

# readme.md for @lossless.zone/objectstorag e

“ ☐ S3-compatible object storage server with clustering, erasure coding, and a slick management UI — powered by [smartstorage](#) .

**objectstorage** gives you a fully featured, self-hosted S3-compatible storage server with a beautiful web-based management interface — all in a single Docker image. No Java, no bloat, no fuss.

Built on Deno for the backend and [@design.estate/dees-catalog](#) for a polished UI, it speaks the S3 protocol out of the box while adding powerful management features on top. Scale from a single node to a distributed cluster with erasure coding and multi-drive support.

## Issue Reporting and Security

For reporting bugs, issues, or security vulnerabilities, please visit [community.foss.global/](#). This is the central community hub for all issue reporting. Developers who sign and comply with our contribution agreement and go through identification can also get a [code.foss.global/](#) account to submit Pull Requests directly.

## ☐ Features

- **Full S3 API compatibility** — Works with any S3 client, SDK, or tool (AWS CLI, boto3, etc.)
- ☐ **Cluster mode** — Distribute storage across multiple nodes with QUIC transport, automatic discovery, and quorum writes/reads

- **Erasure coding** — Reed-Solomon erasure coding (default 4+2) for data durability with minimal overhead
- **Multi-drive support** — Stripe data across multiple disks per node with per-drive health monitoring
- **Self-healing** — Background scanner detects and reconstructs missing or corrupt shards automatically
- **Management UI** — Web dashboard for buckets, objects, policies, credentials, cluster config, and storage drives
- **Finder-style object browser** — Column view with file preview, drag-and-drop upload, move/rename, context menus
- **Inline code editing** — Built-in Monaco editor with syntax highlighting and save-back-to-storage
- **PDF viewer** — Render PDFs inline with page navigation, zoom, and thumbnails
- **Named policy management** — Create reusable IAM-style policies, attach them to multiple buckets
- **Credential management** — Add/remove access keys through the UI with live-reload
- **Single Docker image** — Multi-arch (amd64 + arm64), tiny Alpine-based image
- **Fast** — Rust-powered storage engine via smartstorage, streaming I/O with zero-copy and backpressure
- **Secure by default** — JWT-based admin auth, S3 SigV4 authentication, bucket policies
- **Dark theme** — Automatic dark mode following your system preference

## ⚡ Quick Start

### Docker (recommended)

```
docker run -d \  
  --name objectstorage \  
  -p 9000:9000 \  
  -p 3000:3000 \  
  -v objstdata:/data \  
  -e OBJST_ACCESS_KEY=myadminkey \  
  -e OBJST_SECRET_KEY=mysupersecret \  
  -e OBJST_ADMIN_PASSWORD=myuipassword \  
  code.fossilglobal/lossless.zone/objectstorage:latest
```

Then open **http://localhost:3000** for the management UI and use **http://localhost:9000** as your S3 endpoint.

# Deno (development)

```
# Clone and install frontend dependencies
git clone ssh://git@code.foss.global:29419/lossless.zone/objectstorage.git
cd objectstorage
pnpm install
pnpm run build

# Run in ephemeral mode (data stored in .nogit/objstdata)
deno run --allow-all mod.ts server --ephemeral
```

## Configuration

objectstorage is configured through environment variables, CLI flags, or programmatic config. **Environment variables take precedence** over CLI flags.

## Server Environment Variables

Variable	Description	Default
<code>OBJST_PORT</code>	Storage API port	<code>9000</code>
<code>UI_PORT</code>	Management UI port	<code>3000</code>
<code>OBJST_STORAGE_DIR</code>	Data storage directory	<code>/data</code>
<code>OBJST_ACCESS_KEY</code>	Access key ID	<code>admin</code>
<code>OBJST_SECRET_KEY</code>	Secret access key	<code>admin</code>
<code>OBJST_ADMIN_PASSWORD</code>	Admin UI password	<code>admin</code>
<code>OBJST_REGION</code>	Storage region identifier	<code>us-east-1</code>

## Cluster Environment Variables

Variable	Description	Default
<code>OBJST_CLUSTER_ENABLED</code>	Enable cluster mode ( <code>true</code> / <code>false</code> )	<code>false</code>
<code>OBJST_CLUSTER_NODE_ID</code>	Unique node identifier	auto-generated
<code>OBJST_CLUSTER_QUIC_PORT</code>	QUIC transport port	<code>4433</code>

Variable	Description	Default
<code>OBJST_CLUSTER_SEED_NODES</code>	Comma-separated seed node addresses	<i>(empty)</i>
<code>OBJST_DRIVE_PATHS</code>	Comma-separated drive mount paths	storage dir
<code>OBJST_ERASURE_DATA_SHARDS</code>	Erasure coding data shards	4
<code>OBJST_ERASURE_PARITY_SHARDS</code>	Erasure coding parity shards	2
<code>OBJST_ERASURE_CHUNK_SIZE</code>	Erasure chunk size in bytes	4194304 (4 MB)
<code>OBJST_HEARTBEAT_INTERVAL_MS</code>	Cluster heartbeat interval	5000
<code>OBJST_HEARTBEAT_TIMEOUT_MS</code>	Cluster heartbeat timeout	30000

## CLI Flags

```
deno run --allow-all mod.ts server [options]
```

### Server Options:

```
--storage-port <port>      Storage API port (default: 9000)
--ui-port <port>           Management UI port (default: 3000)
--storage-dir <path>       Storage directory (default: /data)
--ephemeral                Use ./nologit/objstdata for storage (dev mode)
```

### Clustering Options:

```
--cluster-enabled          Enable cluster mode
--cluster-node-id <id>     Unique node identifier
--cluster-quic-port <port> QUIC transport port (default: 4433)
--cluster-seed-nodes <list> Comma-separated seed node addresses
--drive-paths <list>       Comma-separated drive mount paths
--erasure-data-shards <n>  Erasure coding data shards (default: 4)
--erasure-parity-shards <n> Erasure coding parity shards (default: 2)
```

## Cluster Mode

objectstorage supports distributed storage across multiple nodes with automatic failover and data redundancy.

## How it works

1. **Enable clustering** on each node with `OBJST_CLUSTER_ENABLED=true`
2. **Point nodes at each other** using `OBJST_CLUSTER_SEED_NODES` — nodes discover the full cluster from any seed
3. **Configure drives** per node with `OBJST_DRIVE_PATHS` — each drive is independently managed
4. **Erasur coding** splits objects into data + parity shards across drives and nodes

## Example: 3-node cluster

```
# Node 1
docker run -d --name objst-node1 \
  -p 9000:9000 -p 3000:3000 -p 4433:4433/udp \
  -v /mnt/disk1:/drive1 -v /mnt/disk2:/drive2 \
  -e OBJST_CLUSTER_ENABLED=true \
  -e OBJST_CLUSTER_NODE_ID=node-1 \
  -e OBJST_CLUSTER_SEED_NODES=node2:4433,node3:4433 \
  -e OBJST_DRIVE_PATHS=/drive1,/drive2 \
  -e OBJST_ACCESS_KEY=myadminkey \
  -e OBJST_SECRET_KEY=mysupersecret \
  code.foss.global/lossless.zone/objectstorage:latest

# Node 2 and Node 3 – same pattern, different node IDs and seed nodes
```

## Erasur coding presets

Config	Data Shards	Parity Shards	Overhead	Fault Tolerance
Default	4	2	50%	2 failures
High durability	6	3	50%	3 failures
Minimal	2	1	50%	1 failure

## Inter-node transport

Cluster communication uses **QUIC** (UDP port 4433 by default) with:

- Auto-generated TLS certificates
- Multiplexed streams with flow-control backpressure
- Heartbeat-based failure detection (default: 5s interval, 30s timeout)

# Management UI

The web-based management UI is served on the UI port (default: `3000`). Log in with username `admin` and the configured admin password.

## Overview

Dashboard showing server status, uptime, storage usage, bucket count, and connection info.

Overview

## Buckets

Create/delete buckets. View object counts and sizes. Attach/detach named policies per bucket.

Buckets

## Browser

Finder-style column browser for objects. Upload, download, preview, move, rename, and delete files and folders — with syntax-highlighted code preview.

Browser

Right-click any item for quick actions:

Context Menu

## Inline Code Editing

Click Edit on any text file to open the built-in Monaco editor with syntax highlighting, language detection, and save-back-to-storage.

Code Editing

## PDF Viewer

PDF files render inline with a full-featured viewer — page navigation, zoom, fit-to-page, thumbnails, download, and print.

PDF Viewer

# Policies

Create reusable named policies with IAM-style S3 statements. Attach policies to multiple buckets at once.

Policies

## Attaching Policies to Buckets

From the Buckets view, click the policy icon on any bucket to see attached and available policies. Attach or detach with one click.

Attach Policy

# Access Keys

Add and remove access credentials. Secret keys are masked. Changes take effect immediately — no server restart needed.

Access Keys

## Adding Access Keys

Click "Add Key" to create new access credentials. They're immediately available for API authentication.

Add Access Key

# Configuration

View your server's current configuration at a glance — ports, region, storage directory, auth/CORS status, cluster configuration, erasure coding settings, and storage drive paths. The config view also includes an environment variable reference guide for cluster setup.

Configuration

# Dark Theme

Full dark theme support — automatically follows your system preference via `prefers-color-scheme`.

Dark Theme

# ☐☐ Named Policy System

objectstorage adds a **named policy** abstraction on top of standard S3 bucket policies. Instead of editing raw JSON per bucket, you define reusable policy templates and attach them to any number of buckets.

## How it works

1. **Create a named policy** in the Policies view — give it a name, description, and S3 policy statements
2. **Attach it to buckets** — from the Policies view or the Buckets view
3. **objectstorage merges** all attached policy statements into a single S3 policy document and applies it to the bucket automatically

## `${bucket}` placeholder

Use `${bucket}` in your policy's `Resource` ARN and it will be replaced with the actual bucket name when applied:

```
[
  {
    "Sid": "PublicRead",
    "Effect": "Allow",
    "Principal": "*",
    "Action": "s3:GetObject",
    "Resource": "arn:aws:s3:::${bucket}/*"
  }
]
```

This lets one policy like "Public Read" work across many buckets without hardcoding names.

## Lifecycle

- **Updating a policy** automatically recomputes and re-applies the merged S3 policy on all attached buckets
- **Deleting a policy** detaches it from all buckets and recomputes each
- **Deleting a bucket** cleans up its policy attachments automatically

# ☐ S3 API Usage

Use any S3-compatible client to interact with the storage. Here are some examples:

## AWS CLI

```
# Configure AWS CLI
aws configure set aws_access_key_id admin
aws configure set aws_secret_access_key admin
aws configure set default.region us-east-1

# Create a bucket
aws --endpoint-url http://localhost:9000 s3 mb s3://my-bucket

# Upload a file
aws --endpoint-url http://localhost:9000 s3 cp myfile.txt s3://my-bucket/

# List objects
aws --endpoint-url http://localhost:9000 s3 ls s3://my-bucket/

# Download a file
aws --endpoint-url http://localhost:9000 s3 cp s3://my-bucket/myfile.txt ./downloaded.txt
```

## Node.js / TypeScript (AWS SDK v3)

```
import { S3Client, PutObjectCommand, ListObjectsV2Command } from '@aws-sdk/client-s3';

const s3 = new S3Client({
  endpoint: 'http://localhost:9000',
  region: 'us-east-1',
  credentials: {
    accessKeyId: 'admin',
    secretAccessKey: 'admin',
  },
  forcePathStyle: true,
});
```

```
// Upload an object
await s3.send(new PutObjectCommand({
  Bucket: 'my-bucket',
  Key: 'hello.txt',
  Body: 'Hello, S3!',
}));

// List objects in a bucket
const result = await s3.send(new ListObjectsV2Command({
  Bucket: 'my-bucket',
}));
console.log(result.Contents);
```

## Python (boto3)

```
import boto3

s3 = boto3.client(
    's3',
    endpoint_url='http://localhost:9000',
    aws_access_key_id='admin',
    aws_secret_access_key='admin',
    region_name='us-east-1',
)

# Upload a file
s3.put_object(Bucket='my-bucket', Key='hello.txt', Body=b'Hello from Python!')

# List objects
response = s3.list_objects_v2(Bucket='my-bucket')
for obj in response.get('Contents', []):
    print(obj['Key'], obj['Size'])
```



# Build

```
# Build for the native platform
pnpm run build:docker

# Or build and run directly
pnpm run start:docker
```

## Docker Compose (standalone)

```
services:
  objectstorage:
    image: code.foss.global/lossless.zone/objectstorage:latest
    ports:
      - "9000:9000" # S3 API
      - "3000:3000" # Management UI
    volumes:
      - objstdata:/data
    environment:
      OBJST_ACCESS_KEY: myadminkey
      OBJST_SECRET_KEY: mysupersecret
      OBJST_ADMIN_PASSWORD: securepw123

volumes:
  objstdata:
```

## Docker Compose (3-node cluster)

```
services:
  node1:
    image: code.foss.global/lossless.zone/objectstorage:latest
    ports:
      - "9001:9000"
      - "3001:3000"
      - "4433:4433/udp"
    volumes:
```

- node1-drive1:/drive1
- node1-drive2:/drive2

environment:

```
OBJST_CLUSTER_ENABLED: "true"
OBJST_CLUSTER_NODE_ID: node-1
OBJST_CLUSTER_QUIC_PORT: "4433"
OBJST_CLUSTER_SEED_NODES: node2:4433,node3:4433
OBJST_DRIVE_PATHS: /drive1,/drive2
OBJST_ACCESS_KEY: myadminkey
OBJST_SECRET_KEY: mysupersecret
```

node2:

image: code.foss.global/lossless.zone/objectstorage:latest

ports:

- "9002:9000"
- "3002:3000"
- "4434:4433/udp"

volumes:

- node2-drive1:/drive1
- node2-drive2:/drive2

environment:

```
OBJST_CLUSTER_ENABLED: "true"
OBJST_CLUSTER_NODE_ID: node-2
OBJST_CLUSTER_QUIC_PORT: "4433"
OBJST_CLUSTER_SEED_NODES: node1:4433,node3:4433
OBJST_DRIVE_PATHS: /drive1,/drive2
OBJST_ACCESS_KEY: myadminkey
OBJST_SECRET_KEY: mysupersecret
```

node3:

image: code.foss.global/lossless.zone/objectstorage:latest

ports:

- "9003:9000"
- "3003:3000"
- "4435:4433/udp"

volumes:

- node3-drive1:/drive1
- node3-drive2:/drive2

environment:

```
OBJST_CLUSTER_ENABLED: "true"
OBJST_CLUSTER_NODE_ID: node-3
OBJST_CLUSTER_QUIC_PORT: "4433"
OBJST_CLUSTER_SEED_NODES: node1:4433,node2:4433
OBJST_DRIVE_PATHS: /drive1,/drive2
OBJST_ACCESS_KEY: myadminkey
OBJST_SECRET_KEY: mysupersecret
```

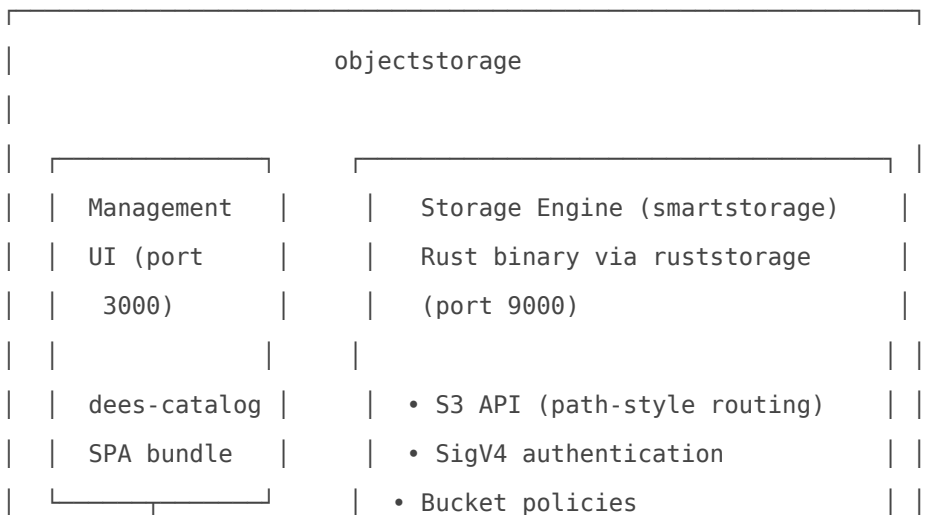
volumes:

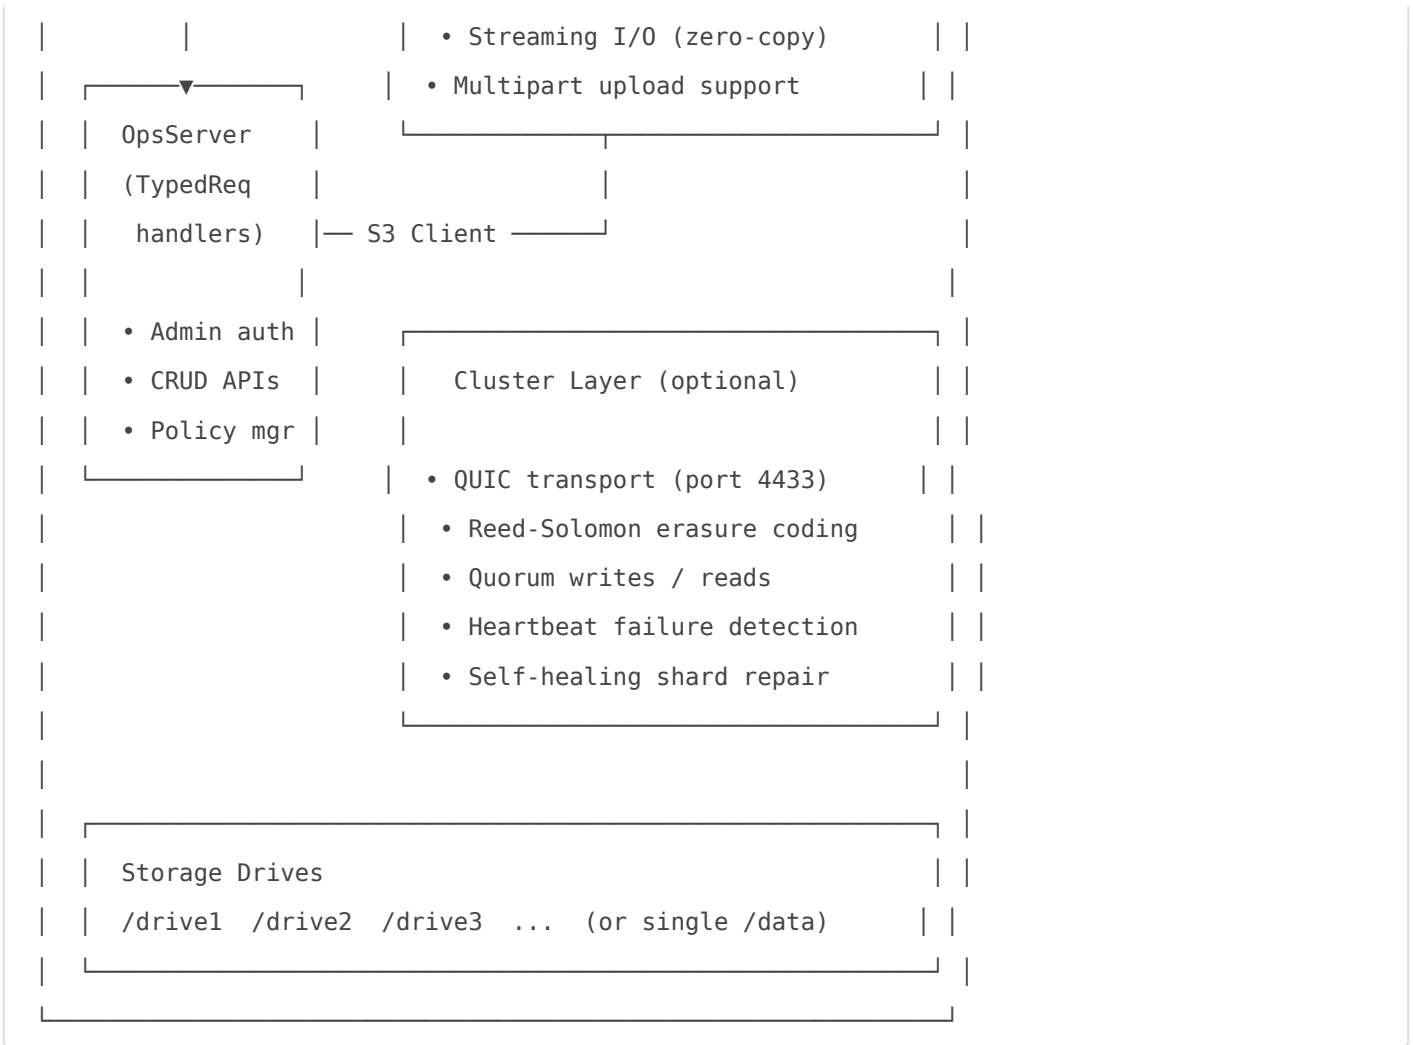
```
node1-drive1:
node1-drive2:
node2-drive1:
node2-drive2:
node3-drive1:
node3-drive2:
```

## Image Details

- **Base:** `alpine:edge` with Deno runtime
- **Architectures:** `linux/amd64`, `linux/arm64`
- **Size:** ~150 MB compressed
- **Init system:** `tini` for proper signal handling
- **Exposed ports:** `9000` (S3), `3000` (UI), `4433` (QUIC cluster transport)
- **Volume:** `/data` — all bucket data and config persisted here

## Architecture





## Tech Stack

Layer	Technology
<b>Storage Engine</b>	<a href="#">@push.rocks/smartstorage</a> (Rust binary via <a href="#">ruststorage</a> )
<b>Cluster Transport</b>	QUIC via <a href="#">quinn</a> (auto-TLS, multiplexed streams, backpressure)
<b>Erasure Coding</b>	Reed-Solomon (configurable data + parity shards)
<b>Runtime</b>	Deno
<b>Management API</b>	<a href="#">@api.global/typedrequest</a> + <a href="#">@api.global/typedserver</a>
<b>Auth</b>	JWT via <a href="#">@push.rocks/smartjwt</a> , S3 SigV4
<b>Frontend</b>	<a href="#">@design.estate/dees-element</a> (LitElement) + <a href="#">@design.estate/dees-catalog</a>
<b>Frontend Build</b>	esbuild via <a href="#">@git.zone/tsbundle</a>
<b>Docker</b>	Multi-stage (Node.js build → Alpine + Deno runtime)

# Development

```
# Install dependencies
pnpm install

# Watch mode – auto-rebuilds frontend + restarts backend
pnpm run watch

# Build frontend bundle only
pnpm run build

# Type check backend
deno check mod.ts

# Run tests
pnpm test

# Run in development mode
deno run --allow-all mod.ts server --ephemeral
```

## License and Legal Information

This repository contains open-source code licensed under the MIT License. A copy of the license can be found in the [license](#) file.

**Please note:** The MIT License does not grant permission to use the trade names, trademarks, service marks, or product names of the project, except as required for reasonable and customary use in describing the origin of the work and reproducing the content of the NOTICE file.

## Trademarks

This project is owned and maintained by Task Venture Capital GmbH. The names and logos associated with Task Venture Capital GmbH and any related products or services are trademarks of Task Venture Capital GmbH or third parties, and are not included within the scope of the MIT license granted herein.

Use of these trademarks must comply with Task Venture Capital GmbH's Trademark Guidelines or the guidelines of the respective third-party owners, and any usage must be approved in writing. Third-party trademarks used herein are the property of their respective owners and used only in a descriptive manner, e.g. for an implementation of an API or similar.

# Company Information

Task Venture Capital GmbH Registered at District Court Bremen HRB 35230 HB, Germany

For any legal inquiries or further information, please contact us via email at [hello@task.vc](mailto:hello@task.vc).

By using this repository, you acknowledge that you have read this section, agree to comply with its terms, and understand that the licensing of the code does not imply endorsement by Task Venture Capital GmbH of any derivative works.

---

Revision #3

Created 2026-03-28 11:09:41 UTC by foss.global Team

Updated 2026-03-28 12:16:17 UTC by foss.global Team