxd3"` and you use the generic mount helper **mount.eosx** ...

eosxd vs eosxd3 starting

eosxd starts the **libfuse2** version

eosxd3 starts the **libfuse3** version

- If you define in `/etc/sysconfig/eosxd`: `EOSXD="/usr/sbin/eosxd3"` and you use the generic mount helper **mount.eosx** ...
 - ... you can switch changing the environment variable between both easily to compare:

eosxd vs eosxd3

starting

eosxd starts the **libfuse2** version

eosxd3 starts the **libfuse3** version

- If you define in `/etc/sysconfig/eosxd`: `EOSXD="/usr/sbin/eosxd3"` and you use the generic mount helper **mount.eosx** ...
 - ... you can switch changing the environment variable between both easily to compare:
 - e.g. by using

```
mount -t eosx eosxd -ofsname=eosinstance.cern.ch:/eos/ /eos/
```


eosxd vs eosxd3

starting

eosxd starts the **libfuse2** version

eosxd3 starts the **libfuse3** version

- If you define in `/etc/sysconfig/eosxd`: `EOSXD="/usr/sbin/eosxd3"` and you use the generic mount helper **mount.eosx** ...

- ... you can switch changing the environment variable between both easily to compare:

- e.g. by using

```
mount -t eosx eosxd -ofsname=eosinstance.cern.ch:/eos/ /eos/
```

- and not

```
mount -t fuse eosxd -ofsname=eosinstance.cern.ch:/eos/ /eos/
```


eosxd3

Performance

- 100 GE client/server
idle instance test

eosxd3

Performance

- 100 GE client/server
idle instance test

eosxd

eosxd3

Performance


- 100 GE client/server
idle instance test



eosxd3

Performance


- 100 GE client/server
idle instance test



eosxd3

Performance

- 100 GE client/server idle instance test




Single Client	eosxd	eosxd3

eosxd3

Performance

- 100 GE client/server
idle instance test




Single Client	eosxd	eosxd3
Seq.Creation		700 Hz

eosxd3

Performance

- 100 GE client/server
idle instance test




Single Client	eosxd	eosxd3
Seq.Creation		700 Hz
Par. Creation		1000 Hz

eosxd3

Performance

- 100 GE client/server
idle instance test




Single Client	eosxd	eosxd3
Seq.Creation		700 Hz
Par. Creation		1000 Hz
w IOPS	16 kHz	500 kHz

eosxd3

Performance

- 100 GE client/server
idle instance test




Single Client	eosxd	eosxd3
Seq.Creation		700 Hz
Par. Creation		1000 Hz
w IOPS	16 kHz	500 kHz
Seq. Read		1.6 GB/s

eosxd3

Performance

- 100 GE client/server
idle instance test



Single Client	eosxd	eosxd3
Seq. Creation		700 Hz
Par. Creation		1000 Hz
w IOPS	16 kHz	500 kHz
Seq. Read		1.6 GB/s
Seq. Write		900 MB/s



Summary





Summary

- Easy FUSE access is **appreciated by users** (when it is reliable) and allows EOS usage for many more applications



Summary

- Easy FUSE access is **appreciated by users** (when it is reliable) and allows EOS usage for many more applications
- Repeated re-implementation and updates have **improved the reliability and usability** of /eos/ access steadily over the past years



Summary

- Easy FUSE access is **appreciated by users** (when it is reliable) and allows EOS usage for many more applications
- Repeated re-implementation and updates have **improved the reliability and usability** of /eos/ access steadily over the past years
- **libfuse3** brings a visible **performance improvement** for shared mounts



Summary

- Easy FUSE access is **appreciated by users** (when it is reliable) and allows EOS usage for many more applications
- Repeated re-implementation and updates have **improved the reliability and usability** of /eos/ access steadily over the past years
- **libfuse3** brings a visible **performance improvement** for shared mounts
- One of the pain points for performance is missing kernel support to cache extended attributes - otherwise most limitations are server-side



Summary

- Easy FUSE access is **appreciated by users** (when it is reliable) and allows EOS usage for many more applications
- Repeated re-implementation and updates have **improved the reliability and usability** of /eos/ access steadily over the past years
- **libfuse3** brings a visible **performance improvement** for shared mounts
- One of the pain points for performance is missing kernel support to cache extended attributes - otherwise most limitations are server-side
- Nevertheless there are still corners client-side to **improve FUSE performance** - which we will also try to exploit!



FO



SXD

BEST OF FUSE

The End!

Questions / Comments ?

<https://indico.cern.ch/event/1227241/>



EOS 2023 Workshop

24–27 Apr 2023
CERN
Europe/Zurich timezone

Enter your search term



Overview

Scientific Programme

Call for Abstracts

Timetable

Contribution List

My Conference

My Contributions

Book of Abstracts

Registration

Participant List

Privacy Information

Videoconference



Videoconference