Operations Guide

Example Deliverable — 4-Node Hyperconverged Cluster

Document	**Example deliverable (sanitized)**
Prepared by	**Joy Automation**
Version	1.0
Date	July 31, 2026

Sanitized example from a completed engagement — client-identifying details removed. Provided to illustrate the deliverables Joy Automation produces.

Contents

> *Sanitized example deliverable from a completed Joy Automation engagement, included to illustrate the operations documentation you would receive with a delivered system. It describes that project's 4-node, two-site cluster; your delivered guide will reflect your own as-built architecture. Client-identifying details have been removed.*

Audience: the on-site team that operates this cluster day-to-day. **Scope:** how to check health, perform routine operations, and recover from every failure mode we tested. Recovery procedures triggered by a failure are written here as **Runbooks** (§7); the Testing Report references them.

This guide is the operator's front door and is **self-contained for day-2 operation**. The exact as-delivered state — each switch's running-config and every node's network configuration — is captured in your handover package (snapshot at delivery on 2026-06-09, plus a final pre-ship update on 2026-06-11 covering the external-uplink trunks and mgmt NIC bond). Deeper engineering documentation — architecture decisions and rationale, rebuild guide, and Ceph setup/upgrade runbooks — can be provided by Joy Automation on request.

---

1. What this cluster is (60-second version)

Four hyperconverged nodes across **two sites (A: node1/2 · B: node3/4)**, joined by a **10 G ERPS ring** over four FS S3900-R switches. Every node runs:

- **Ceph** (Squid 19.2.4) — shared storage for everything (`incus`, `kube-rbd`, `cnpg`, RGW/S3).
- **Incus** — clusters the VMs (Kubernetes control plane + workers, Ignition SCADA, Windows servers).
- **Kubernetes** (Talos VMs) — runs containerized workloads (TimescaleDB/CNPG, MetalLB, Headlamp).

Design tolerance: any single node can fail with no data loss and automatic recovery. Two-node / full-site loss is a manual-recovery event — see §7.7 and the Testing Report.

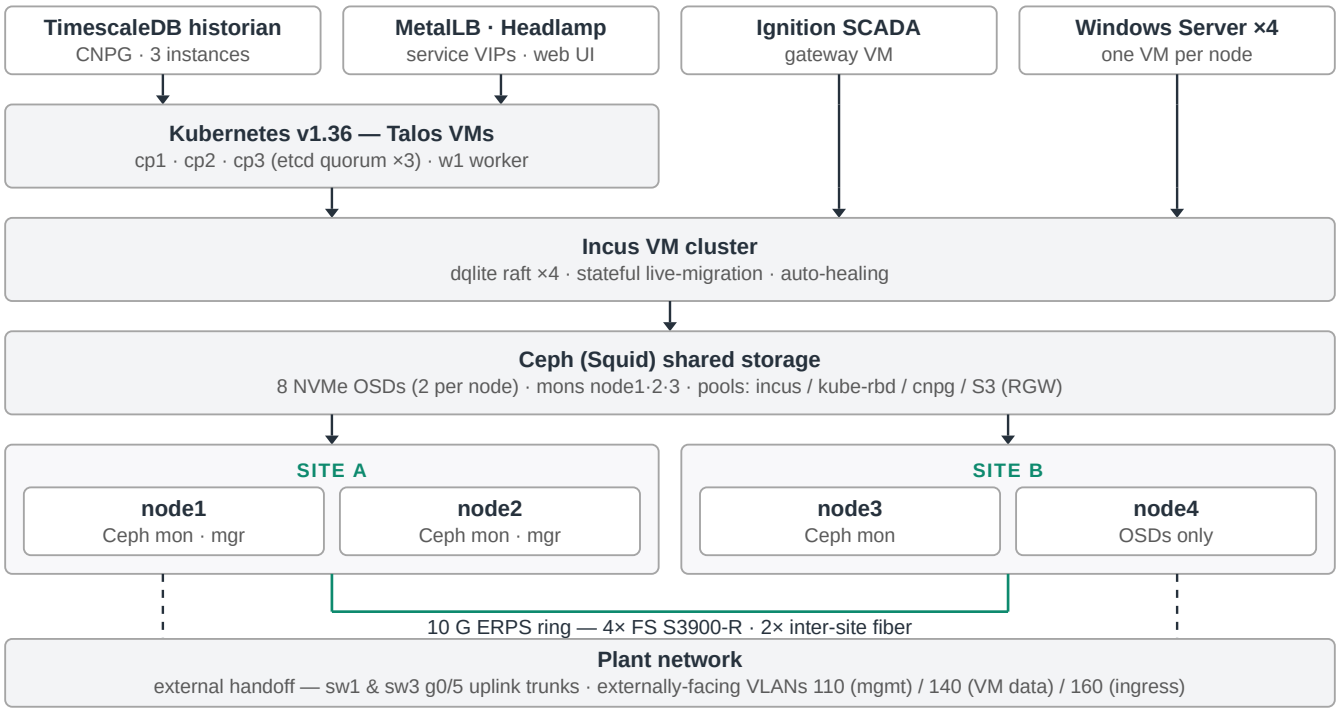

Figure 1: Logical architecture — every node carries every shared layer (hyperconverged); any single node can fail with automatic recovery.

Plane	Hardware / version
Nodes	4× Supermicro **SYS-112B-FWT**, Intel **Xeon 6515P**, board X14SBW-TF
Storage	Ceph **Squid 19.2.4** (cephadm/Docker), 8 OSDs, ~7 TiB raw
Virtualization	**Incus** cluster (dqlite/raft)
Kubernetes	**Talos v1.13.x / k8s v1.36.x** on 4 VMs (3 control-plane + 1 worker)
Switches	4× **FS S3900-24T4S-R**, FSOS 2.2.0F, single ERPS ring

2. Daily / periodic health checks (“is it green?”)

Run from any node (Ceph/Incus) or the admin workstation (k8s). All should be quiet.

```
# --- Ceph (run on node1/2/3) ---
sudo ceph -s                      # HEALTH_OK; mon quorum 3/3 (node1,2,3); 8 osds up & in
sudo ceph osd df tree             # no OSD near-full; site-a/site-b balanced
sudo ceph osd crush tree          # hosts under correct datacenter buckets (no mixed site)

# --- Incus (run on any node) ---
sudo incus cluster list           # 4 members, all ONLINE, none EVACUATED
sudo incus list                   # ignition, talos-cp1/2/3, talos-w1, winsrv01-04 RUNNING

# --- Kubernetes (admin workstation, KUBECONFIG set) ---
kubectl get nodes -o wide         # 4 Ready
kubectl get pods -A | grep -vE 'Running|Completed'   # (empty = healthy)
kubectl -n tsdb1 get cluster tsdb1                   # 3/3 instances healthy
```

```
talosctl -n 172.16.140.51 etcd members              # 3 members

# --- Switch ring (SSH to the FSOS CLI) ---
ssh admin@172.16.150.41          # then at the prompt:  show erps
#   expect: RPL Owner, State Idle, tg0/28 BLK (ring closed & healthy)
```

Green baseline (as delivered): Ceph `HEALTH_OK`, 8/8 OSDs, ~3.8 % raw used; ring `Idle` with RPL blocked on sw1 tg0/28; all VMs and pods Running.

Hardware/thermal spot checks (monthly or on alert): `sudo nvme list` + `sudo smartctl -A /dev/nvmeN` (NVMe < 70 °C), `sudo ipmitool sensor list` (fans/temps), BMC web UI per node.

3. Access & credentials

Plane	How to reach it	Auth	Notes
Node OS (SSH)	`ssh nodeN` (aliases) → `172.16.110.41–.44`	SSH key	Ceph `_admin` on node1/2/3
Node BMC/IPMI	`https://172.16.110.51–.54` (web) or `ipmitool`	password manager	Out-of-band power/console
Ceph dashboard	`https://172.16.110.41:8443` (active mgr; fails to `.42`)	`admin` / **password manager**	Self-signed cert
Ceph S3 (RGW)	`http://172.16.120.50` (ingress VIP)	per-user keys / password manager	Internal (VLAN 120) only
Kubernetes API	`https://172.16.140.50:6443` (Talos VIP)	`~/.kube/config`	Floats across cp1/2/3
Headlamp (k8s viewer)	`http://172.16.160.211`	per Headlamp	Read-only on-site view
TimescaleDB	`172.16.160.210:5432` (CNPG rw VIP)	DB creds / password manager	`tsdb1-rw-lb`
Switches	`ssh admin@172.16.150.41–.44` (FSOS CLI)	`admin` / **password manager**	Password set off the default; stored in the password manager
Windows servers	`sudo ssh -i /etc/winserver/id_ed25519 Administrator@172.16.140.60–.63`	key (root-only)	From a node **other than** the host (macvlan isolation)

Secrets live in the password manager. A few credential files also exist locally on the nodes / admin workstation: `.env` (BMC creds), `k8s/talos/{controlplane,worker,talosconfig,kubeconfig}`, `/etc/winserver/*` (Windows keys).

4. Network reference

Authoritative live capture: the switch running-configs and per-node network configuration in your handover package (snapshot at delivery on 2026-06-09, with a final pre-ship update on 2026-06-11 covering the g0/5 uplink trunks and `bond-mgmt`). **Design rationale (the *why*):** the ring uses a **single ERPS instance** — dual-

instance load-balancing isn't supported on the S3900-**R** firmware (no per-instance protected VLAN). A single instance still keeps the RPL on an inter-site link, with the second inter-site fiber as hot standby (§4.4).

4.1 VLANs

VLAN	Name	Subnet	Carries
110	mgmt	172.16.110.0/24	Node OS mgmt + BMC/IPMI (1 G copper)
120	ceph-public	172.16.120.0/24	Ceph client traffic + RGW ingress VIP `.50`
130	ceph-cluster	172.16.130.0/24	Ceph OSD replication/heartbeat
140	vm-data	172.16.140.0/24	VM & pod data plane (Talos, Ignition, Windows, k8s API VIP)
150	switch-mgmt (OOB)	172.16.150.0/24	Switch management SVIs only — **never on ring/node trunks**
151	erps-raps	—	ERPS R-APS control VLAN
152	reserved	—	Unused (spare ERPS instance)
160	k8s-ingress	172.16.160.0/24	MetalLB LoadBalancer pool (`.200–.250`)

4.2 IP allocations (key reservations)

Address	Role
172.16.110.41–.44	Node OS mgmt (node1–4)
172.16.110.51–.54	Node **BMC/IPMI** (node1–4)
172.16.120.50	**Ceph RGW/S3 ingress VIP**
172.16.140.41–.44	Node host IPs on `vm-data`
172.16.140.50	**Kubernetes API VIP** (Talos) — never reuse
172.16.140.51–.54	Talos VMs (cp1, cp2, cp3, w1)
172.16.140.60–.63	**Windows servers** (winsrv01–04)
172.16.140.64–.199	free for additional Windows/VM guests
172.16.150.41–.44	Switch mgmt SVIs (sw1–4)
172.16.160.210	TimescaleDB rw VIP (`tsdb1-rw-lb`)
172.16.160.211	Headlamp
172.16.160.200–.250	MetalLB pool range

> **VLAN 140 caution:** Windows/VM static IPs share the segment with the k8s data plane. `create-winserver` hard-refuses `.1–.59`. Never reuse `.50` (API VIP) or `.51–.54` (Talos).

4.3 Switch port map (all 4 switches identical)

Port	Use	VLAN
g0/1	OOB switch-mgmt uplink	150
g0/2, g0/3	Node NIC1 / NIC2	110
g0/4	Node BMC	110
g0/25–26	Node SFP28 **LACP bond** (`p0`/`p1`)	trunk 120/130/140/160
tg0/27, tg0/28	**Ring trunk** ports	ring
g0/5	**External-uplink trunk** (sw1 & sw3, one per site) — plant-network handoff; carries only the externally-facing VLANs. mira1 interim router rides sw1's today (**REMOVE at site**)	trunk 110/140/160

4.4 ERPS ring

Single instance, control VLAN **151**. Ring order **SW1 → SW2 → SW3 → SW4 → SW1**:

Link	Ports	Path
L1	SW1:tg27 ↔ SW2:tg28	intra-site A (DAC/short SR)
L2	SW2:tg27 ↔ SW3:tg28	**inter-site fiber #1**
L3	SW3:tg27 ↔ SW4:tg28	intra-site B
L4	SW4:tg27 ↔ SW1:tg28	**inter-site fiber #2** (RPL)

- **RPL Owner = sw1** (`tg0/28` blocked in steady state). **RPL Neighbour = sw4** (`tg0/27`).
- Steady state: `show erps` → **State Idle, tg0/28 BLK**. Convergence on a break: <50 ms intra-site, ~220 ms inter-site (measured, Testing Report §T3).
- **At the site, keep L2 and L4 as the inter-site fibers** so the RPL stays on an inter-site link.

4.5 Node networking (per node)

- **Bond:** 2× SFP28 in **LACP 802.3ad** (`bond-data`, `xmit_hash_policy layer3+4`, `lacp_rate fast`). VLAN sub-interfaces ride the bond: `ceph-public` (120), `ceph-cluster` (130), `vm-data` (140), `bond-data.160` (k8s ingress). `br-ceph` bridges Ceph to the host.
- **mgmt:** `bond-mgmt` — `eno1`+`eno2` **active-backup** (primary `eno1`, presents eno1's MAC) on VLAN 110 (`.41`–`.44`). Failover/failback verified ≲1 s; a single mgmt NIC, cable, or switch port can no longer isolate Incus liveness from a node whose storage is healthy.
- VM NICs attach via **macvlan** on `vm-data`. (A host cannot reach *its own* macvlan guest — always manage a VM from a different node.)

> **Interim router (mira1):** today VLAN gateways + NAT egress + DHCP reservations run on the `mira1` bring-up router (an interim bench device). **Remove it at the production site** and provide VLAN gateways / DHCP / internet egress from plant infrastructure. VLAN 140 needs internet egress + working DNS for Windows MAK activation (the cluster uses public DNS 1.1.1.1 for that — the VLAN gateway is **not** a DNS server).

---

5. Storage layout (Ceph)

`fsid <cluster-fsid>` · mon on **node1/2/3** (not node4) · 8 OSDs · CRUSH datacenter buckets **site-a (node1,2) / site-b (node3,4)**. Deployed with cephadm on Docker, size=3, host failure domain.

Pool	size/min	Used by
`incus`	3 / 2	Incus VM disks
`kube-rbd`	3 / 2	Kubernetes PVs (ceph-csi, SC `ceph-rbd`)
`cnpg`	**2 / 1**	TimescaleDB PVs (SC `ceph-rbd-cnpg`) — **deliberately size=2**: the DB already keeps 3 streaming replicas, so size=3 would triple-pay on top of DB replication; size=1 is unsafe (Ceph scatters every volume across all OSDs, so one disk loss corrupts every volume at once)
`.rgw.*`	3 / 2	S3 object store (TimescaleDB backups → `s3://cnpg-backups`)

RBD images use **exclusive-lock** as a storage-level fence — an *unresponsive* stale writer is blocklisted when the survivor force-breaks the lock. (A still-reachable writer yields the lock cooperatively instead, which is why

a healed-from node that never rebooted must be power-cycled — §7.1 step 4.) VMs also carry rolling snapshots (`vm-snapshots` profile).

---

6. Routine operations

6.1 Planned node maintenance (reboot, hardware, patching)

```
sudo incus cluster evacuate <node>     # live-migrates stateful VMs off (~1.5 s each)
# ... do the work / reboot ...
sudo incus cluster restore <node>      # migrates guests back
```

Wait for `ceph -s` HEALTH_OK between nodes. **One node at a time.** Same pattern for switches (reboot one at a time; the ring tolerates a single switch loss).

6.2 Graceful full shutdown (e.g., to relocate / box up)

1. Shut down Windows guests cleanly (`incus stop winsrv0N`).
2. `sudo incus stop --all` (or stop Talos/Ignition VMs) on each node.
3. `sudo ceph osd set noout` (prevent rebalancing while down).
4. Power off nodes, then switches. If a node hangs in shutdown (>5 min — observed once on node3), power it off at the BMC: with VMs stopped and `noout` set this is safe.
5. Reverse to bring up — **see §7.8 cold-start**; `ceph osd unset noout` once all OSDs are back.

6.3 Adding a Windows server

```
ssh node1
sudo create-winserver winsrvNN --node nodeX --ip 172.16.140.6X/24 --gateway 172.16.140.1
# ~2 min later it's reachable; reboot once to apply the hostname
```

Stay at **≤1 Windows VM per node** (no Software Assurance → host-bound licenses). See §9.

---

7. Runbooks (failure response)

> Decision tree: **one node/disk/switch down** → automatic, just verify & restore (7.1–7.5). **Two nodes / a whole site down** → manual recovery (7.7). **Everything off** → cold start (7.8).

7.1 Unplanned single-node loss

Symptom: a node is unreachable/dead. **Expected:** Ceph stays `HEALTH_OK` (degraded then self-heals); Incus auto-heals its VMs onto survivors (~40 s) with the same IPs; k8s reschedules. **Do:**

1. Confirm the rest is green (§2). Workloads should already be back.
2. The node **self-reboots** reliably in our testing; when it returns: `sudo incus cluster restore <node>` to move guests home.
3. If it doesn't return: power-cycle via BMC (`ipmitool -H 172.16.110.5X ... chassis power cycle`).
4. **If the node never went down** — it's still up but Incus healed VMs off it (a management-path partition, not a crash): **power-cycle it via BMC *before* it rejoins.** A healed-from node that didn't reboot may still be running stale copies of the relocated guests; never let it rejoin without a reboot.

7.2 Planned maintenance — see §6.1.

7.3 VM corruption / bad change — roll back a snapshot

```
sudo incus snapshot list <vm>
sudo incus snapshot restore <vm> <auto-YYYY-MM-DD-HHMM>
```

7.4 Single OSD / disk failure

Symptom: `ceph -s` shows an OSD down; `HEALTH_WARN`, data re-replicates automatically. **Do:** replace the drive, then `sudo ceph orch osd rm <id> --replace` (the all-available-devices spec re-creates the OSD on the new disk). No data loss (size≥2).

7.5 Switch / ring / inter-site link failure

Symptom: one switch or ring link down. **Expected:** ERPS reconverges (<50 ms intra / ~220 ms inter-site); the directly-attached node loses its bond path and behaves like 7.1. **Do:** verify the ring re-formed (`ssh admin@172.16.150.4N → show erps`); replace/restore the switch or fiber; ring auto-recovers (WTR).
Firmware (FSOS): upgrade **one switch at a time** (the ring tolerates a single switch down) — back up its config first (`copy running-config startup-config`), flash the `.bin` via the web UI (File Management → Copy → Boot) or CLI, reboot, and confirm `show erps` rejoins Idle before moving to the next. Leave the bootrom/Loader alone.

7.6 Kubernetes API VIP / MetalLB LoadBalancer unreachable

- **API VIP (`.50`) down:** check `talosctl -n 172.16.140.51 etcd members` (need 3). If a lone CP survivor holds the VIP but the API is wedged: `talosctl -n <cp-ip> service kubelet restart` (the apiserver-stuck-on-VIP failure mode — see §7.7).
- **LB VIP (160.x) down:** `kubectl -n metallb-system get pods` (speaker on all nodes). CP nodes are kept as announcers via the `lb-announce-cp` patch — if reverted, only the worker announces and failover stalls.

7.7 Two-node / full-site loss (manual recovery)

This is the cluster's hard limit — Ceph loses quorum/PGs and Incus loses dqlite quorum, so workloads stop until nodes return. **RTO ~5–15 min.** Recovery:

1. Wait for the downed nodes to self-reboot (~3 min) **or** power-cycle via BMC (`ipmitool -I lanplus -H 172.16.110.5N -U <user> -E chassis power cycle`).
2. As they rejoin: `sudo ceph -s` until HEALTH_OK / mon 3/3, then `sudo incus cluster restore <node>` for each evacuated node (migrates VMs home). If a down-site VM stays STOPPED with "storage pool unavailable", wait for HEALTH_OK then `sudo incus start <vm>`.
3. **Check the k8s API VIP** (`172.16.140.50:6443`). If it stays down *even though* etcd is healthy (`talosctl -n 172.16.140.51 etcd members` shows 3), the VIP-holding CP node's apiserver is wedged — restart its kubelet: `talosctl -n <vip-holder> service kubelet restart` (rebuilds the static pod; the VIP serves immediately).
4. Clear stale pods: `kubectl delete pod --field-selector status.phase=Failed -A` (and any `ContainerStatusUnknown`).

> Healing may relocate the downed site's Talos VMs onto survivors during the outage; `incus cluster restore` puts them home (Talos co-location is temporary and safe).

7.8 Full cold start (all power lost / after a planned shutdown)

VMs auto-start once Ceph + the Incus cluster are ready (hardened via the `incus-startup-wait-for-ceph` drop-in, which waits for both `rbd ls` and `incus list` before autostarting guests). Expect **~4–10 min** to green. Steps:

1. Power on all 4 nodes (BMC `... chassis power on` if needed); they reach SSH in ~30–60 s.
2. Watch `ssh node1 'sudo ceph -s; sudo incus list'` — HEALTH_OK and VMs RUNNING ~40 s after the cluster forms. **Only if** members show EVACUATED (`incus cluster list`): `for n in node1 node2 node3 node4; do sudo incus cluster restore $n --force; done`.
3. k8s forms as Talos boots (~2–3 min); if the API VIP stays down despite healthy etcd, do the kubelet restart from §7.7. Force-delete any `ContainerStatusUnknown`/`Error` pods.
4. `ceph osd unset noout` if it was set.

> **Planned shutdown tip:** *before* powering off, set `incus config set cluster.healing_threshold=0` so healing doesn't mark every member evacuated on the way down (restore it to `20` after bring-up).

7.9 Data restore drills

- **TimescaleDB (CNPG) from S3 backup:** bootstrap a recovery CNPG cluster from `s3://cnpg-backups` (base + WAL); in testing a fresh DB restored from backup alone in ~100 s.
- **Ignition gateway:** `gwcmd` backup/restore of the `.gwbk` gateway backup.
- Both mechanisms are **verified**; re-validate fidelity once real production data exists (Testing Report §T5).

8. Monitoring & dashboards

- **Ceph dashboard** — `https://172.16.110.41:8443` (health, OSDs, pools, RGW, capacity).
- **Headlamp** — `http://172.16.160.211` (read-only k8s view for the on-site team).
- **CLI** is the source of truth (§2). No external alerting is wired yet — **recommended follow-up:** point plant monitoring at `ceph -s`, `incus cluster list`, and the ERPS state, and alert on NVMe temp / OSD near-full / mon-quorum-loss.

9. Handover / security actions (do at commissioning)

#	Action	Why
1	**Configure plant NTP** — set the `server` in `/etc/chrony/chrony.conf` to the plant time source (node1 is the interim master)	Ceph mons need <0.05 s skew; no good source → `MON_CLOCK_SKEW`
2	**Remove the mira1 interim router**; provide VLAN gateways, DHCP, DNS, internet egress from plant infra (plant uplinks land on the **sw1 & sw3 g0/5** external trunks, VLANs 110/140/160; VLAN 140 needs egress + DNS for Windows activation)	Interim only (mira1 bring-up router)
3	**Patch all four switch g0/1 OOB ports into the plant management network** and route `172.16.150.0/24` at the FortiGate	Remote switch management that survives a fabric fault — VLAN 150 is deliberately **not** on the ring or uplink trunks (out-of-band by design)
4	**Confirm Windows CALs** with your reseller (CDW/MPSA)	Core licenses ≠ CALs — see §9.1
5	Confirm BMC, dashboard, DB passwords are set/rotated and stored in the password manager	Production hygiene

#	Action	Why
6	Set Ceph `mon` capacity alerting; wire plant monitoring (§8)	Day-2 visibility
✓	~~Change switch admin password~~ — **done** (random pw stored in the password manager; default `admin` no longer works)	
✓	~~Rotate the Ceph dashboard RGW key~~ — **done** (rotated to a fresh key; leaked key removed)	
✓	~~Remove the `failure-test` demo app~~ — **done**	

> **Time sync (chrony):** all 4 nodes run chrony; node1 is the interim master (`local stratum 8`, serves `172.16.110.0/24`), node2/3/4 sync to it (`server 172.16.110.41`). With no reachable upstream the cluster stays *mutually* consistent (clears Ceph skew); add a plant NTP `server` at the site for accurate, redundant time. The mira1 router's NAT now auto-detects its uplink (wifi/wired) so an uplink failover no longer silently kills cluster egress.

9.1 Windows licensing & CALs (turnover note)

- **Server licenses (have):** 6× 16-core Windows Server 2025 **Standard**, **MAK**-activated, no Software Assurance. This 1-Windows-VM-per-node layout needs **4** (one fully-core-licensed node per Windows host); you have 6 (2 more on order to reach the 8 for 2/node). **No SA → licenses are host-bound; keep ≤1 Windows VM per node** so a single failover never pushes a survivor past the 2-OSE Standard limit.
- **CALs (separate, likely needed once in use):** server licenses cover *running* the OS; **CALs cover clients *accessing* it** (file/print/AD auth, etc.) — User or Device CALs.
  - **Exempt:** up to **2 administrators** managing the box; purely **anonymous internet** access.
  - **RDS CALs** are a *separate* license for Remote-Desktop *session use* (not the built-in 2-session admin RDP) — and unlike general CALs they are **technically enforced** (120-day grace, then non-admin RDP is refused).
  - **If unlicensed:** general CALs have **no technical enforcement** (audit/compliance risk only); RDS CALs **do** (service stops after grace). For these admin-only boxes you likely need **none today** — buy CALs when a real workload with authenticated users lands, via CDW/MPSA.
  - *This describes standard Microsoft terms, not a compliance ruling — confirm count with your reseller.*

---

10. Known limitations & follow-ups

Item	Status / action
No two-node / full-site survival	By design — two equal sites can't form a quorum (Ceph mons, etcd, Incus dqlite each need a majority). Removing the limit needs a **3rd failure domain** (a small tiebreaker box/VM in a separate location) plus Ceph stretch mode (`size=4`, 2 copies/site, 5th tiebreaker mon, ~25 % capacity cost), a 3rd-site etcd member, and a 3rd-site Incus voter. A documented forward path if full-site HA is later required.
Restore fidelity unvalidated	Mechanisms tested; re-run CNPG + Ignition restores against real historian data post-commission.

Item	Status / action
CNPG Barman native deprecated	Migrate backups to the Barman Cloud **Plugin** before bumping the CNPG operator to 1.30 (the in-tree Barman backup is deprecated upstream).
No external alerting	Wire plant monitoring (§8).
Default switch creds / RGW key / mira1	See §9 handover actions.

11. References — manufacturer & project documentation

Authoritative documentation for every hardware component and third-party software package in the platform. Versions listed are as-shipped; check the vendor site for the current release before upgrading.

Compute & operating system

Component (as-shipped)	Role	Documentation
Supermicro **SYS-112B-FWT** (board X14SBW-TF, onboard BMC/IPMI)	Cluster node server ×4	Support / download center · Product page
Intel **Xeon 6515P**	Node CPU	Intel ARK specifications
Ubuntu Server 24.04 LTS	Node operating system	Ubuntu Server documentation
chrony	Time synchronization (NTP)	chrony documentation

Network

Component (as-shipped)	Role	Documentation
FS **S3900-24T4S-R** (FSOS 2.2.0F)	Access + ring switches ×4	Configuration guide · CLI reference · Web management
ERPS (ITU-T G.8032)	Ring protection protocol	ITU-T G.8032 recommendation
Fortinet **FortiGate 121G**	Site firewall / routed HA gateway	Fortinet Document Library · 120G/121G QuickStart

Storage (Ceph)

Component (as-shipped)	Role	Documentation
Ceph Squid 19.2.4	Distributed block/object storage	Ceph Squid docs
cephadm	Ceph deployment / orchestration	cephadm docs
ceph-csi (RBD)	Kubernetes storage driver	ceph-csi (GitHub)
Docker Engine	Ceph container runtime	Docker Engine docs

Virtualization

Component (as-shipped)	Role	Documentation
Incus	VM / container cluster manager	Incus documentation
distrobuilder	Windows golden-image builder	distrobuilder docs
cloudbase-init	Windows guest first-boot config	cloudbase-init docs

Kubernetes platform

Component (as-shipped)	Role	Documentation
Talos Linux v1.13	Immutable Kubernetes node OS (VMs)	Talos v1.13 docs
Kubernetes v1.36	Container orchestration	Kubernetes documentation
MetalLB	LoadBalancer (L2 mode)	MetalLB documentation
Headlamp	Kubernetes web UI	Headlamp documentation

Database & applications

Component (as-shipped)	Role	Documentation
CloudNativePG	PostgreSQL Kubernetes operator	CloudNativePG documentation
TimescaleDB	Time-series extension (historian)	Timescale documentation
PostgreSQL 17	Database engine	PostgreSQL 17 docs
Inductive Automation **Ignition 8.3**	SCADA platform	Ignition 8.3 user manual
Microsoft **Windows Server 2025**	Windows guest OS	Windows Server documentation

---

Companion deliverable: the Testing Report (failure-injection campaign results). Full engineering documentation — architecture decisions and rationale, rebuild guide, Ceph setup/upgrade runbooks, detailed network design, and the settings register — can be provided by Joy Automation on request.