

# MOTR API

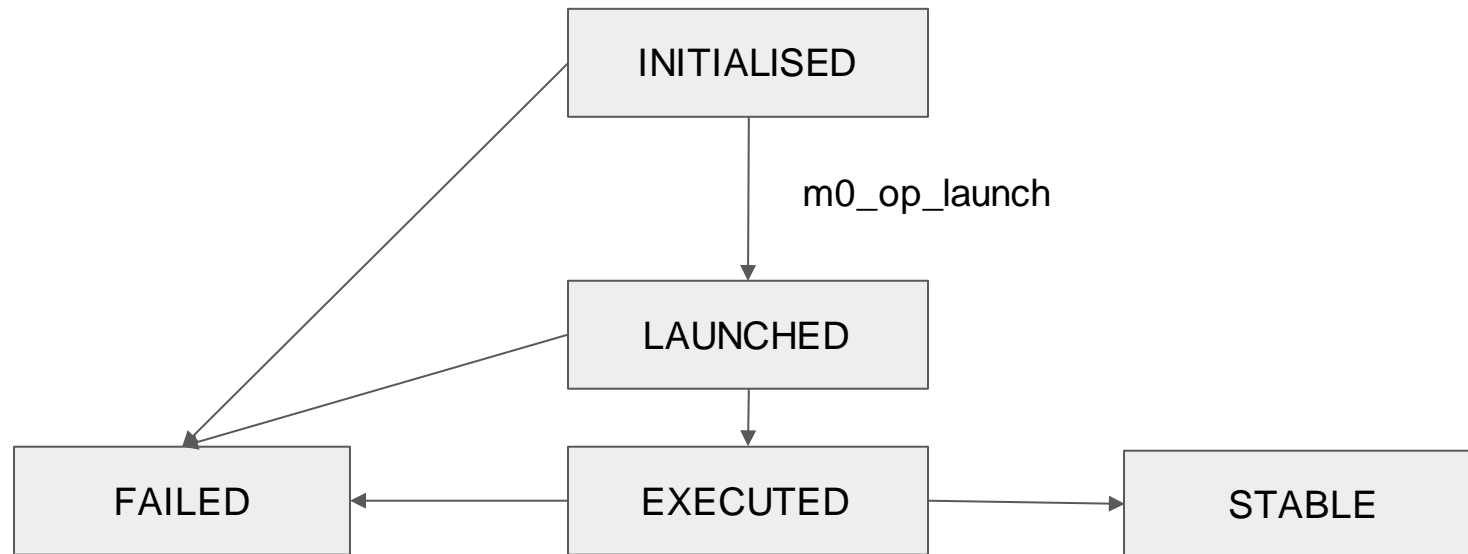


# Motr Access Interface Overview

- Operation
  - asynchronous
- Object
  - 128-bit persistent identifiers, assigned by user
  - Create/Delete an Object.
  - Read/Write/Sync an Object.
- Index: distributed key-value store
  - 128-bit persistent identifiers, assigned by user
  - Create/Delete and Lookup an Index.
  - GET/PUT/DEL/NEXT key-value pairs of an Index.
- Others such as APIs for composite objects (useful for applications such as HSM)

# Part I: Operation

- A state machine that represents an ongoing request to Motr
- Asynchronous



<b>Operation</b>	m0_op_launch()  m0_op_wait()  m0_op_setup()  m0_op_cancel()  m0_op_fini/free()	Launch an operation  Wait on an operation  Set callback functions for an operation  Cancel operations  Finalise/free an operation
------------------	--	---

<p><b>Entity Operations</b> (for both object and index)</p>	<p>M0_EO_CREATE M0_EO_DELETE M0_EO_OPEN</p> <p>M0_EO_SYNC</p> <p>M0_EO_GETATTR M0_EO_SETATTR</p> <p>M0_EO_LAYOUT_GET M0_EO_LAYOUT_SET</p>	<p>Create/delete an entity</p> <p>Open an entity</p> <p>Flush entity data to storage devices.</p> <p>For Motr internal uses only</p> <p>For composite layout only</p>
<p><b>Object Operations</b></p>	<p>M0_OC_READ M0_OC_WRITE</p>	<p>IOs in scatter-gather-scatter fashion</p>
<p><b>Index Operations</b></p>	<p>M0_IC_GET M0_IC_PUT M0_IC_DEL M0_IC_NEXT</p> <p>M0_IC_LOOKUP M0_IC_LIST</p>	<p>Index queries, including GET, PUT, DEL and NEXT</p> <p>Check an index for an existence Given an index id, get the list of next indices</p>

## Part II: Object

- An object is viewed as an array of data-blocks,  $[0, 2^{64})$ , initially a hole
  - Flat name space
  - Network striped (Parity declustered data layout)
- No usual metadata (attributes, etc.) exposed to users, such as object size
  - Maximize IO performance by avoiding updating object attributes.
  - Mero has internal object attributes such as data layout.
- CREATE, DELETE, READ, WRITE, ALLOC, FREE operations
- scatter-gather-scatter READ and WRITE operations

<h2>Object Access</h2>	<p> m0_obj_init/fini()  m0_entity_create()  m0_entity_delete()  m0_entity_open()    m0_obj_op()    m0_sync_op_init()  m0_sync_entity_add()  m0_sync_op_add() </p>	<p> Initialise/finalise in-memory object data structure.  Create and initialise operation for object creation and deletion.  Once opened, the object data structure can be used in later READ/WRITE ops. </p> <p> Create and initialise an object operation specified by the opcode </p> <p> Create and initialise an SYNC operation.  Add an entity to SYNC operation  Add an `op` to SYNC operation. </p>
<h2>Object Lock</h2>	<p> m0_obj_lock_init()  m0_obj_lock_fini()  m0_obj_lock_get_sync()  m0_obj_lock_get()  m0_obj_lock_put() </p>	<p> Currently Motr only supports exclusive “whole object” lock in a group. </p>

## Part III: Index

- Distributed key-value store
- Key and value
  - Keys and values are bit-strings
  - Values accessed as a whole
  - Iteration in the order of keys
- Operations
  - CREATE, DELETE, LOOKUP and LIST index
  - GET, PUT, DEL and NEXT for index query
  - Scatter-gather-scatter operations
- User uses indices to build meta-data structures: namespace, file attributes



# Motr Client Index APIs

<b>Index (Key-value) Access</b>	<p>m0_idx_init/fini()</p> <p>m0_entity_create() m0_entity_delete()</p> <p>m0_idx_op()</p> <p>m0_sync_op_init m0_sync_entity_add m0_sync_op_add</p>	<p>Initialise/finalise in-memory object data structure</p> <p>Create and initialise index creation/deletion operations.</p> <p>Create and initialise index query operation. Clovis support GET, PUT, DEL and NEXT queries.</p> <p>Flush index key and value to persistent storage devices.</p>
---------------------------------	--	--

# Motr Access Interface Specifications

<https://github.com/Seagate/cortex-motr/blob/main/motr/client.h>

## Examples

<https://github.com/Seagate/cortex-motr-apps>

<https://github.com/Seagate/cortex-motr/tree/main/hsm>

<https://github.com/Seagate/cortex-mio>

# Questions?



# Motr Internal Architecture

