PUBLIC TECHNICAL DOCUMENTATION

Self-Hosted Private Cloud Infrastructure

A resilient NAS, private file service, and role-authorised Drive platform

Designed, built, tested, and maintained by

Md Samsudduha Shawon

Release	Version 1.0 - Public edition
Published	September 2026
Project page	https://meetshawon.com/projects/private-cloud-infrastructure
Service gateway	https://drive.meetshawon.com
Project status	Operational - authorised users only

Infrastructure / Cybersecurity / Storage Engineering

DOCUMENT GUIDE

About this documentation

This public document explains the completed NAS and private-cloud platform from requirements through operation while intentionally excluding secrets and sensitive infrastructure details.

Audience	Recruiters, hiring managers, technical reviewers, collaborators, and authorised stakeholders
Document type	Public technical case study, architecture overview, and operational summary
Scope	Hardware, TrueNAS, ZFS, application services, identity, provisioning, remote access, assurance, and maintenance
Excluded	Credentials, tokens, private addresses, internal paths, exact dataset identifiers, and recovery secrets
Owner	Md Samsudduha Shawon

> **Security boundary:** The documentation is deliberately sanitised. It explains engineering decisions and validated outcomes without providing information that would materially reduce the security of the live service.

Contents

01 / EXECUTIVE SUMMARY

A private cloud moved from concept to operation

The project delivers a working, self-hosted storage platform that combines resilient local storage, controlled remote access, automated per-user provisioning, and a portfolio-integrated Drive experience.

The original requirement was straightforward: provide a small number of authorised users with a familiar, private file-storage service while retaining control of the hardware, data location, access rules, and operational lifecycle. Meeting that requirement required more than installing a NAS operating system. The completed solution integrates hardware selection, ZFS storage design, application isolation, identity and role checks, secure external connectivity, automated provisioning, monitoring, maintenance, and recovery planning.

The final system is based on an HP EliteDesk 800 G3 SFF running TrueNAS SCALE. A dedicated NVMe device holds the operating system, a separate SATA SSD supports application workloads, and two enterprise-class 2 TB hard disks form a mirrored ZFS storage pool. The public portfolio application provides the user-facing Drive gateway and authorisation boundary; the NAS administration surface remains private.

> **Current outcome:** The service is operational, role-protected, remotely accessible over encrypted connections, and validated for isolated user storage, quotas, file operations, monitoring, and recovery-oriented tasks.

Platform	HP EliteDesk 800 G3 SFF with TrueNAS SCALE
Compute	Intel Core i5-7500 and 16 GB DDR4-2666 memory
Storage	128 GB NVMe boot, 128 GB SATA applications SSD, and 2 x 2 TB mirrored HDDs
Usable capacity	Approximately 1.76 TiB before user quotas and operational overhead
Access	Authorised admin and partner users through the Drive gateway

02 / CONTEXT AND MOTIVATION

Why self-hosted storage

The project was designed as both a useful private service and a practical infrastructure exercise in storage, security, automation, and operations.

Commercial cloud storage is convenient, but it abstracts away many of the engineering decisions that determine how data is stored, protected, accessed, monitored, and recovered. Building a private alternative created an opportunity to learn those decisions directly while producing a service with genuine day-to-day value.

Primary motivations

- Maintain direct control over the storage hardware and service configuration.
- Provide separate accounts and storage areas for a limited group of approved users.
- Use mirrored storage to continue operating after a single data-drive failure.
- Support remote file access without exposing the TrueNAS administration interface.
- Automate repetitive provisioning tasks to reduce manual configuration errors.
- Develop demonstrable skills across Linux, ZFS, networking, access control, APIs, monitoring, and recovery planning.

Engineering principles

Layer	Technology	Engineering role
Separation	Dedicated storage roles	Keep boot, applications, and user data on distinct devices or pools.
Least privilege	Role and ACL controls	Grant only the access required for each user and workflow.
Resilience	ZFS mirror	Tolerate one data-disk failure while maintaining service availability.
Defence in depth	Layered controls	Combine identity, HTTPS, private administration, isolation, and monitoring.
Repeatability	Automated provisioning	Create user storage using consistent, testable procedures.

03 / REQUIREMENTS

Scope, constraints, and success criteria

Requirements were defined before implementation so that hardware decisions, security controls, and validation activities could be assessed against explicit outcomes.

Functional requirements

- Authenticated access for approved users from inside and outside the home network.
- Separate storage allocation for each user with configurable quota enforcement.
- Common file operations including upload, download, organisation, and deletion within authorised storage.
- Administrative control over approval, provisioning, suspension, retention, and eventual deletion.
- A user-facing Drive experience integrated with the professional portfolio platform.

Non-functional requirements

- Encrypted external sessions and no direct publication of the NAS management interface.
- Single-drive fault tolerance for the primary user-data pool.
- Predictable permissions and isolation between users.
- Maintainable hardware with acceptable cooling, power, cabling, and drive serviceability.
- Monitoring and maintenance routines that make storage health and capacity visible.
- Sanitised public documentation suitable for professional review.

> **Out of scope:** The platform is not positioned as an enterprise backup provider, high-availability cluster, or unlimited public storage service. Access remains deliberately limited to approved users.

04 / HARDWARE ENGINEERING

Purpose-selected physical platform

The final hardware was selected only after verifying physical compatibility, storage connectivity, cooling, power delivery, and serviceability rather than relying on processor and memory specifications alone.

Selection criteria

- Physical accommodation for two dedicated 3.5-inch data drives without unsafe mounting or blocked airflow.
- Sufficient SATA connectivity and power delivery for boot, application, and mirrored-data devices.
- A business-class chassis that could be opened, inspected, cleaned, cabled, and serviced predictably.
- Enough processor and memory capacity for TrueNAS, ZFS, encryption, application services, and light multi-user activity.
- Acceptable acquisition and running cost for a small-scale home infrastructure project.

> **Decision gate:** The platform was commissioned only after the drive bays, mounting, cabling, clearance, airflow, and service procedure had been physically verified.

Figure 1. Sanitised view of the completed internal NAS hardware layout.

System	HP EliteDesk 800 G3 SFF
Processor	Intel Core i5-7500
Memory	16 GB DDR4-2666 configured as 4 x 4 GB
Boot device	128 GB NVMe SSD dedicated to TrueNAS
Application device	128 GB SATA SSD separated from the data mirror
Data devices	2 x 2 TB Seagate Exos SATA hard disks

Why this design works

- Business-class hardware provides a stable platform and useful serviceability features.
- The processor is sufficient for storage, application, encryption, and light multi-user workloads.
- Sixteen gigabytes of memory supports TrueNAS, ZFS caching, and the current application footprint.
- Dedicated boot and application devices prevent routine system and application I/O from competing with user data.
- The two data disks provide a simple mirrored topology that matches the small-scale capacity and resilience requirement.

05 / STORAGE DESIGN

Physical separation and mirrored user data

The storage topology assigns a clear purpose to each device and protects the primary user-data pool with ZFS mirroring.

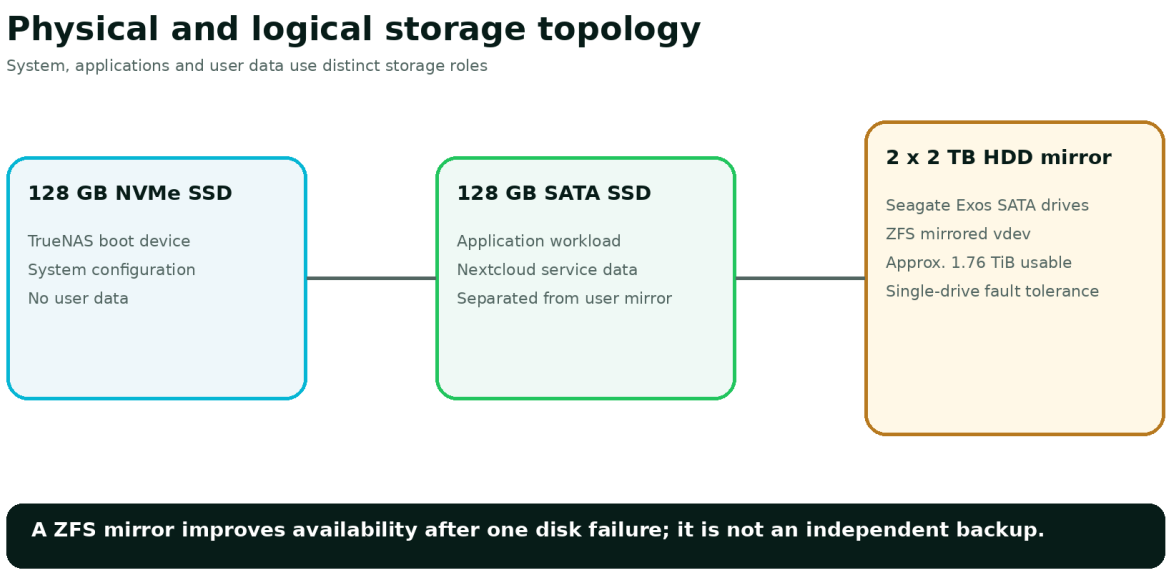

Figure 2. Public storage topology showing device roles and the mirrored user-data boundary.

ZFS mirror behaviour

The two 2 TB hard disks form a mirrored vdev. Each block written to the data pool is stored on both disks, allowing the pool to remain available if one disk fails. Usable capacity is therefore approximately the capacity of one drive rather than the sum of both drives. The reported usable capacity is approximately 1.76 TiB after unit conversion and filesystem overhead.

> **Important distinction:** Mirroring improves availability; it does not protect against accidental deletion, ransomware, corruption replicated to both disks, theft, fire, or administrative mistakes. Independent backups remain necessary.

Logical organisation

- System configuration remains on the dedicated boot device.
- Application workloads use the separate SATA SSD.
- User data is stored in isolated ZFS datasets on the mirrored pool.
- Per-user quotas limit consumption and support predictable capacity management.
- ACLs and service permissions restrict access to the correct identity and workflow.

06 / ARCHITECTURE

Public experience, private infrastructure

The design separates the public application, identity checks, provisioning workflow, application service, storage management, and physical data boundary.

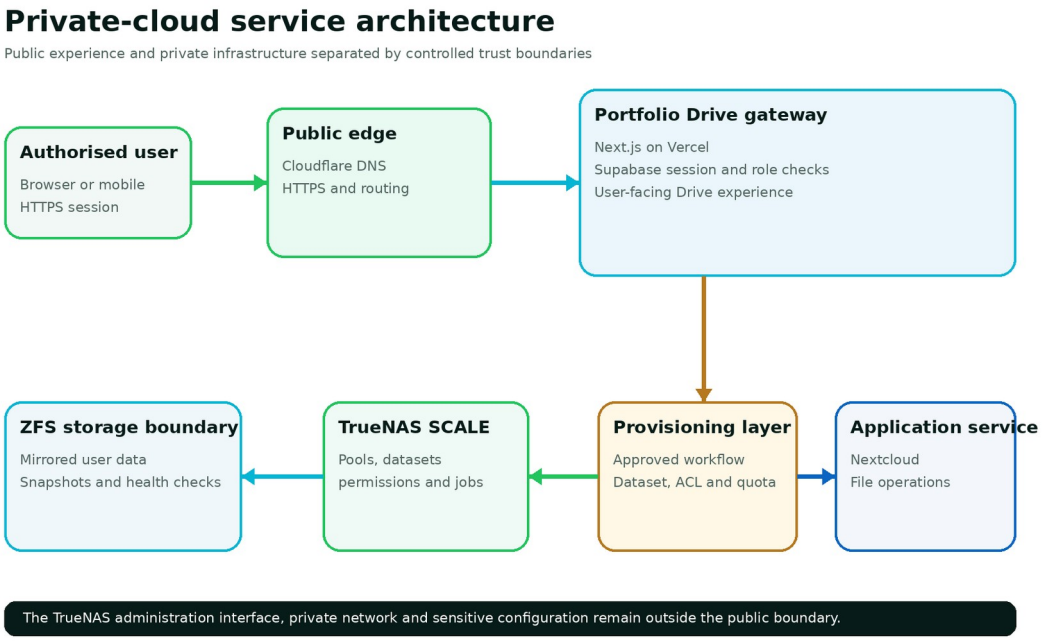

Figure 3. Sanitised high-level architecture and trust boundaries.

Layer	Technology	Engineering role
Domain edge	Cloudflare	DNS, HTTPS routing, and the public hostname boundary.
Gateway	Next.js on Vercel	User-facing Drive experience and protected application routes.
Identity	Supabase	Authentication, session handling, profile roles, and authorisation context.
Application	Nextcloud	Familiar web-based file operations for approved users.
Storage OS	TrueNAS SCALE	Pools, datasets, ACLs, applications, health, and maintenance jobs.
Filesystem	ZFS	Mirroring, datasets, quotas, checksums, snapshots, and storage visibility.

07 / IDENTITY AND ACCESS

Authentication is not authorisation

A valid login establishes identity, but server-side role and account checks determine whether the user may enter the Drive service or trigger storage operations.

Role model

Admin	Manages approvals, provisioning, service controls, suspension, retention, and guarded deletion.
Partner	Approved Drive user with access to assigned storage and permitted file operations.
Other authenticated roles	May use relevant portfolio features but are denied Drive access.
Unauthenticated visitor	Redirected to sign-in and cannot reach protected Drive functions.

Access decision sequence

1. The user opens drive.meetshawon.com and is redirected to the controlled authentication flow if no valid session exists.
2. The application validates the session on the server and retrieves the associated profile role.
3. Only an authorised role is permitted to continue into the Drive experience.
4. Storage operations are restricted to the user and storage allocation represented by the approved account state.
5. Suspended, unapproved, or otherwise ineligible users receive a controlled denial outcome rather than access to the private service.

> **Security property:** Client-side visibility is never treated as sufficient authorisation. Protected decisions are enforced through server-side checks and controlled backend workflows.

08 / PROVISIONING

Repeatable isolated user storage

Approved accounts are processed through a controlled workflow that creates consistent storage allocations without granting users direct administrative access.

Provisioning workflow

1. An administrator approves the account for Drive use through the protected management interface.
2. The backend records a provisioning request with the minimum information required for the operation.
3. A restricted workflow creates the user's isolated storage dataset and applies the configured quota.
4. The required ownership and ACL rules are applied to align the storage identity with the application account.
5. The workflow validates the result before marking the storage account active.
6. The user receives a controlled notification and can access only the assigned storage area.

Lifecycle controls

- Default allocations use a controlled quota profile; the implemented workflow supports approximately 101 GiB per standard approved user unless an administrator assigns another value.
- Suspension removes active access while retaining data for administrative review and recovery.
- Retention separates account access decisions from irreversible data deletion.
- Deletion is guarded and occurs only after prerequisite checks and explicit administrative confirmation.
- Operational logs avoid unnecessary exposure of private identifiers and secrets.

09 / REMOTE CONNECTIVITY

Encrypted access with limited exposure

Remote access is designed around the user-facing service rather than publishing the TrueNAS administration surface or raw file-sharing protocols to the internet.

The drive.meetshawon.com hostname provides the public entry point. Cloudflare manages the domain boundary, while the portfolio platform handles authentication and role authorisation. External sessions use HTTPS so credentials and file transfers are encrypted in transit. The storage system and its administrative interface remain separated from the public experience.

Exposure controls

- No public link to the TrueNAS administration interface.
- No publication of SMB or other internal storage protocols for general internet access.
- Role-restricted application routes and server-side session validation.
- Public documentation omits internal addresses, hostnames, dataset names, credentials, and tokens.
- The remote-access route can be disabled independently of the underlying local storage service.

> **Operational principle:** The public hostname provides a controlled service boundary; it is not a shortcut around storage permissions, application authentication, or administrative separation.

10 / SECURITY ENGINEERING

Defence in depth across the service

The security model combines identity, role checks, encryption, isolated storage, restricted automation, private administration, monitoring, and recoverable lifecycle controls.

Layer	Technology	Engineering role
Identity	Supabase sessions	Establish the authenticated account using managed session controls.
Authorisation	Server-side role checks	Limit Drive entry and management functions to approved roles.
Transport	HTTPS	Protect credentials and data in transit over external networks.
Isolation	ZFS datasets and ACLs	Separate user storage and prevent cross-user access.
Automation	Restricted workflow	Provision storage without exposing administrative credentials to users.
Administration	Private TrueNAS surface	Keep infrastructure management outside the public application boundary.
Recovery	Retention and backup planning	Reduce the impact of mistakes, suspension, or service failure.

Security limitations

- The platform depends on correct configuration across multiple services and must be reviewed after changes.
- A mirror cannot substitute for an independent backup or tested restoration procedure.
- Home internet and power availability remain single-site dependencies.
- Authorised users still require strong passwords, safe devices, and appropriate account handling.
- Public details are intentionally high level; operational documentation is maintained separately.

11 / RESILIENCE AND RECOVERY

Availability, snapshots, backups, and restoration

Recovery planning treats disk redundancy, point-in-time recovery, off-system backup, configuration protection, and tested restoration as separate controls.

Protection layers

Disk availability	Two-drive ZFS mirror continues operating after one data-drive failure.
Integrity visibility	ZFS checksums, pool status, and scheduled maintenance expose storage-health problems.
Point-in-time recovery	Snapshots support recovery from selected accidental or unwanted changes when retained.
Configuration recovery	TrueNAS configuration and documented build information support service reconstruction.
Independent backup	Required for protection beyond the NAS; off-system copies and restoration tests remain an ongoing priority.

Example recovery priorities

1. Protect people and hardware, then stop unsafe writes if storage health is uncertain.
2. Record the observed fault and confirm whether it affects a disk, pool, application, network path, or account workflow.
3. Preserve available evidence and avoid destructive repair actions until the recovery path is understood.
4. Restore service from the least disruptive valid layer: account correction, snapshot, application recovery, disk replacement, or independent backup.
5. Validate data access and permissions after recovery, then document the incident and corrective action.

12 / MONITORING AND MAINTENANCE

Keeping a live storage service trustworthy

Operational readiness depends on detecting degradation early, applying maintenance deliberately, and retaining enough evidence to investigate failures.

Monitored conditions

- ZFS pool state, disk errors, scrub outcomes, and capacity consumption.
- System CPU, memory, temperature, storage I/O, and application resource use.
- Application availability, failed jobs, and provisioning workflow results.
- Drive gateway health, authentication failures, and controlled backend errors.
- Backup or replication task outcomes and evidence of successful recovery testing.

Maintenance rhythm

Routine	Review alerts, failed tasks, service availability, and unusual capacity changes.
Scheduled	Run and review ZFS maintenance, update supported software, and inspect storage health.
Periodic	Audit users, roles, ACLs, quotas, external exposure, and recovery procedures.
After change	Validate login, authorisation, provisioning, file operations, monitoring, and rollback readiness.
After incident	Preserve evidence, restore carefully, verify data and permissions, and record lessons learned.

13 / VALIDATION

Testing the complete service path

The platform was not considered operational until storage health, user separation, provisioning, remote access, and recoverability were tested as connected workflows.

Layer	Technology	Engineering role
Hardware	Drive and system checks	All installed devices detected; temperatures, links, and capacity reviewed.
Storage	Mirror validation	Healthy pool state and no reported disk errors during final validation.
Identity	Role scenarios	Approved roles allowed; unapproved and unauthenticated scenarios denied.
Provisioning	End-to-end account test	Dataset, quota, ACL, activation, and notification workflow completed.
Isolation	Cross-user access test	Users limited to their assigned storage boundary.
Files	Remote operations	Upload, download, organisation, and deletion completed through the user experience.
Operations	Jobs and monitoring	Health visibility, maintenance tasks, and backup-oriented controls reviewed.

Acceptance result: The completed service met its original small-scale requirements: resilient primary storage, approved remote users, isolated allocations, automated provisioning, encrypted access, and maintainable operational controls.

14 / ENGINEERING CHALLENGES

Problems that materially changed the design

The most valuable learning came from correcting assumptions about hardware, permissions, automation, routing, and the difference between resilience and backup.

Hardware compatibility correction

An earlier computer appeared suitable from its headline processor, memory, and storage specifications. Physical inspection showed limitations in drive bays, mounting, power, SATA connectivity, clearance, airflow, and future serviceability. Replacing that system increased time and cost but prevented an unreliable final build.

Drive gateway and authentication

The Drive subdomain required hostname-aware routing, controlled cross-domain authentication, safe return destinations, shared session handling, and server-side role enforcement. Treating the subdomain as a distinct service boundary avoided exposing the NAS administration interface.

Repeatable permissions

Creating storage manually for each user would have introduced inconsistent dataset properties, quotas, ownership, and ACLs. A restricted provisioning workflow made the operation repeatable and testable while keeping privileged credentials away from the public user interface.

Operational truth

A healthy mirror can create false confidence if backup, restoration, monitoring, and power-loss scenarios are ignored. The final design therefore treats resilience as one layer of a wider recovery model rather than evidence that data is fully protected.

15 / OPERATIONAL PROCEDURES

Controlled administration without publishing secrets

Public procedures describe decision points and validation requirements; exact commands, credentials, identifiers, and private paths remain in restricted operational records.

Approve and provision a user

1. Confirm the identity, business purpose, required role, quota profile, and retention expectation.
2. Grant the minimum approved application role through the protected administrator interface.
3. Submit the controlled provisioning request and monitor its result.
4. Verify the dataset, quota, ownership, ACL, application association, and active account state.
5. Confirm that the user can reach only the assigned storage and record completion.

Suspend access

1. Confirm the reason, authority, scope, and expected retention period.
2. Remove active Drive access without deleting the underlying user data.
3. Verify that existing sessions and new access attempts are handled correctly.
4. Record the action and schedule the next retention or restoration review.

Respond to a degraded mirror

1. Confirm the reported device and current pool state before touching hardware.
2. Protect current data and verify the independent backup position.
3. Replace only the confirmed failed device using a compatible disk and controlled procedure.
4. Monitor resilvering and system health until the mirror returns to a verified healthy state.
5. Run post-recovery checks and update maintenance records.

16 / RISK REGISTER

Known risks and treatment priorities

The platform is intentionally small, but it is operated with explicit awareness of availability, security, capacity, recovery, and lifecycle risks.

Layer	Technology	Engineering role
Power loss	Unexpected shutdown	Add UPS integration and test graceful shutdown and restart.
Single site	Theft, fire, or major local failure	Maintain independent off-system backups and restoration evidence.
Disk failure	Degraded mirror	Monitor health, retain compatible replacement capability, and follow a controlled resilver procedure.
Capacity	Service disruption or failed jobs	Apply quotas, monitor growth, and set capacity thresholds.
Misconfiguration	Permission or exposure error	Use repeatable workflows, peer-style review, backups, and post-change validation.
Account compromise	Unauthorised data access	Strong authentication, role minimisation, session controls, and periodic access review.
Software change	Regression or incompatibility	Review releases, back up configuration, stage changes, and retain rollback options.

17 / OUTCOMES AND LEARNING

What the project demonstrates

The completed platform provides operational value while evidencing practical engineering skills that extend beyond a laboratory-only build.

Technical outcomes

- Built and commissioned a dedicated TrueNAS SCALE platform with separated system, application, and user-data storage.
- Configured and validated a healthy two-drive ZFS mirror with dataset and quota controls.
- Integrated a branded Drive gateway with authentication and server-side role authorisation.
- Implemented automated isolated-user provisioning and guarded account lifecycle controls.
- Enabled encrypted remote file access without publishing the NAS administration interface.
- Established monitoring, maintenance, backup-oriented tasks, and recovery documentation.

Professional competencies

Layer	Technology	Engineering role
Research	Hardware and platform comparison	Converted requirements into compatible, supportable design decisions.
Problem solving	Physical and software corrections	Identified invalid assumptions and changed the design before operational failure.
Security	Layered access controls	Applied authentication, authorisation, isolation, encryption, and private administration.
Automation	Provisioning workflow	Reduced repetitive privileged work and improved consistency.
Operations	Monitoring and recovery	Treated maintenance and restoration as part of the build, not an afterthought.
Communication	Sanitised documentation	Explained architecture and evidence without exposing operational secrets.

18 / ROADMAP

Continuous improvement

Operational does not mean permanently finished; the next phase focuses on stronger recovery assurance, power protection, alerting, and capacity planning.

- Deploy UPS monitoring and verify graceful shutdown and automatic recovery behaviour.
- Strengthen independent off-system backups and perform scheduled restoration tests.
- Expand alerts for pool health, disk errors, temperatures, capacity thresholds, failed jobs, and service availability.
- Review long-term hardware expansion against chassis, controller, power, cooling, and ZFS topology constraints.
- Maintain supported TrueNAS and application versions through controlled, recoverable updates.
- Conduct periodic access, ACL, quota, external-exposure, and recovery-procedure reviews.
- Update this public document when major architecture or operational milestones change.

> **Architecture decision:** Jellyfin and other temporary media-server experimentation are not part of this private-cloud deliverable. Future media services will be treated as a separate project and infrastructure boundary.

19 / CONCLUSION

A working service and reusable infrastructure platform

The project's value is not limited to approximately 1.76 TiB of usable storage; it lies in the complete process of designing, correcting, securing, automating, validating, and operating the service.

The Self-Hosted Private Cloud Infrastructure demonstrates how a modest business-class computer can become a disciplined, role-authorised storage service when hardware compatibility, filesystem behaviour, permissions, remote access, monitoring, and recovery are treated as connected engineering concerns.

The platform now provides approved users with isolated storage through a professional Drive experience while keeping infrastructure administration private. Its mirrored storage, automated provisioning, encrypted access, operational checks, and documented limitations create a credible foundation for continued learning and carefully controlled self-hosted services.

> **Final assessment:** The original requirements were achieved: the platform is tested, operational, remotely usable by authorised users, and maintainable as a long-term private-cloud and cybersecurity project.

Public links

Project case study	https://meetshawon.com/projects/private-cloud-infrastructure
Portfolio platform	https://meetshawon.com
Drive gateway	https://drive.meetshawon.com - authorised users only
GitHub profile	https://github.com/Shawon1024
Professional profile	https://linkedin.com/in/shawon1024

APPENDIX A

Public glossary

Concise definitions for the storage, identity, and operational concepts used throughout this document.

ACL	Access control list defining permissions for identities and groups.
Dataset	A ZFS-managed logical storage area with its own properties, permissions, and quota.
HTTPS	Encrypted HTTP used to protect browser sessions and file transfers in transit.
Mirror	A storage layout that maintains identical data on two devices for single-device fault tolerance.
NAS	Network-attached storage: a system providing managed storage services over a network.
Nextcloud	A self-hosted application providing browser-based file access and collaboration capabilities.
Quota	A limit applied to storage consumption for a user or dataset.
Resilver	The ZFS process that reconstructs mirrored data onto a replacement device.
Snapshot	A point-in-time ZFS reference useful for selected recovery scenarios.
TrueNAS SCALE	The storage operating platform used to manage pools, datasets, applications, jobs, and health.
ZFS	A checksummed storage filesystem and volume manager supporting datasets, snapshots, quotas, and mirroring.

Publication note

This is the public edition of the technical documentation. Sensitive operating procedures, secrets, private identifiers, exact internal paths, and recovery credentials are intentionally excluded. The document describes the system at a high level as of September 2026 and should be updated after material architecture changes.