

Secure, API based Control Plane for Storage Scale

Luis Teran

<luis.i.teran@ibm.com>



Disclaimer

- The information in this document is **IBM CONFIDENTIAL**.
- **USE OF THIS DOCUMENT IS LIMITED TO SELECT IBM PERSONNEL AND TO BUSINESS PARTNERS WHO HAVE A CURRENT SIGNED NONDISCLOSURE AGREEMENT ON FILE WITH IBM. THIS INFORMATION CAN ALSO BE SHARED WITH CUSTOMERS WHO HAVE A CURRENT SIGNED NONDISCLOSURE AGREEMENT (NDA) ON FILE WITH IBM, BUT THIS DOCUMENT SHOULD NOT BE GIVEN TO A CUSTOMER EITHER IN HARDCOPY OR ELECTRONIC FORMAT.**
- This information is provided on an "AS IS" basis without warranty of any kind, express or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Some jurisdictions do not allow disclaimers of express or implied warranties in certain transactions; therefore, this statement may not apply to you.
- This information is provided for information purposes only as a high-level overview of possible future products. PRODUCT SPECIFICATIONS, ANNOUNCE DATES, AND OTHER INFORMATION CONTAINED HEREIN ARE SUBJECT TO CHANGE AND WITHDRAWAL WITHOUT NOTICE.
- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.
- IBM reserves the right to change product specifications and offerings at any time without notice. This publication could include technical inaccuracies or typographical errors. References herein to IBM products and services do not imply that IBM intends to make them available in all countries.

It's Here!

- GA'd on May 9th!
- Tech Preview since 5.1.9.0



Security Improvements

Removal of SSH /
SCP dependency

Removal of root
requirement for
control plane

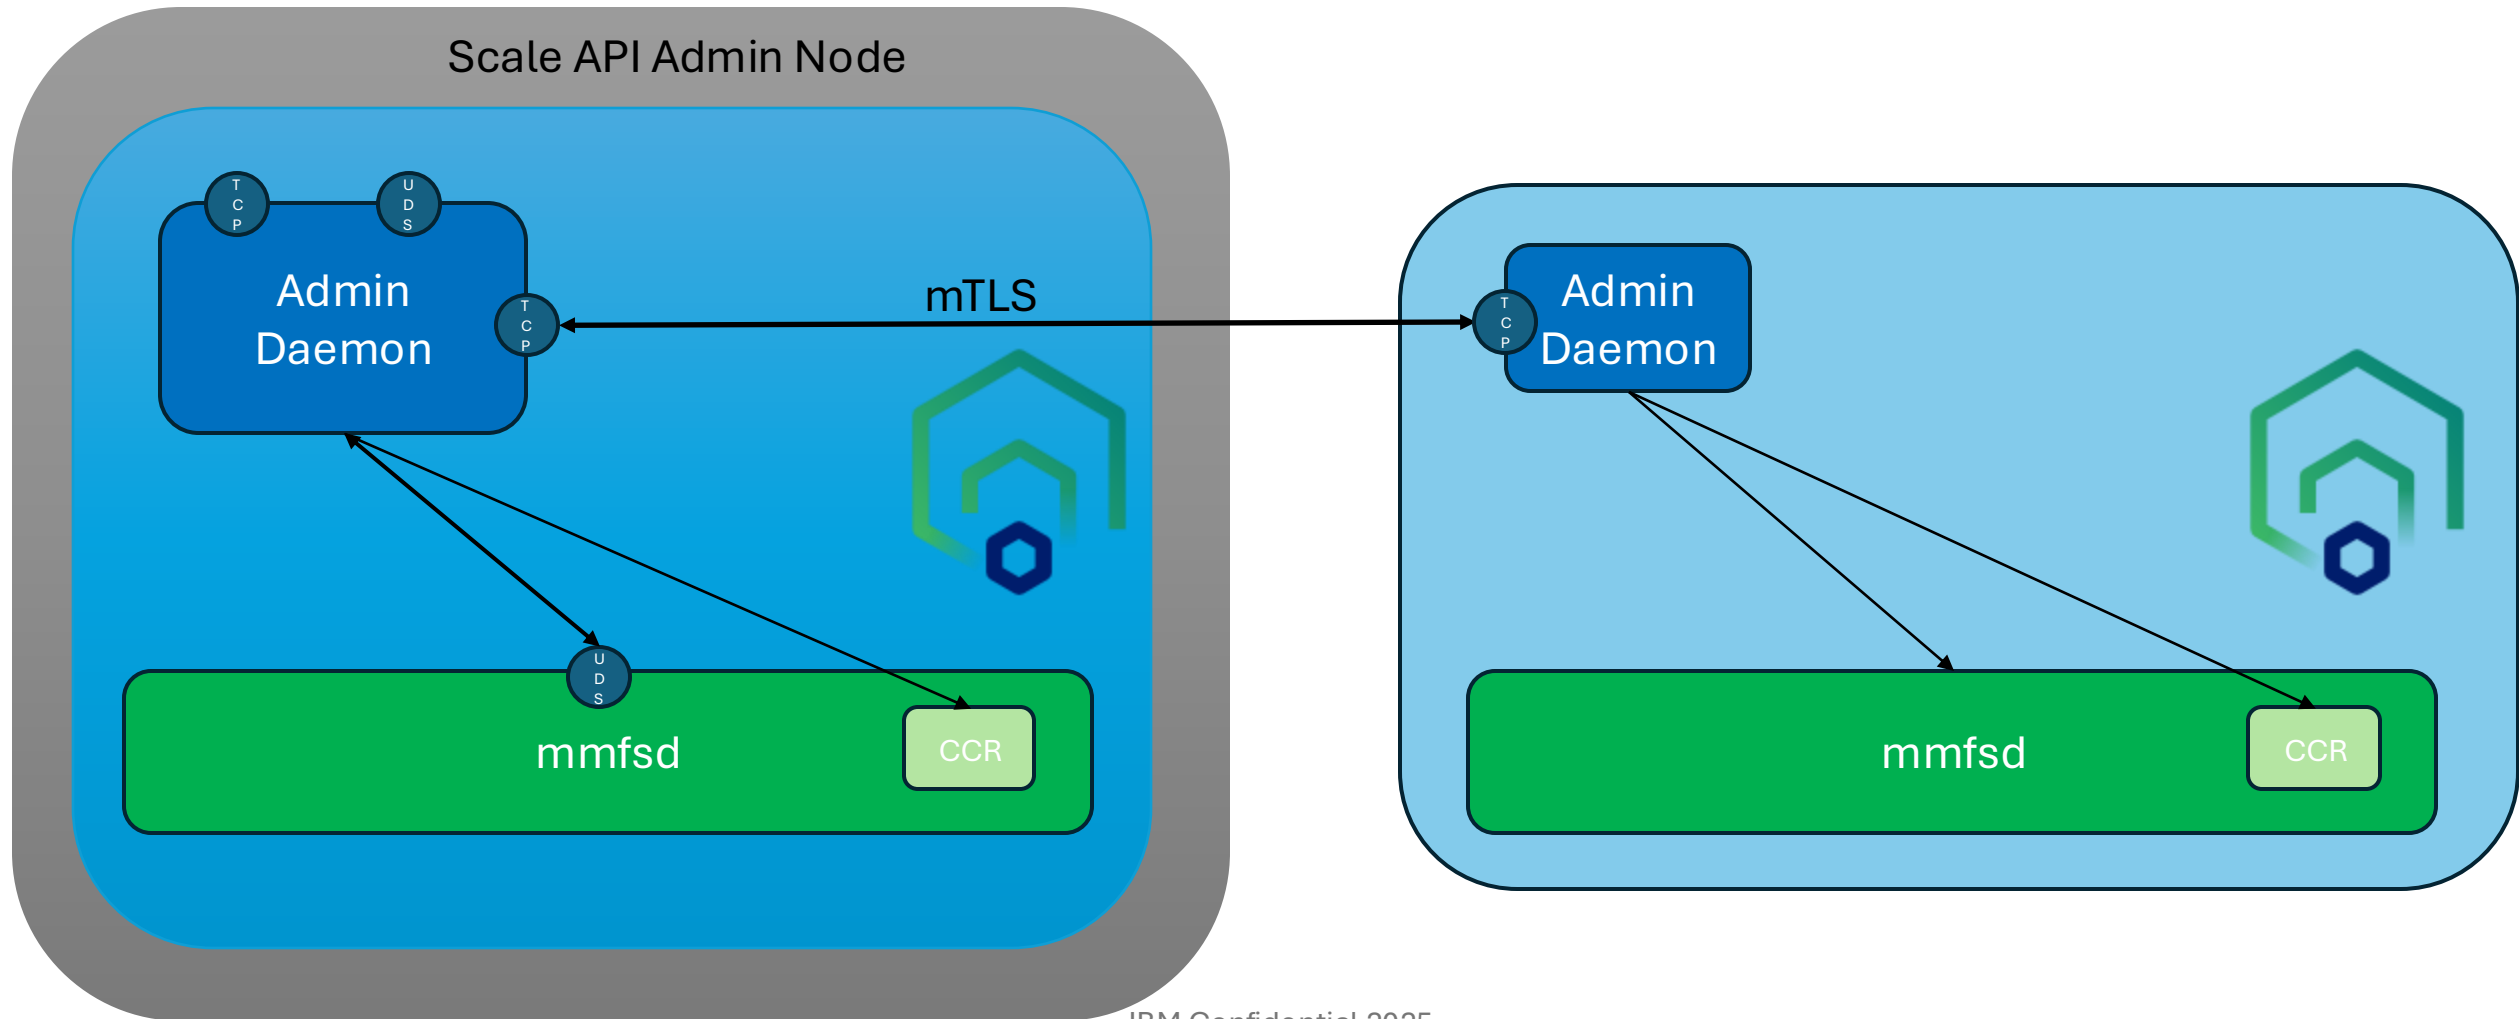


Remote
Administration

Fine-Grained
Access Control

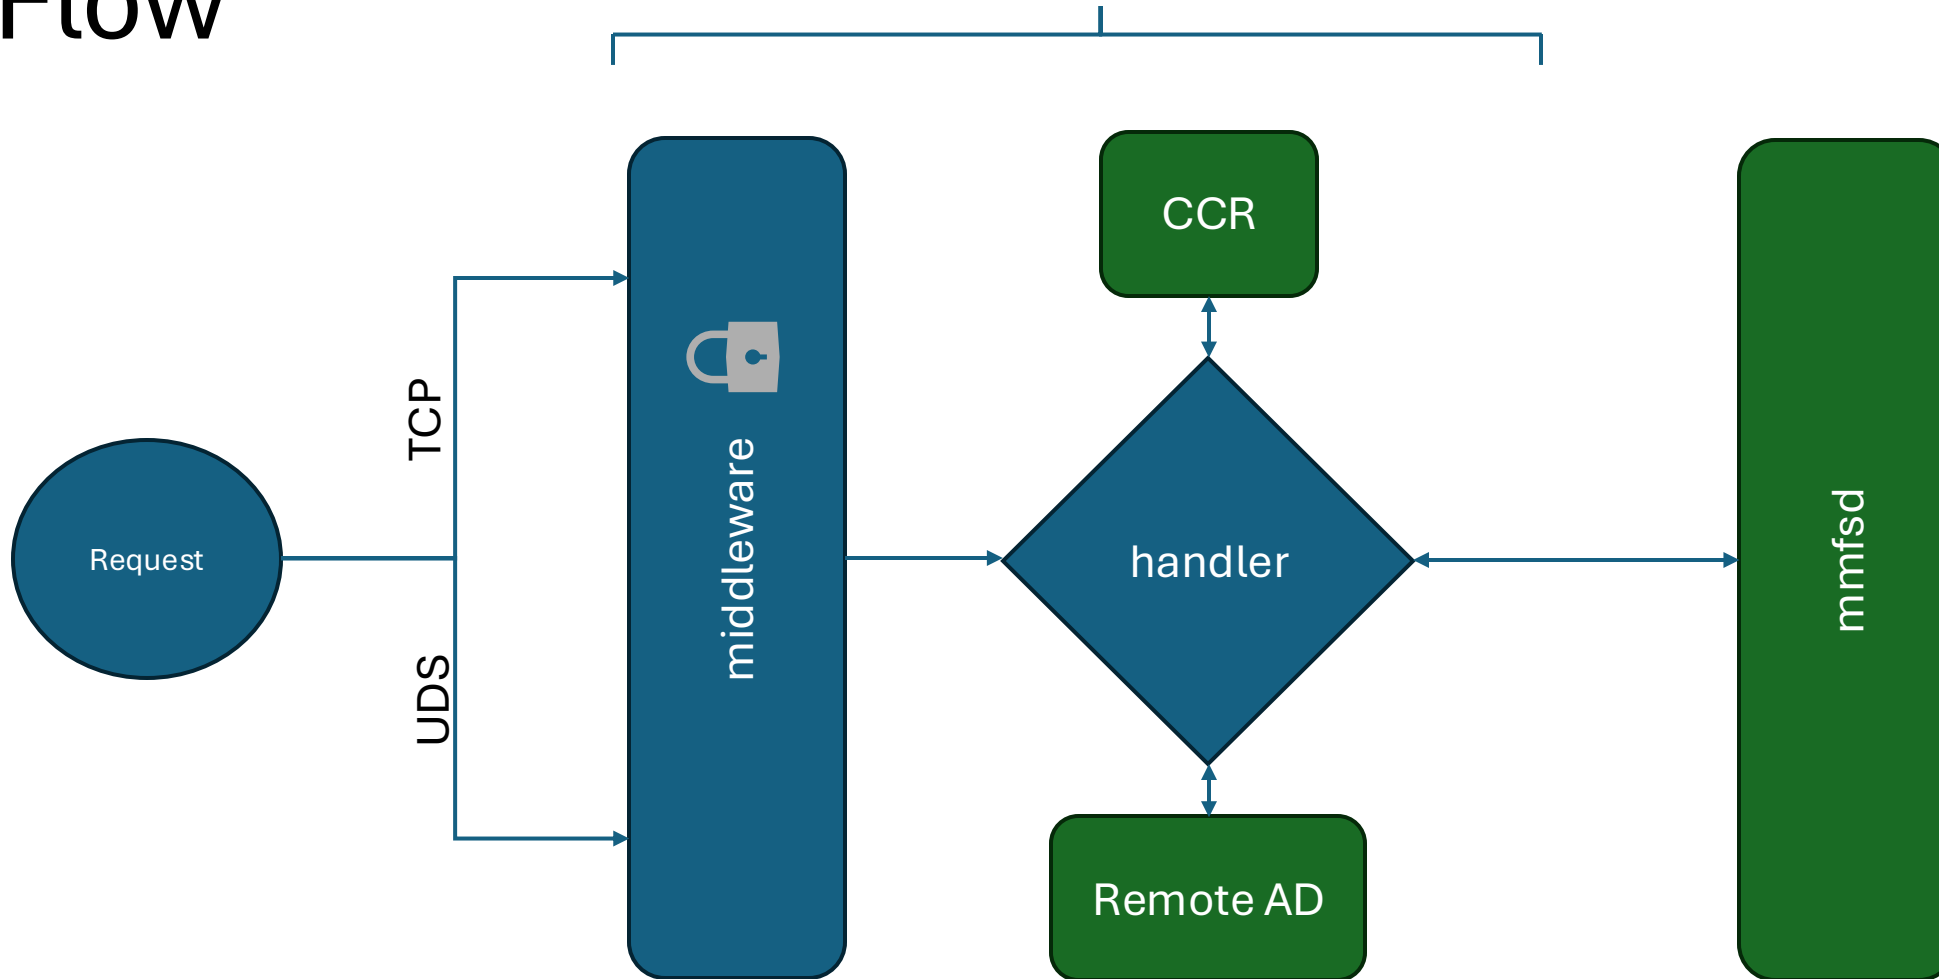


Simplified High Level Overview



Flow

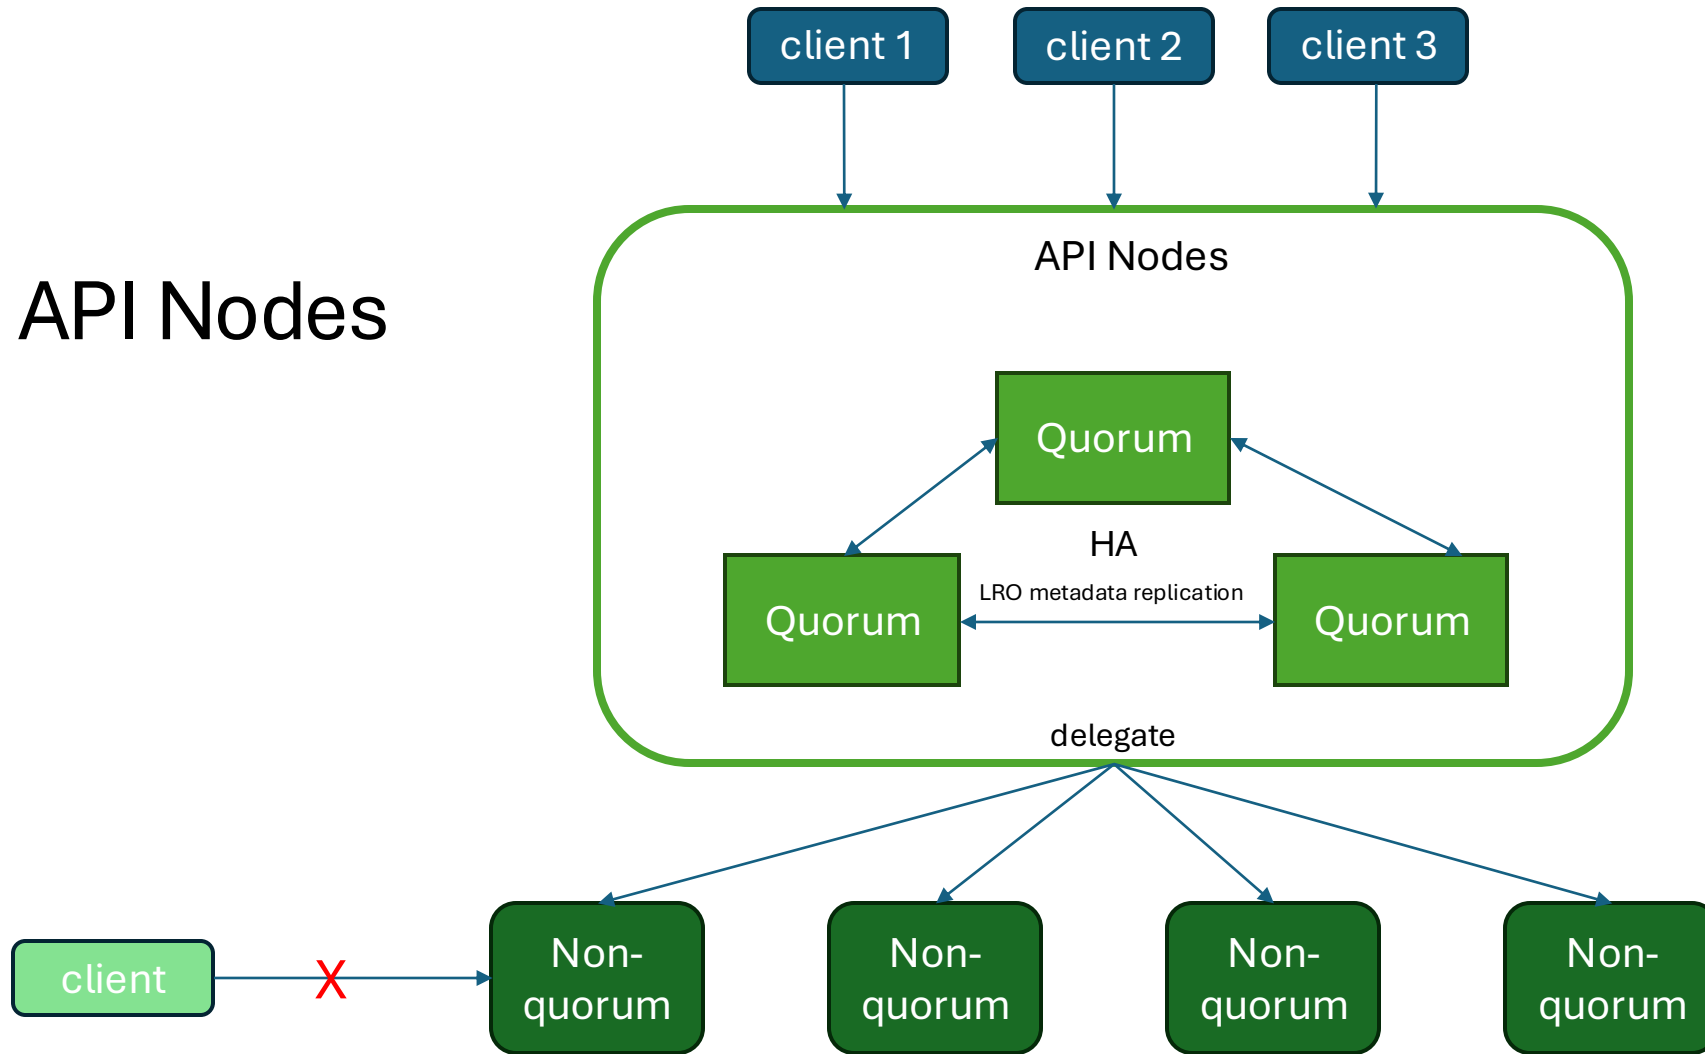
Admin Daemon



Request:
scalectl command
curl
Client Application
...

Middleware:
Authentication
Authorization
Request logging

API Nodes



New CLI

Storage Scale Admin CLI interface

Usage:

```
scalectl [flags]  
scalectl [command]
```

Available Commands:

apihealth	apihealth commands
authorization	authorization commands
cluster	cluster commands
config	cluster configuration commands
diagnostics	node diagnostics commands
fileset	fileset commands
filesystem	file system commands
node	node commands
nodeclass	node class commands
nsd	NSD commands
operation	operation commands
policy	policy commands
troubleshooting	trouble shooting commands

5.2.3.0 GA for Scale-API

Cluster Management

Cluster Create / Node add / Node list
Configuration management
Nodeclass management

Node Management

Start / Stop (mmstartup/mmshutdown)
Status (mmgetstate)

Misc

SSH + SCP Replacement
Non-Root Admin Daemon via privilege escalation (linux capabilities)
Health monitoring of admin daemons

File System Management

CRUD
Mount / Unmount
Mount status
Policy (mmlspolicy / mmchpolicy)
Snapshots
Storage pools
Disk management
Fileset management
FS Manager

RBAC Management

Authorization Management

Authentication

PAM (default)
OIDC
LDAP/AD
Certificates

NSD Management

CRUD

Long Running Operations

Long Running Operations management
Cancellation

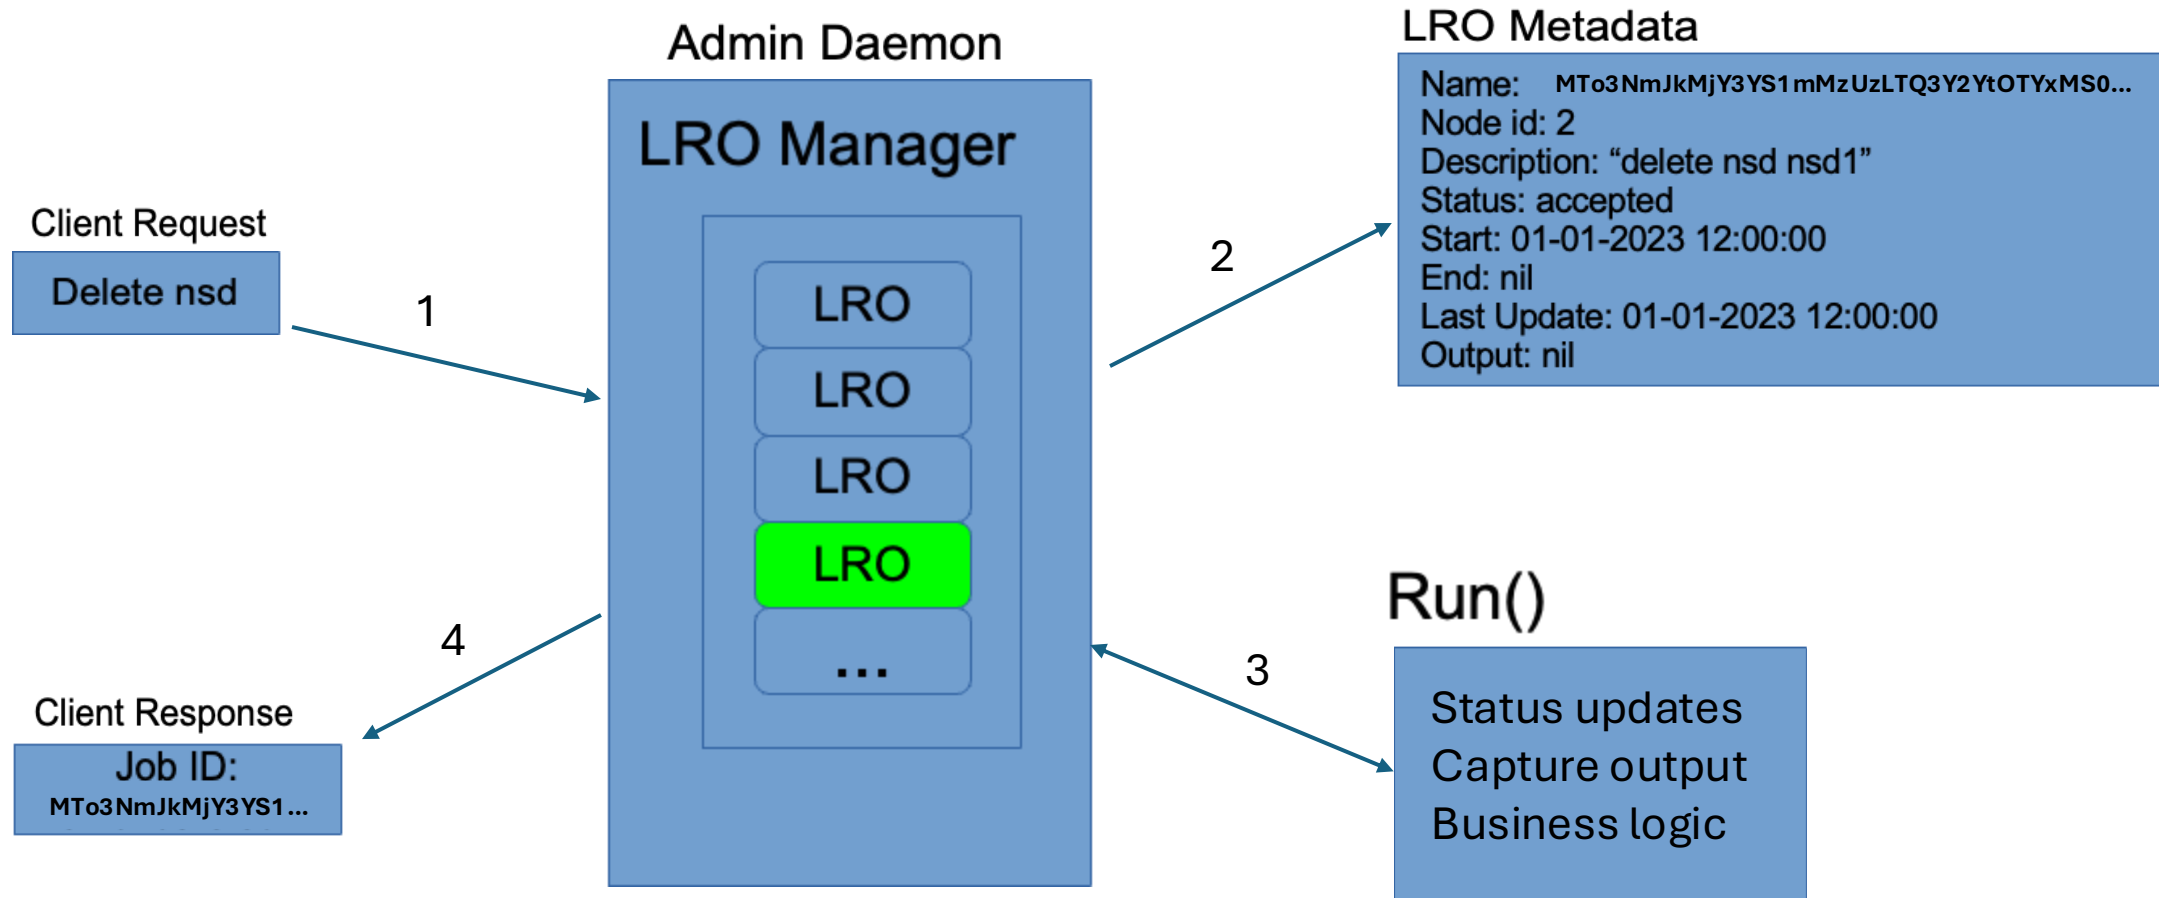
Remote Clusters

Remote Cluster Setup
Interoperability with existing functionality

Toolkit Support

Long Running Operations (LRO)

LRO life cycle



Long Running Operations

```
[root@nopriv-41 ~]# scalectl fileset create gpfs0 -n fileset1
Successfully created fileset fileset1 in filesystem gpfs0 with id 2 root inode 35587.
[root@nopriv-41 ~]# scalectl operations list
```

Job Id	Operation	Status	Request Time	Last Update Time	Completion Time
MT03NmJkMjY3YS1mMzUzLTQ3Y2YtOTYxMS0yMjYxODMxZjEwMDE=	CreateFilesetRequest	Done	2025-05-05T22:31:45Z	2025-05-05T22:31:45Z	2025-05-05T22:31:45Z

- Added cancellation support
 - Currently need to do a request to `/scalemgmt/v3/operations/{id}:cancel`
- Planned support for `scalectl`
 - Discussing if `ctrl-c` on `scalectl` will propagate the cancellation to the operation or just for `scalectl`. Open to feedback!

```
[root@nopriv-41 ~]# scalectl operations get MT03NmJkMjY3YS1mMzUzLTQ3Y2YtOTYxMS0yMjYxODMxZjEwMDE=
```

Job Id	MT03NmJkMjY3YS1mMzUzLTQ3Y2YtOTYxMS0yMjYxODMxZjEwMDE=
Operation	CreateFilesetRequest
Status	Done
Request Time	2025-05-05T22:31:45Z
Last Update Time	2025-05-05T22:31:45Z
Completion Time	2025-05-05T22:31:45Z
Response:	
attribute	value
Status	unlinked
Inode Owner	false
Perm Change Flag	chmodAndSetAcl
Prevent Snapshot Restore	false
Perm Inherit Flag	inheritAclOnly
Fal Status	Unspecified
Inode Space Designation	root
Fileset Name	fileset1
Id	2
Root Inode	35587
Comment	
Max Inodes	0
IAM Mode	off
Filesystem Name	gpfs0
Path	--
Parent Id	
Creation Time	2025-05-05T22:31:45Z
Inode Space	0
Snap Id	0
Domain Id	0
Allocated Inodes	0

SSH & SCP Replacement

Eliminate the requirement of passwordless SSH across cluster nodes
by leveraging new Scale-API Architecture

Note: mm-commands are still compatible during this phase

```
# scalectl cluster list
```

Storage Scale cluster information

GPFS cluster name : scale-cluster-4

GPFS cluster id : 15748700724538103684


GPFS UID domain : scale-cluster-4

Remote shell command : /usr/lpp/mmfs/bin/scaleadmremotexecute

Remote file copy command : /usr/lpp/mmfs/bin/scaleadmremotetransfer

Repository type : CCR

Manager : nopriv-41



Configuration Management

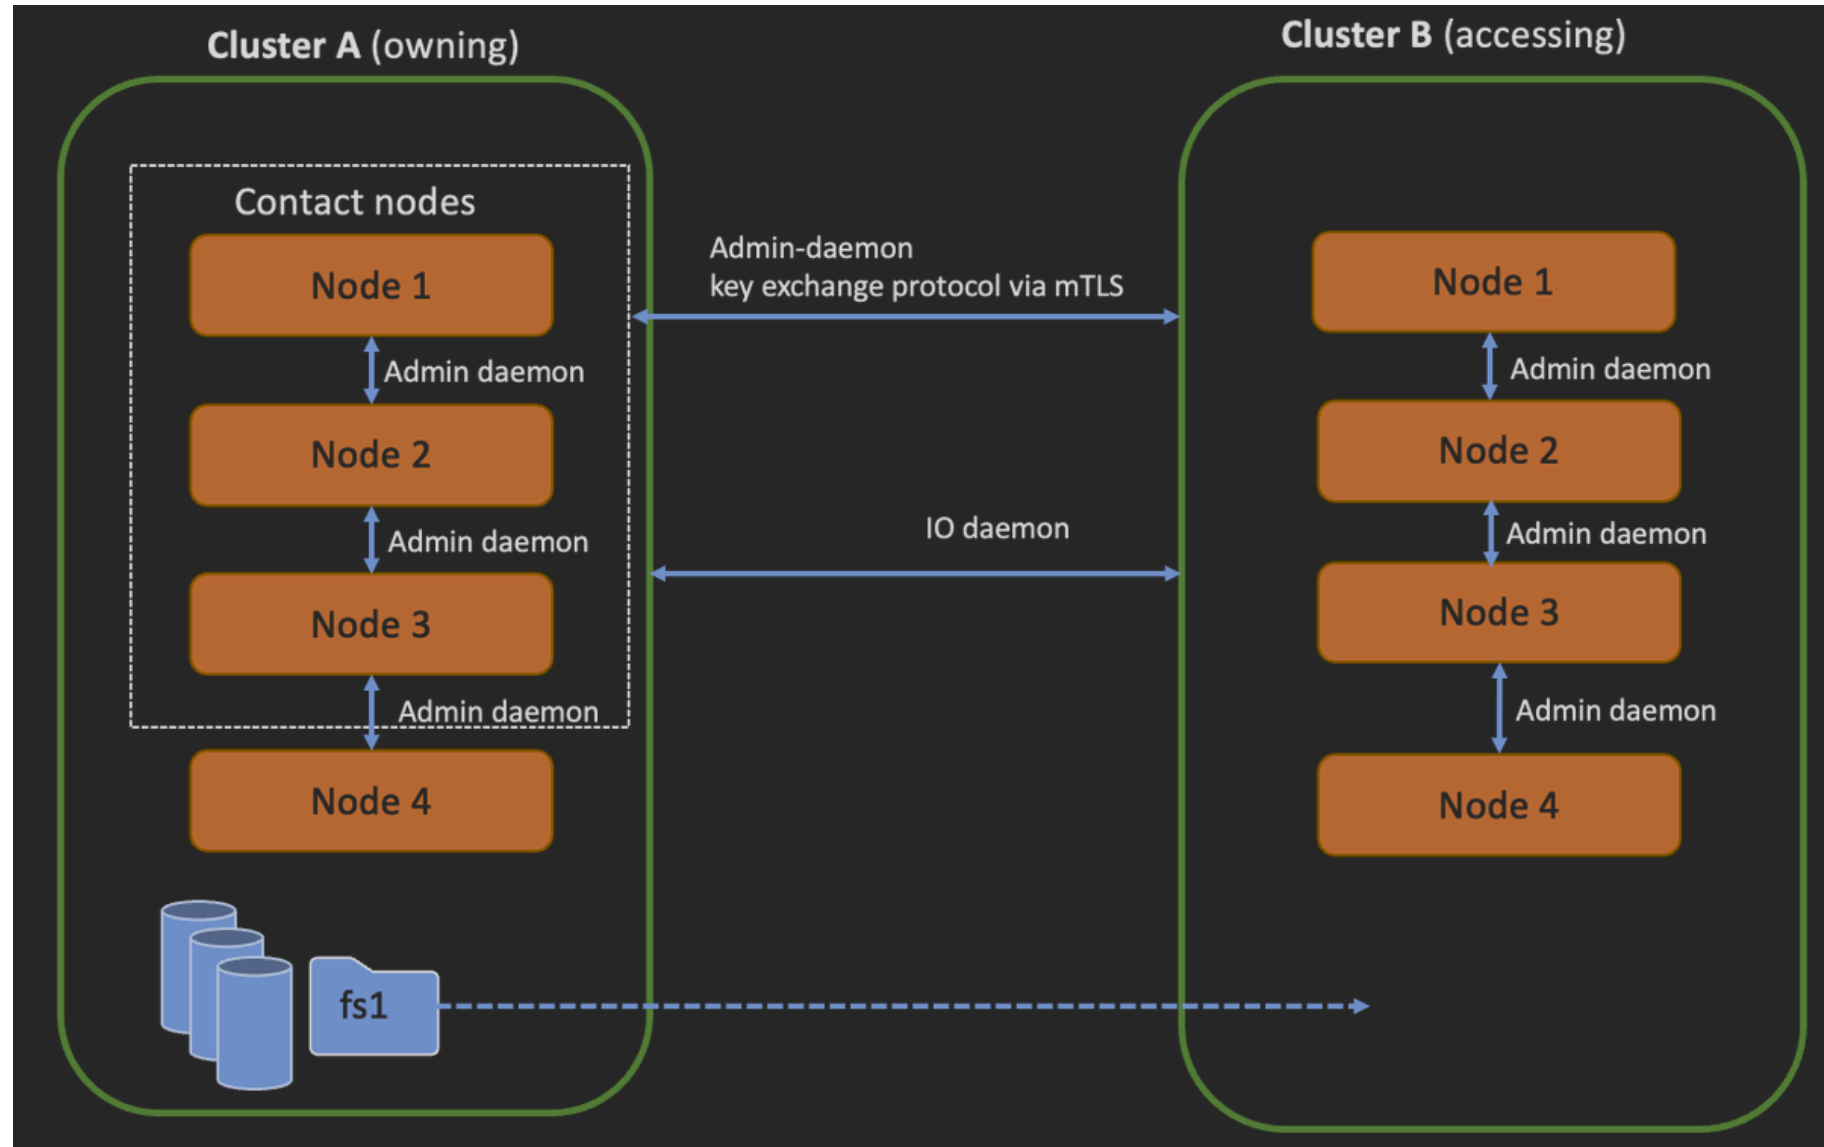
```
[root@nopriv-41 ~]# scalectl config admin list
```

Attribute Name	Value	Nodes
log.file	scaleadmd.log	
log.level	debug	
log.path	/var/mmfs/scaleadm/log	
log.utc	true	
lro.cleanup_job_age	240h	
lro.cleanup_number_jobs	5000	
lro.dir	/var/mmfs/scaleadm/lro	
lro.output_limit	32	
lro.replicate_num_jobs	1000	
lro.replicate_wait_time	30s	
lro.rest_wait_time	2s	
lro.stale_wait_time	300s	
nodeid.cachain	/var/mmfs/ssl/cachain.pem	
nodeid.certificate	/var/mmfs/ssl/cert.pem	
nodeid.key	/var/mmfs/ssl/key.pem	
nsd.list_cache_age	30s	
remote_grpc_server.host	0.0.0.0	
remote_grpc_server.port	50052	
remote_grpc_server.scheme	https	
rest_server.authentication_type	pam	

```
[root@nopriv-41 ~]# scalectl config cluster list
```

Attribute Name	Value	Nodes
clusterName	scale-cluster-4	
clusterId	15748700724538103684	
scaleApiEnabled	yes	
dmapiFileHandleSize	32	
minReleaseLevel	3700 5.2.3.0	
tscCmdAllowRemoteConnections	no	
ccrEnabled	yes	
sdrNotifyAuthEnabled	yes	
autoload	yes	
maxblocksize	16384K	
maxFilesToCache	100010	
maxStatCache	0	
cipherList	AUTHONLY	
traceRecycle	local	
trace	all 0	
adminMode	central	

Remote Cluster



Key Exchange Protocol

```
scalectl cluster remote add --name owningcluster --contact-nodes 9.1.2.3,...
```

Accessing cluster

Send accessing cluster keys

Store owning cluster keys
for communication

Success

Owning cluster

Receive accessing cluster keys

Verify accessing cluster identity and verify
the accessing cluster is authorized

Store accessing cluster keys for
communication

Load and reply with owning cluster keys

Outlook

5.2.3.3	Future Priority Items
Batch node add	Policy (mmapplypolicy) <ul style="list-style-type: none">• Encryption
CA certs (trusts)	Quotas
	ESS Integration
	AFM Integration
	CSI/CNSA Integration
	Protocols Integration <ul style="list-style-type: none">• NFS priority
	Delete node
	QOS
	Caching

Thank you for
choosing Storage
Scale!

