

Integrated Technical Design, Configuration, and Handover Report

*Company 1, Company 2, MSP Operations, Shared Storage, and
Cross-Site Protection*

Document Type: Formal Integrated Technical Handover

Course: 26W_CST8248 Emerging Technologies

Professor: Denis Latremouille

Prepared By: Raspberry Pioneers

Team Members: Bailey Kulla, Elyazid Sidelkheir, Ru Wang,
Justin Rosseleve, Yiqin Huang, Omer Deniz

Prepared For: Client IT Administrators, MSP Support Staff, and
Successor Operators

Submission Type: Group Submission

Demo Date: April 11, 2026

Contents

List of Figures	vi
List of Tables	ix
Executive Summary	xii
1. Introduction	1
2. Background.....	2
2.1 Project Context and Intended Audience.....	2
2.2 Design Context and Operating Model.....	2
2.3 Evidence Basis and Reporting Method	3
3. Discussion	4
3.1 Shared MSP Boundary, Network Architecture, and Platform	4
Private-cloud topology, segmentation, and virtualization baseline	4
Network Design and IP Addressing Rationale	6
Identity Infrastructure and Tenant Separation	7
Platform and Service Choice Rationale	8
Physical Server Baseline	9
Physical Switching and Uplink Baseline.....	11
Compute and Virtualization.....	11
Public-cloud boundary, MSP entry, and gateway model.....	15
Network Design, IP Addressing, and Segmentation Rationale	17
MSP Entry, Remote Access, and Gateway Design.....	20
3.2 Company 1 Services.....	23
Configuration reference retained from the private-cloud handover	23
Observed service state, client access, and platform evidence.....	26
Service Overview.....	26
Architectural Rationale.....	28
Observed Service State.....	29

Service Composition	34
3.3 Company 2 Services	35
Configuration reference retained from the private-cloud handover	35
Identity, DNS, client delivery, and namespace evidence.....	38
Service Overview.....	38
Architectural Rationale.....	40
Observed Service State	40
Service Composition	48
3.4 Shared Storage, SAN, and File Services	49
Private-cloud storage and SAN baseline	49
Public-cloud storage and isolated SAN evidence.....	56
3.5 Backup, Recovery, and Cross-Site Protection	58
Private-cloud backup baseline	58
Agentless System-Level VM Backup	60
File-Based Client Backup	61
Offsite Backup Copy to Site 2	61
Public-cloud offsite and repository evidence	63
Inter-site VPN protection path.....	66
3.6 Administrative Access, Dependencies, Maintenance, and Triage	70
Administrative access, risks, and support reading from the private-cloud handover	70
Maintenance, Change Workflow, and Failure Domains	70
Standard Change Workflow	71
Failure Domains and Recovery Priority.....	71
Risks, Deferred Improvements, and Support Assumptions.....	72
Service dependency, routine checks, and troubleshooting from the Site 2 handover	73

Maintenance, Daily Duties, and Operational Checks	74
Troubleshooting Guide	75
3.7 Value-Added Operational Tooling and Public Service Extension	78
Operational tooling and management support from the private-cloud handover	78
Standardized Endpoint Policy Control (Company 1).....	78
Infrastructure Monitoring Dashboard (Grafana)	79
Browser-Based Linux Administration (Cockpit).....	81
Centralized Windows Administration (Windows Admin Center)	81
Cloud-hosted public service extension from the Site 2 handover.....	83
Value-Add Overview	83
Architectural Rationale.....	83
Infrastructure Components and Deployment Flow	85
Observed Deployment State.....	86
Dependency Notes and Operational Considerations	87
4. Conclusion	89
5. Appendices	90
Appendix A. Addressing, Inventory, and Endpoint Reference.....	91
Appendix B. DHCP, Storage, and Backup Reference.....	96
Appendix C. Operational Handover, Verification, and Triage.....	98
Network and Shared Infrastructure	98
Company 1	98
Company 2	99
Cross-Site and Remote Access	99
Operational Handover Use	100
Appendix D. C2FS Samba Configuration (sanitized excerpt).....	104
Appendix E. Resolved Gaps and Troubleshooting History.....	105

Appendix F. Supplemental Validation and Escalation Notes.....	106
Appendix G. Requirements Traceability Matrix	107
6. References	109

List of Figures

Figure 1. Overall Site 1 network topology	5
Figure 2. VLAN routing and firewall flow design.....	6
Figure 3. Directory services architecture.....	8
Figure 4. Site 1 rack installation (rear view)	10
Figure 5. Site 1 rack installation (front view)	11
Figure 6. Proxmox VE management interface	13
Figure 7. Virtual machine logical layout.....	15
Figure 8. Site 2 network topology and service-role alignment.....	16
Figure 9. OPNsense edge publication, interface zones, and alias structure	21
Figure 10. OPNsense OpenVPN inter-site tunnel and backup copy firewall rules.....	22
Figure 11. C1-DC1 Windows PowerShell validation for Active Directory, DNS, and DHCP	26
Figure 12. Company 1 services visible from the MSP management path (MSPUbuntuJump port checks)	28
Figure 13. Jump64 Windows bastion baseline	29
Figure 14. C1DC1 and C1DC2 directory service state from authenticated tenant sessions.....	30
Figure 15. C1FS storage volume, SMB shares, and iSCSI session	31
Figure 16. C1WebServer local workgroup status and IIS binding evidence	32
Figure 17. C1WindowsClient local domain membership, DNS, and dual-web access	33
Figure 18. C1UbuntuClient user-session realm state and dual-web access	34
Figure 19. C2-DC1 primary-node DHCP failover, workflow accounts, and group evidence ..	36
Figure 20. C2-DC1 local DNS resolution and recursive lookup validation	37
Figure 21. C2-DC2 secondary-node DHCP failover and DNS resolution validation	37
Figure 22. C2IdM1 Active Directory, DNS, and DHCP evidence.....	41
Figure 23. C2IdM2 Active Directory, DNS, and DHCP evidence.....	42

Figure 24. Cross-domain DNS visibility without AD trust.....	43
Figure 25. C2FS iSCSI-backed storage, mounted volume, and share definitions	44
Figure 26. C2FS synchronization log and directory structure.....	45
Figure 27. C2LinuxClient domain identity, resolver state, and corrected service access	48
Figure 28. C2WebServer nginx hostname binding and raw-IP behavior	48
Figure 29. Server2 storage volume presentation	53
Figure 30. Server2 network interface layout.....	53
Figure 31. C1-Client2 domain-user session and Company 1 share mount validation	54
Figure 32. C2-Client1 domain-user session and Company 2 share mount validation	54
Figure 33. C1-DC1 Windows PowerShell validation for DFS namespace and SMB share publication.....	55
Figure 34. C1-Client1 Windows PowerShell validation for iSCSI session and mounted SAN volume.....	56
Figure 35. C1SAN isolated storage bridge interface evidence	56
Figure 36. C2SAN isolated storage bridge interface evidence	57
Figure 37. Veeam backup overview	60
Figure 38. Primary Veeam job status	61
Figure 39. Site 1 to Site 2 backup copy job	62
Figure 40. Remote backup copy inventory.....	63
Figure 41. S2Veeam repository, job inventory, and offsite-copy configuration	64
Figure 42. Site 1 OpenVPN client status.....	66
Figure 43. Site 2 OpenVPN server status	66
Figure 44. Firewall-centered security segmentation.....	69
Figure 45. Group Policy desktop branding evidence	79
Figure 46. Grafana infrastructure dashboard	80
Figure 47. Ubuntu jumpbox CLI validation for Grafana and InfluxDB endpoints.....	80
Figure 48. Cockpit administration overview.....	81

Figure 49. Cockpit web terminal.....	81
Figure 50. Windows Admin Center connection inventory	82
Figure 51. Windows Admin Center server overview.....	83

List of Tables

Table 1. VLAN and gateway design	6
Table 2. Identity services summary	7
Table 3. Platform and service choice rationale	8
Table 4. Service stacking rationale	8
Table 5. Physical server hardware summary	9
Table 6. Management endpoints	12
Table 7. Company 1 virtual machine inventory	13
Table 8. Company 2 virtual machine inventory	13
Table 9. Shared administrative systems	14
Table 10. Site 2 network segments and gateways	17
Table 11. OPNsense firewall policy and alias summary	18
Table 12. MSP systems and platform summary	22
Table 13. Company 1 domain controllers	23
Table 14. Company 1 DHCP scope summary	23
Table 15. Company 1 service and platform summary	26
Table 16. Company 2 domain controllers	35
Table 17. Company 2 DHCP scope summary	35
Table 18. Company 2 service stack summary	35
Table 19. Company 2 identity, DNS, and DHCP summary	38
Table 20. Company 2 service summary	38
Table 21. Company 1 file service access control summary	49
Table 22. Company 2 replicated file service summary	51
Table 23. Server2 volume layout	52
Table 24. Company 1 iSCSI targets	52
Table 25. Company 2 iSCSI targets	52

Table 26. SAN VLAN design.....	52
Table 27. Storage server SAN interfaces.....	52
Table 28. Storage architecture summary.....	57
Table 29. Backup infrastructure components.....	58
Table 30. Backup repository layout.....	59
Table 31. Backup scope by class.....	59
Table 32. Backup and protection summary.....	64
Table 33. Inter-site VPN routing and firewall control summary.....	66
Table 34. Recommended operational checks.....	70
Table 35. Primary failure domains.....	71
Table 36. Risks, deferred improvements, and support assumptions.....	72
Table 37. Service dependency map.....	73
Table 38. Routine operational checks.....	74
Table 39. Fast triage guide.....	76
Table 40. Group Policy branding configuration summary.....	78
Table 41. Grafana dashboard metrics.....	79
Table 42. Cockpit management features.....	81
Table 43. Windows Admin Center managed systems.....	82
Table 44. Windows Admin Center capabilities.....	82
Table 45. Management endpoints.....	91
Table 46. Full VLAN addressing matrix.....	91
Table 47. Appendix A addressing reference.....	92
Table 48. Documented DHCP scope.....	96
Table 49. Storage server interfaces.....	96
Table 50. iSCSI target mappings.....	96
Table 51. Backup repository volumes.....	96

Table 52. Veeam backup file types	96
Table 53. Routine operational checks	100
Table 54. Fast triage guide	101
Table 55. Project requirement to integrated-report evidence map.....	107

Executive Summary

This report combines the updated Company 1 and Company 2 handover sources into one support-ready record while preserving the configuration detail, validation evidence, tables, and figures needed for ongoing operations. The combined environment should be understood as one managed service offering with two client environments, approved remote management channels, shared support dependencies, and data protection across both sites.

Company 1 and Company 2 use different service platforms and delivery methods, but both are supported through the same operating model: separated client networks, controlled administration paths, protected data handling, and automated backup and recovery processes. The merged report keeps the detail needed for troubleshooting without turning the document into a build log.

For decision-making purposes, the receiving support team should continue to protect client separation first, keep remote support paths controlled second, and treat continuity of data protection as a shared business priority before expanding the environment further.

1. Introduction

The purpose of this document is to merge the two updated handover packages into one professional technical report without losing the service-by-service configuration detail that support staff need after project delivery. It is written to satisfy the final technical-report rubric while remaining practical for onboarding, troubleshooting, and day-two operations.

The discussion section first establishes the shared MSP boundary, network design, and platform baseline. It then separates Company 1 services, Company 2 services, shared storage and SAN handling, backup and cross-site protection, operational access plus triage guidance, and the value-added services that extend the environment beyond the core project scope. The appendices gather the addressing, reference tables, verification aids, and supplemental notes that are most useful during support handover.

Sensitive credentials, private keys, and other secrets are intentionally excluded. This document is a configuration and operations handover, not a credential vault.

2. Background

2.1 Project Context and Intended Audience

The project brief defines the team as the managed service provider responsible for delivering a hybrid private/public service environment for two different client organizations. The intended reader is therefore a client administrator, MSP support technician, or successor operator who needs enough design context to maintain the environment without relying on build-time memory.

Company 1 is presented in the project brief as a newly launched, mission-driven organization that prefers a professionally supported environment, predictable pricing, centralized identity and access control, and hybrid-cloud operations that remain aligned with Canadian compliance expectations. By retaining the primary storage and backup control path on the Site 1 physical infrastructure located in Canada, the delivered design also aligns with Company 1's preference for Canadian data handling and controlled administrative governance.

Company 2 is presented as a budget-conscious NGO that prefers open-source or license-free platforms where possible, wants to minimize vendor lock-in, and still requires secure hybrid-cloud operations, fault-tolerant core services, and supportable file, identity, and backup workflows.

For readability, Lumora is referred to as Company 1 and ClearRoots is referred to as Company 2 throughout the remainder of this report.

Before exploring the detailed service sections, the reader should treat tenant isolation, approved management entry points, and recoverable operations as the assumptions that every later chapter depends on. Those assumptions matter more than any single VM or host because they explain why routing, identity, storage, and backup were arranged the way they were.

2.2 Design Context and Operating Model

The combined environment is not a flat inventory of disconnected machines. It is a managed service model with an MSP boundary, two tenant stacks, a shared recovery path, and explicit separation between production LANs, administrative interfaces, DMZ publication, and storage network paths. The private-cloud portion carries the heavier Windows-service baseline, while the second location extends the design through Linux-heavy tenant services and an additional public-cloud delivery component.

A reader approaching this report for the first time should treat VLAN segmentation, tenant identity isolation, approved bastion entry, and recoverable storage handling as the load-bearing assumptions that every later section depends on. Those choices explain why some services are centralized, why others are duplicated per tenant, and why the report keeps storage and backup discussion close to operational access and failure-domain guidance.

2.3 Evidence Basis and Reporting Method

The merged document preserves both source handovers because they serve different strengths. The private-cloud source handover contributes the richer shared-platform, storage, backup, and Windows-service configuration record. The second source handover contributes the stronger source-based observation method, management-path evidence, dependency mapping, and troubleshooting guidance.

Where a configuration or operating statement appears below, it is drawn from the delivered environment evidence already captured in the updated source documents unless it is explicitly marked as a risk, deferred improvement, or operating assumption. That approach keeps the report aligned to the rubric requirement that the reader be told what was actually built, why it matters, and how it should be supported.

3. Discussion

3.1 Shared MSP Boundary, Network Architecture, and Platform

This section keeps the shared platform baseline, segmentation model, and MSP-facing boundary material together so the reader can understand the environment as one managed service architecture before moving into the tenant-specific sections.

Private-cloud topology, segmentation, and virtualization baseline

The environment is organized around a single Proxmox virtualization host, a centralized OPNsense firewall, and a dedicated Windows storage server (Server2). Architecturally, Site 1 follows a centralized control model in which OPNsense provides routing and policy enforcement, Proxmox hosts the compute layer, and Server2 delivers shared storage and backup services. Tenant isolation is achieved through separate client, server, DMZ, and SAN VLANs for each company, with a restricted management VLAN used for administrative access. Storage access for Company 1 is presented from Server2 over dedicated SAN interfaces, and systems that require block storage use dedicated SAN NICs rather than relying on the routed client path. In addition to the local Site 1 design, Site 1 and Site 2 are connected by a site-to-site OpenVPN tunnel that supports controlled cross-site management and offsite Veeam backup copy traffic. Although only one tunnel instance is used, Company 1 and Company 2 remain logically separated across that inter-site path through routed subnet definitions and OPNsense firewall policy.

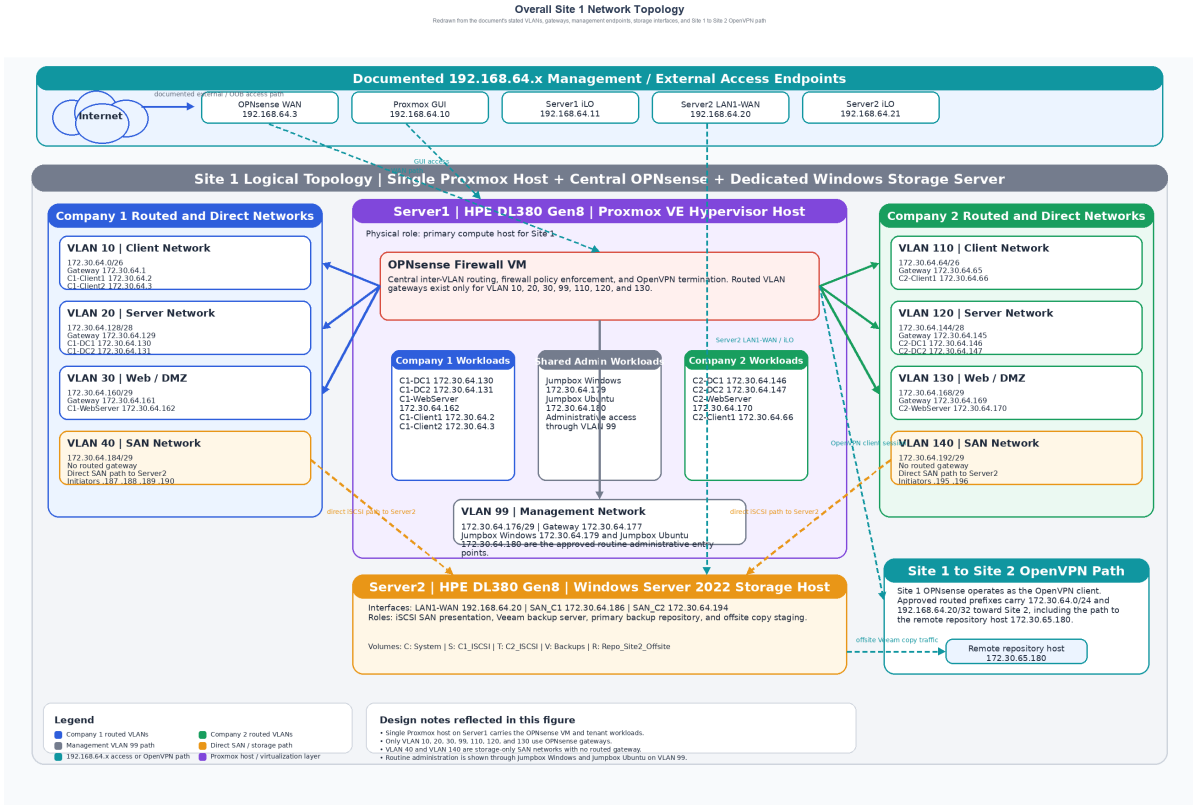


Figure 1. Overall Site 1 network topology

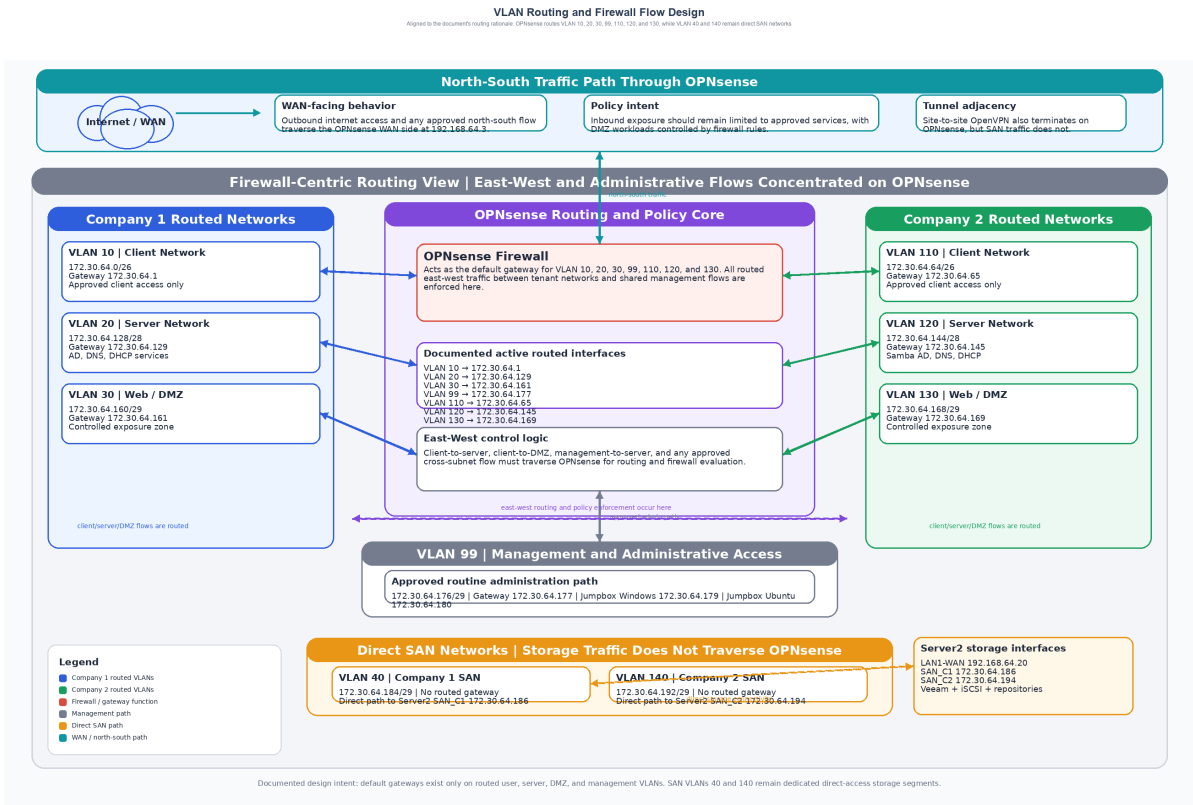


Figure 2. VLAN routing and firewall flow design

Network Design and IP Addressing Rationale

Routed east-west and north-south traffic traverses OPNsense. Default gateways are assigned only to the routed user, server, DMZ, and management VLANs, and management access is constrained to VLAN 99 jump hosts. The separate 192.168.64.0/24 segment is the shared infrastructure and WAN-side management network used for the OPNsense WAN interface, the Proxmox VE management interface, Server2 management connectivity, and the hardware iLO endpoints. SAN traffic remains on dedicated storage segments to avoid contention with client or server workloads, and the active Company 1 SAN data path is direct from the initiator SAN NICs to Server2 rather than through OPNsense.

Table 1. VLAN and gateway design

VLAN	Network	Gateway	Purpose	Traffic Type
VLAN 10	172.30.64.0/26	172.30.64.1	Company 1 Client Network	Routed
VLAN 20	172.30.64.128/28	172.30.64.129	Company 1 Server Network	Routed
VLAN 30	172.30.64.160/29	172.30.64.161	Company 1 Web / DMZ	Routed
VLAN 40	172.30.64.184/29	None (direct SAN)	Company 1 SAN Network	Direct storage
VLAN 99	172.30.64.176/29	172.30.64.177	Management Network	Restricted admin
VLAN 110	172.30.64.64/26	172.30.64.65	Company 2 Client Network	Routed
VLAN 120	172.30.64.144/28	172.30.64.145	Company 2 Server Network	Routed
VLAN 130	172.30.64.168/29	172.30.64.169	Company 2 Web / DMZ	Routed
VLAN 140	172.30.64.192/29	None (direct SAN)	Company 2 SAN Network	Direct storage
Infrastructure / WAN	192.168.64.0/24	192.168.64.1	Shared infrastructure and WAN-side management segment for OPNsense WAN, Proxmox VE, Server2 management, and iLO access	Shared infrastructure

Client networks are isolated by tenant and permitted to reach only approved services.

DMZ networks provide controlled exposure for web workloads.

SAN networks are dedicated to storage access and should not be used for general traffic.

The management VLAN is the only approved path for routine administrative access.

Identity Infrastructure and Tenant Separation

Table 2. Identity services summary

Organization	Directory Platform	Domain	Domain Controllers	Core Functions
Company 1	Windows Server 2022 AD DS	c1.local	C1-DC1 (172.30.64.130), C1-DC2 (172.30.64.131)	AD authentication, DNS, DHCP, Group Policy
Company 2	Samba AD on Ubuntu Server	c2.local	C2-DC1 (172.30.64.146), C2-DC2 (172.30.64.147)	AD-compatible auth, DNS, DHCP, SMB

Company 1 uses a traditional Microsoft Active Directory deployment with two Windows Server 2022 domain controllers. Company 2 uses Samba Active Directory on Ubuntu Server to provide a Linux-based domain environment with Microsoft-compatible protocols. No cross-domain trust relationship exists between c1.local and c2.local, which preserves administrative independence between the two tenants.

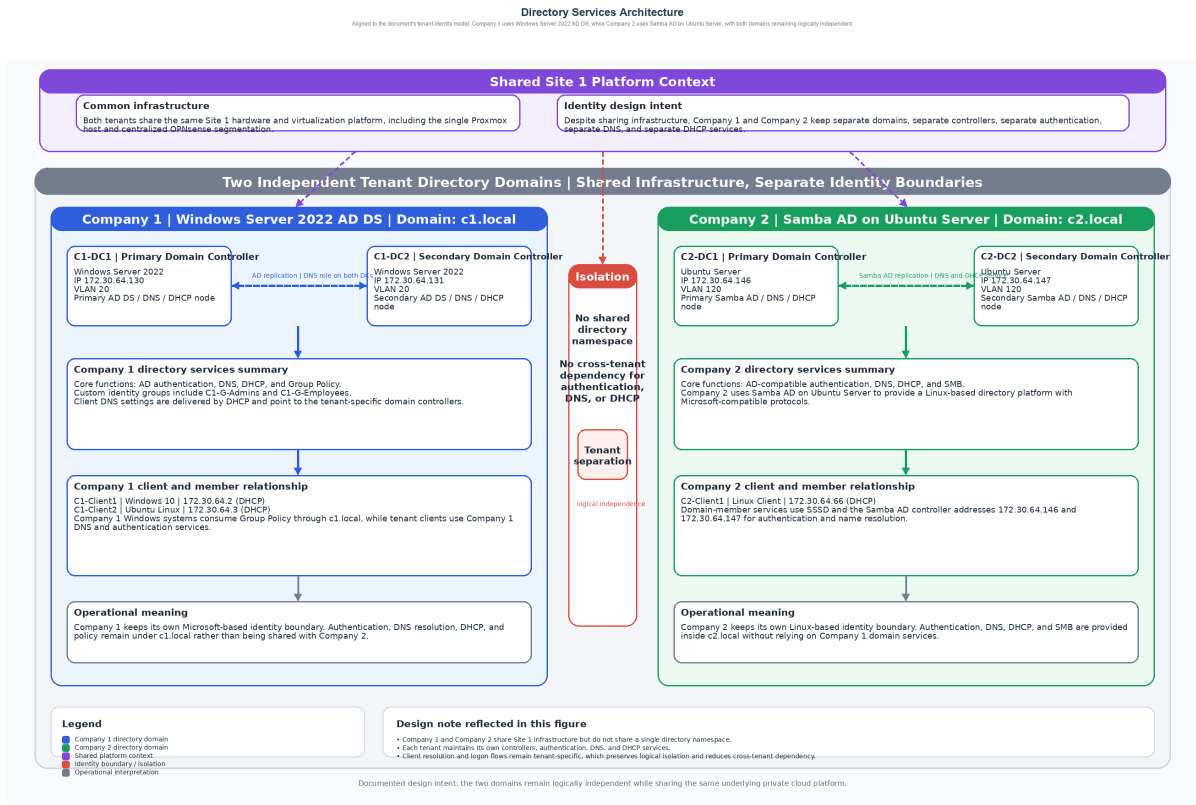


Figure 3. Directory services architecture

Platform and Service Choice Rationale

This section records the principal platform and service-composition decisions so that a receiving support team understands not only what was deployed, but also what operational trade-offs those choices introduce. The goal is to preserve rationale, not just inventory.

Table 3. Platform and service choice rationale

Platform / Service	Rationale
Proxmox VE	Selected as the hypervisor because it supports enterprise-style virtualization with low licensing overhead and strong lab flexibility [1].
OPNsense	Chosen to centralize VLAN routing, firewall policy, and gateway control with an interface suitable for support teams [2].
Windows Server AD DS	Used for Company 1 because it provides native Microsoft identity, DNS, DHCP, Group Policy, and integration with Windows clients [3].
Samba AD	Used for Company 2 to demonstrate a Linux-based directory platform with Microsoft-compatible protocols and lower platform cost [4].
Veeam Backup & Replication	Adopted to provide centralized backup operations, restore workflows, and recognizable enterprise backup practices [5].

Table 4. Service stacking rationale

Service Stack	Rationale
AD DS + DNS + DHCP + Group Policy on Company 1 domain controllers	These services are tightly related, reduce server count, and are common in smaller enterprise deployments where directory services control address assignment and policy.
Samba AD + DNS + DHCP + SMB on Company 2 domain controllers	Combining these Linux-based services reduces infrastructure overhead while preserving identity, name resolution, and file access for the second tenant.
iSCSI + backup repositories + Veeam on Server2	Storage and backup were grouped on the same physical system to conserve lab hardware and keep storage-adjacent services close to the data path.
Jumpbox + Windows Admin Center / Grafana / browser-based tooling	Management interfaces were grouped behind controlled administrative systems

	to reduce direct access from user networks and simplify support workflows.
--	--

In practice, the most demanding aspect of these stacking decisions was not the individual service builds, but keeping Windows AD DS and Samba AD cleanly separated on shared hardware without allowing DNS or DHCP scope to bleed across tenant boundaries. That constraint influenced many of the network and policy decisions described in the sections that follow.

Physical Server Baseline

Below the service and platform layer, Site 1 is physically hosted on two HP ProLiant DL380 Gen8 servers. This hardware baseline is important because the virtualization, storage, and backup design described later in the report depends directly on how these two hosts are divided between compute and shared storage roles.

Server1 is dedicated to Proxmox VE and acts as the primary compute host for the virtual environment. Server2 runs Windows Server 2022 directly on the hardware and provides centralized SAN storage for shared use. This split keeps compute and storage responsibilities separate, which makes the architecture easier to explain, operate, and troubleshoot.

The physical hardware allocation is summarized in the following table so that a receiving administrator can quickly see which disks are reserved for the operating systems, which arrays are used for workload storage, and how each server contributes to the overall platform under the fixed lab hardware delivered for this project.

Table 5. Physical server hardware summary

Server	Model	Base Role	Rear Drives / RAID / Use	Front Drives / RAID / Use	Design Intent
Server1	HP ProLiant DL380 Gen8	Proxmox VE hypervisor host	2 x 2 TB drives, RAID 1, used for the Proxmox operating system	2 x 2 TB drives, RAID 1, used for Proxmox VM storage	Keeps the hypervisor OS separate from guest workload storage while preserving basic disk redundancy.
Server2	HP ProLiant DL380 Gen8	Windows Server 2022 storage host	2 x 2 TB drives, RAID 1, used for the Windows Server 2022 operating system	6 x 2 TB drives, RAID 10, used for SAN shared storage	Uses the available eight-disk lab configuration to preserve separate OS and shared-storage

					roles while providing centralized SAN capacity and fault tolerance.
--	--	--	--	--	---

From an operational perspective, Table 12 confirms that Server1 should be described as the compute foundation of Site 1. The mirrored rear disks protect the Proxmox installation itself, while the mirrored front disks provide a separate and redundant storage area for virtual machine deployment. This makes it clear that the host was built first to run Proxmox reliably and second to provide local VM storage for the workloads it carries.

Server2 should be described as the storage foundation of Site 1. Installing Windows Server 2022 directly on mirrored rear disks isolates the operating system from the high-capacity front storage pool. In the lab-delivered eight-disk configuration, the front six-disk RAID 10 array then supports SAN storage for shared infrastructure use, preserving separate OS and workload storage roles within the available hardware envelope.



Figure 4. Site 1 rack installation (rear view)




Figure 5. Site 1 rack installation (front view)

Physical Switching and Uplink Baseline

A Cisco Catalyst lab switch provides the physical Layer 2 aggregation point for Site 1. It connects Server1, Server2, both iLO interfaces, the debug/admin port, and the upstream teacher-switch trunk, so it forms the handoff between the on-rack infrastructure and the wider lab network.

The Proxmox trunk on Gi8/0/1 carries VLANs 10, 20, 30, 40, 64, 65, 99, 110, 120, 130, and 140 with native VLAN 64 for shared infrastructure management. Gi8/0/5 and Gi8/0/6 are storage-only trunks that carry only VLANs 40 and 140 with native VLAN 999 as a blackhole native VLAN, while Gi8/0/24 uplinks only VLANs 1, 64, and 65 to the teacher switch. Server2 management and both iLO interfaces are placed on VLAN 64 access ports, and the switch itself is managed over SSH on 192.168.64.2 with HTTP disabled and rapid-PVST enabled.

Compute and Virtualization

Proxmox VE is the primary compute platform for all workloads in the environment. It hosts tenant servers, clients, jump systems, and shared services. Administrative access should be performed from the management VLAN through approved jump systems. The documented Site 1 deployment is a single-node Proxmox platform on Server1, and template virtual machines for Windows Server 2022, Windows 10, and Linux clients

were created to support repeatable provisioning and consistent configuration across the environment.

Table 6. Management endpoints

System	IP / URL	Access Method	Operational Use
Proxmox VE	https://192.168.64.10:8006	Web console	Hypervisor management
Grafana	http://172.30.64.180:3000	Web console	Monitoring dashboard
Cockpit	https://172.30.64.146:9090	Web console	Linux administration for Company 2
Windows Admin Center	https://172.30.64.179:6600	Web console	Windows infrastructure management
Jumpbox Windows	172.30.64.179	RDP / local admin tooling	Primary administrative entry point
Jumpbox Ubuntu	172.30.64.180	SSH / browser-based access	Linux administration plus Grafana and InfluxDB monitoring services
Proxmox Host iLO	https://192.168.64.11	Web console	Out-of-band hardware management for Server1
Server2 iLO	https://192.168.64.21	Web console	Out-of-band management endpoint for Server2
Lab Switch	ssh admin@192.168.64.2	SSH	Physical switching, VLAN trunk control, and upstream lab-network handoff

3. Discussion

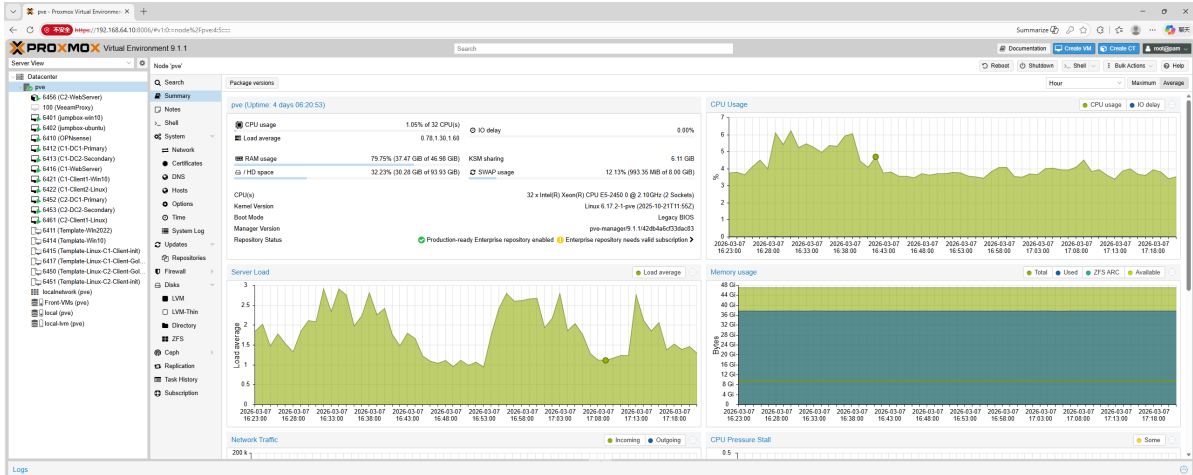


Figure 6. Proxmox VE management interface

Table 7. Company 1 virtual machine inventory

Virtual Machine	Operating System	IP Address	VLAN	Purpose
C1-DC1	Windows Server 2022	172.30.64.130	VLAN 20	Primary Domain Controller
C1-DC2	Windows Server 2022	172.30.64.131	VLAN 20	Secondary Domain Controller
C1-WebServer	Windows Server 2022	172.30.64.162	VLAN 30	Standalone IIS server in the Company 1 DMZ with hostname-restricted HTTPS access; only https://c1-webserver.c1.local is supported.
C1-Client1	Windows 10	172.30.64.2 (DHCP reservation); 172.30.64.189 (SAN)	VLAN 10 / VLAN 40	Client workstation with SAN iSCSI disk, Veeam Agent file-level backup to Server2, and managed RDP access through the jumpbox path
C1-Client2	Ubuntu Linux	172.30.64.3 (DHCP reservation); 172.30.64.190 (SAN)	VLAN 10 / VLAN 40	Client Workstation with SAN iSCSI disk

Table 8. Company 2 virtual machine inventory

Virtual Machine	Operating System	IP Address	VLAN	Purpose
C2-DC1	Ubuntu Server	172.30.64.146	VLAN 120	Primary Domain Controller
C2-DC2	Ubuntu Server	172.30.64.147	VLAN 120	Secondary Domain Controller
C2-WebServer	Linux Container	172.30.64.170	VLAN 130	Nginx-based Linux container web server in the Company 2 DMZ with hostname-restricted HTTPS access; only https://c2-webserver.c2.local is supported.
C2-Client1	Linux Client	172.30.64.66 (DHCP)	VLAN 110	Client workstation (realm member via SSSD; admin local, employee1@c2.local and employee2@c2.local validated as successful domain-user shell and home-directory sessions)

Table 9. Shared administrative systems

System	IP Address	VLAN	Purpose
Jumpbox Windows	172.30.64.179	VLAN 99	Administrative Access
Jumpbox Ubuntu	172.30.64.180	VLAN 99	Administrative Access
OPNsense Firewall	192.168.64.3	WAN Network	Network Gateway

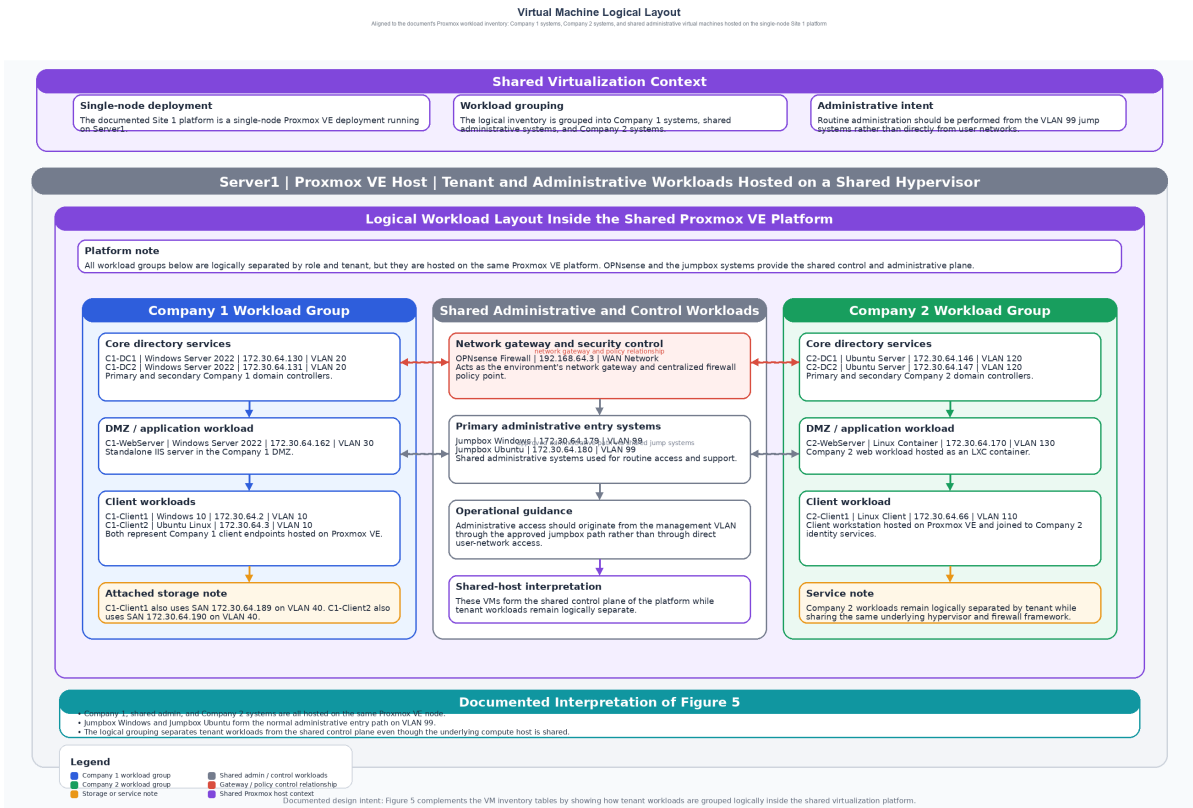


Figure 7. Virtual machine logical layout

Public-cloud boundary, MSP entry, and gateway model

Site 2 is organized around three clearly separated scopes: the MSP management scope, Company 1, and Company 2. The MSP scope contains the gateway, the bastions, and the backup platform. Company 1 contains a Windows-centric directory, file, web, and client stack. Company 2 contains a Linux-centric identity, file, web, and client stack. The result is a site that mixes platforms on purpose while keeping support entry points consistent.

Taken as service blocks, Site 2 delivers the full operating stack expected of this environment: a virtual router and policy edge, controlled remote access, tenant account administration, DNS and DHCP authority, internal HTTPS publication, file services, isolated iSCSI-backed storage, cross-site copy and synchronization paths, and Veeam-based protection. That matters for handover because every major service category has both an implementation and an owner. The site is therefore documented as a complete operating environment rather than as a partial lab sample.

Although the underlying hypervisor platform is outside the documentation scope, the virtual hardware baseline is not. Site 2 runs as 17 virtual machines allocated a combined 64 GiB RAM (65,536 MB). The MSP layer carries the gateway, dual bastions, and recovery platform, while each tenant carries its own identity, file, web, client, and storage roles. For support planning, that VM allocation is the practical baseline: capacity

questions at Site 2 begin with service placement and virtual sizing rather than with a physical chassis inventory.

OPNsense host rp-msp-gateway sits at the center of the topology. Its WAN side faces 172.20.64.1/16. Internally it anchors the MSP segment at 172.30.65.177/29, C1LAN at 172.30.65.1/26, C1DMZ at 172.30.65.161/29, C2LAN at 172.30.65.65/26, and C2DMZ at 172.30.65.169/29. SITE1_OVPN extends the site toward Site 1 for cross-site web and backup traffic. C1SAN and C2SAN are deliberately outside the routed OPNsense path, so block storage does not ride on the same path as tenant user traffic.

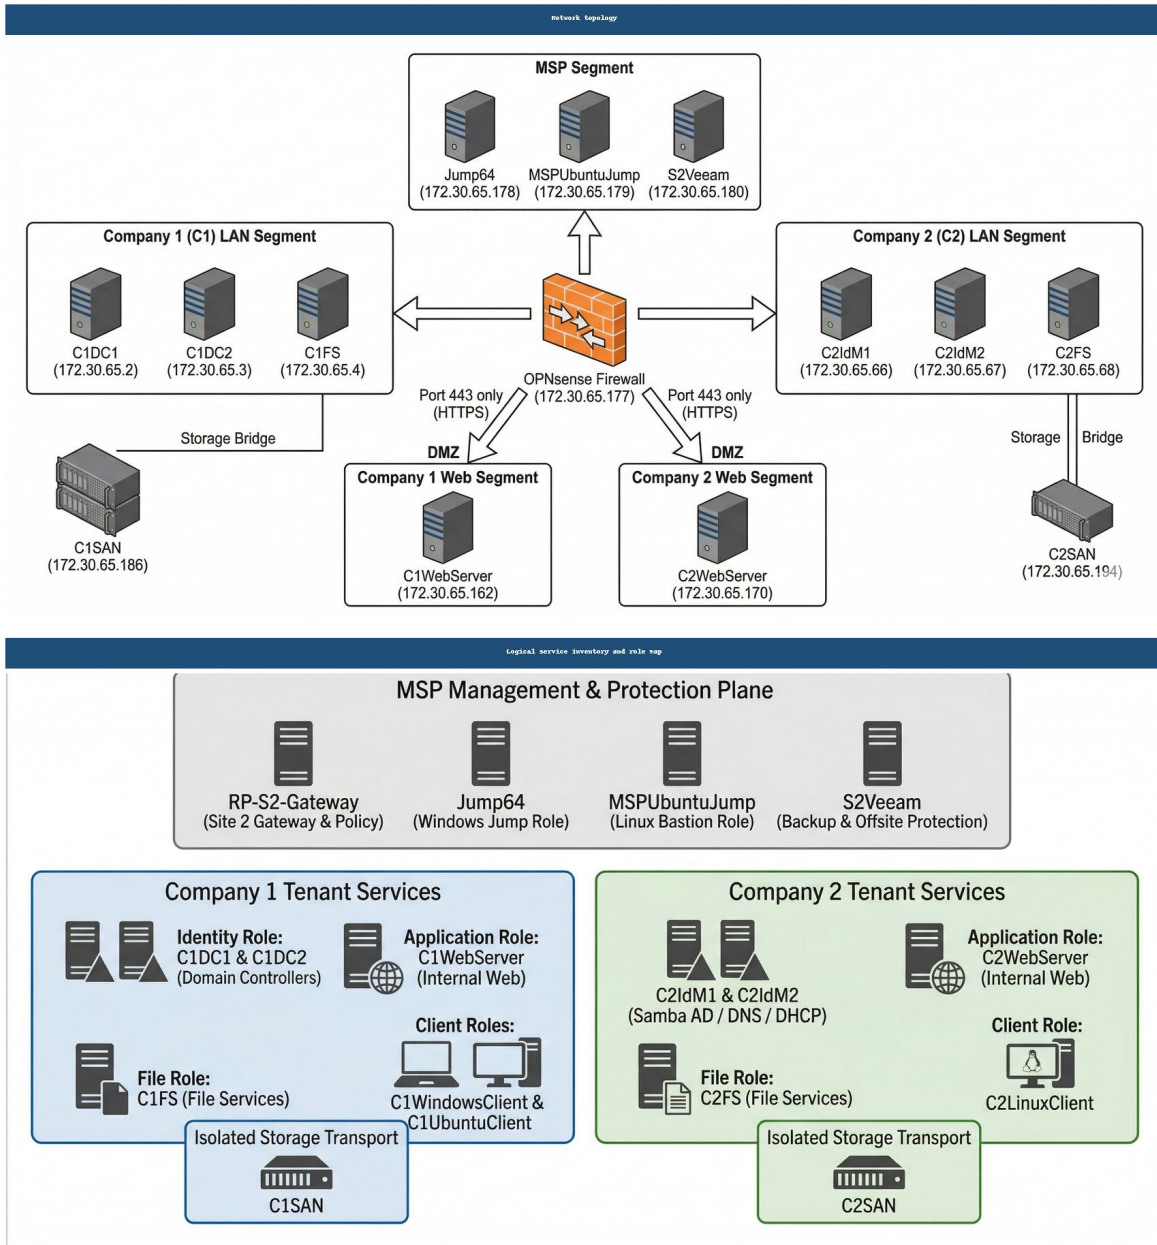


Figure 8. Site 2 network topology and service-role alignment

That topology matters because it makes fault domains visible. If the problem lives on the MSP segment, troubleshooting should begin with rp-msp-gateway, Jump64, MSPUbuntuJump, or S2Veeam. If the problem lives on C1DMZ, it is probably a Company 1 publication issue rather than a Company 2 identity issue. If the storage path breaks, the routed network can still be healthy while the file service fails. The diagram is therefore not just a map. It is a triage aid.

Network Design, IP Addressing, and Segmentation Rationale

Site 2 was not segmented only for security. It was segmented for operational clarity. Each network tells a support engineer where to look first. The MSP segment contains management and recovery tooling. C1LAN and C2LAN contain tenant identity, file, and endpoint workloads. C1DMZ and C2DMZ contain web publication tiers that should stay close to the edge but still behind OPNsense policy. The SAN bridges stay outside routed tenant paths so block transport and user traffic do not compete for the same trust boundary [20]-[22].

Table 10. Site 2 network segments and gateways

Segment	Addressing and Interface	Gateway or Routing Point	Purpose	Key Member Systems
WAN	172.20.64.1/16 on rp-msp-gateway	Upstream provider network	Provider-facing ingress and egress for the site edge	rp-msp-gateway
MSP	172.30.65.177/29 on rp-msp-gateway	Local gateway is rp-msp-gateway	Shared management and recovery segment for bastions and backup	Jump64 172.30.65.178, MSPUbuntuJump 172.30.65.179, S2Veeam 172.30.65.180
C1LAN	172.30.65.1/26 on rp-msp-gateway	Local gateway is rp-msp-gateway	Company 1 directory, file, and endpoint network	C1DC1 172.30.65.2, C1DC2 172.30.65.3, C1FS 172.30.65.4, C1WindowsClient 172.30.65.11, C1UbuntuClient 172.30.65.36
C1DMZ	172.30.65.161/29 on rp-msp-gateway	Local gateway is rp-msp-gateway	Company 1 web publication zone	C1WebServer 172.30.65.162
C2LAN	172.30.65.65/26 on rp-msp-gateway	Local gateway is rp-msp-gateway	Company 2 identity, file, and endpoint	C2IdM1 172.30.65.66, C2IdM2

			network	172.30.65.67, C2FS 172.30.65.68, C2LinuxClient 172.30.65.70
C2DMZ	172.30.65.169/29 on rp-msp-gateway	Local gateway is rp-msp-gateway	Company 2 web publication zone	C2WebServer 172.30.65.170
SITE1_OVPN	OpenVPN inter-site tunnel interface	Used for traffic to Site 1 and the static route 192.168.64.20/32	Controlled inter-site path for cross-site web access and Veeam copy traffic	Site 1 remote services and SITE1_VEEAM 192.168.64.20
C1SAN	172.30.65.186/29 with gateway 172.30.65.185	Not routed through OPNsense	Isolated Company 1 block storage bridge	C1SAN, consumed by C1FS
C2SAN	172.30.65.194/29 with gateway 172.30.65.193	Not routed through OPNsense	Isolated Company 2 block storage bridge	C2SAN, consumed by C2FS over ens18 172.30.65.195/29

Table 11. OPNsense firewall policy and alias summary

Type	Name or Rule	Value or Target	Operational Meaning
WAN NAT	33464 -> 172.30.65.178:3389	Jump64 RDP publication	Only the Windows bastion is exposed for remote RDP entry.
WAN NAT	33564 -> 172.30.65.179:22	MSPUbuntuJump SSH publication	Only the Linux bastion is exposed for remote shell entry.
Alias	C1_Nets	172.30.65.0/26 and 172.30.65.160/29	Company 1 addressing is grouped as one operational scope.
Alias	C2_Nets	172.30.65.64/26 and 172.30.65.168/29	Company 2 addressing is grouped as one operational scope.
Alias	C1_REMOTE and C2_REMOTE	Inter-site remote access aliases	Cross-site allowances are driven by named

			remote scopes instead of one-off IP rules.
Alias	ALL_WEBS	Combined web server alias	Shared web reachability is granted without opening unrelated tenant services.
Alias	ALL_DNS	Combined C1_DCs and C2_DCs alias	Tenant DNS authorities are reachable where name resolution is required.
Alias	S2_VEEAM	172.30.65.180	The Site 2 backup target has a fixed policy identity.
Alias	SITE1_VEEAM	192.168.64.20	The Site 1 copy endpoint has a fixed policy identity.
Alias	VEEAM_COPY_PORTS	135, 445, 6160, 6162, 2500-3000, 10005, 10006	Only the Veeam data and control ports needed for copy traffic are opened.
Policy	C1LAN allow and block set	Allow to C1_GLOBAL, ALL_WEBS, ALL_DNS; block to C2_GLOBAL	Company 1 can reach its own scope and shared dependencies, but not Company 2.
Policy	C2LAN allow and block set	Allow to C2_GLOBAL, ALL_WEBS, ALL_DNS; block to C1_GLOBAL	Company 2 can reach its own scope and shared dependencies, but not Company 1.
Cross-site web	C1_REMOTE -> 172.30.65.170/32 on HTTP/HTTPS	Company 1 remote users can reach the Company 2 web node only on the published web ports	Cross-site web consumption stays narrow and named.
Cross-site web	C2_REMOTE -> 172.30.65.162/32 on HTTP/HTTPS	Company 2 remote users can reach the Company 1 web node only on the published web ports	Cross-site web consumption stays narrow and named.
Backup copy	SITE1_VEEAM -> S2_VEEAM on VEEAM_COPY_PORTS	Veeam copy traffic from Site 1 to Site 2	Recovery traffic has an explicit firewall path.

Static route	192.168.64.20/32 via Site 1 OpenVPN gateway	Route to SITE1_VEEAM across SITE1_OVPN	The offsite copy target is reachable through a deterministic inter-site route.
--------------	---	--	--

The SAN isolation deserves special attention. C1SAN at 172.30.65.186/29 with gateway 172.30.65.185 and C2SAN at 172.30.65.194/29 with gateway 172.30.65.193 are not carried through OPNsense. That means storage failure analysis begins at the consumer - C1FS or C2FS - rather than at the routed firewall. That pattern keeps block transport private and makes the storage problem space smaller when a share disappears.

The WAN side is equally deliberate. Only the two jump hosts are published externally. Jump64 is reachable through 33464 -> 172.30.65.178:3389 and MSPUbuntuJump is reachable through 33564 -> 172.30.65.179:22. No tenant server is published directly at the edge. That decision narrows the attack surface, but just as importantly it prevents ad hoc administrative habits. If support work must begin on a bastion, the audit path stays predictable. For the support team, that segmentation identifies the correct starting point before time is lost searching the wrong tenant or the wrong layer.

MSP Entry, Remote Access, and Gateway Design

Site 2 is operated through a dual-bastion model. Jump64 is the primary Windows inspection platform. MSPUbuntuJump is the primary Linux inspection platform. Both were retained because this site genuinely needs both. Company 1 deep inspection and S2Veeam administration are easiest from Windows tooling. Company 2 identity, file, and web validation are fastest from Linux tooling. Splitting the bastions by operating model reduces the temptation to overextend one host with every tool in the environment.

The Linux bastion baseline was confirmed on March 25, 2026. Hostname mspubuntujump and user admin were visible on the host. Interface ens18 carried 172.30.65.179/29 with default gateway 172.30.65.177. A check to the OPNsense management plane returned HTTP 403 on port 80, which confirmed that the interface was present but not anonymously accessible. TCP 53 on rp-msp-gateway was also reachable from the jump host. That, combined with the alias inventory, NAT publication rules, and OpenVPN firewall policy confirmed through GUI screenshot evidence captured during the final pass, was sufficient to validate the full OPNsense management surface.

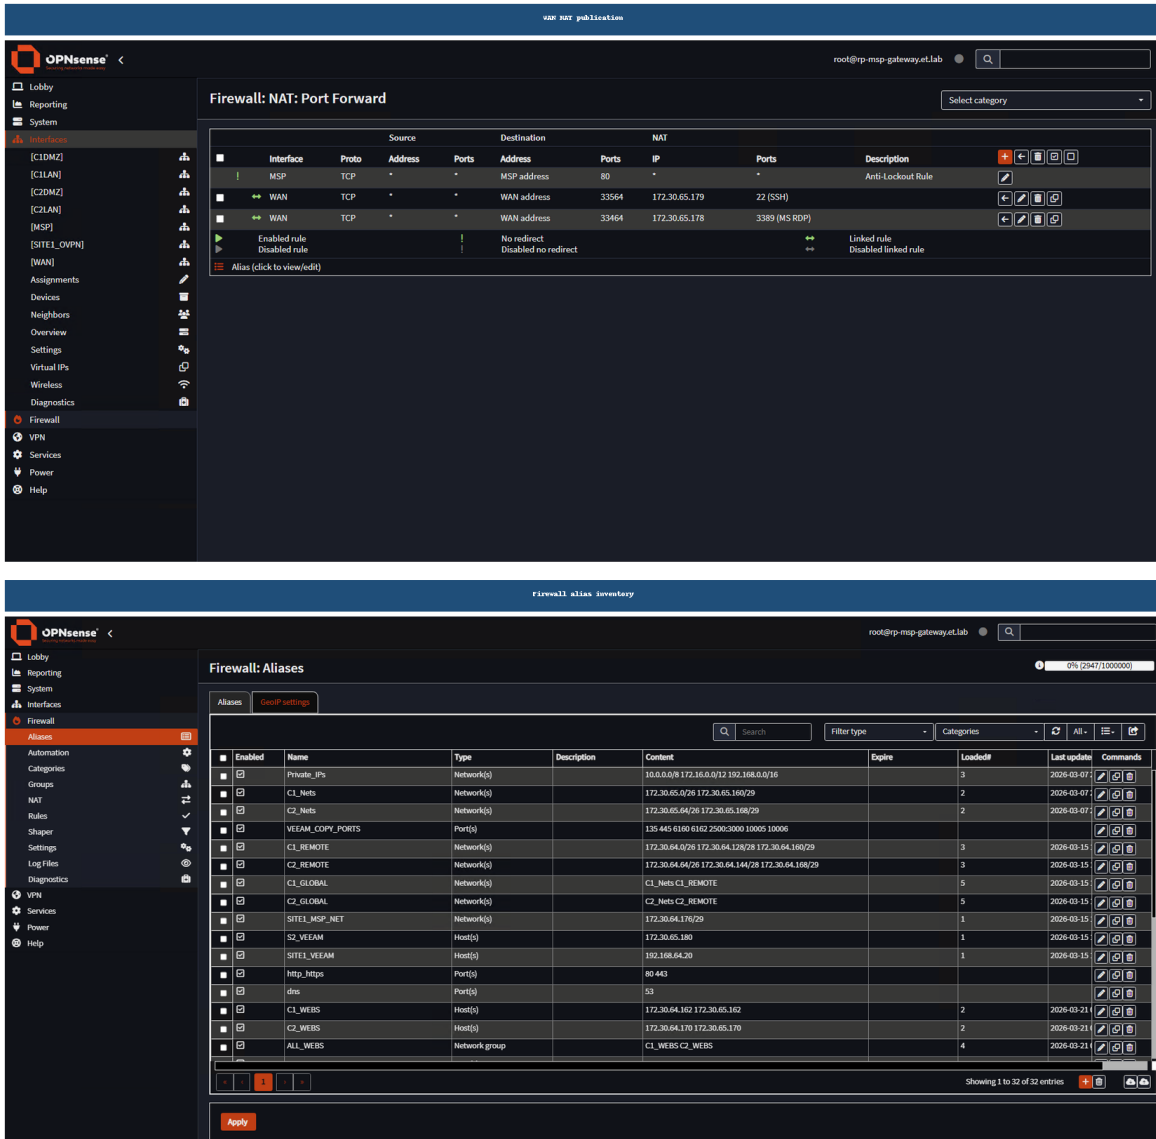


Figure 9. OPNsense edge publication, interface zones, and alias structure

The firewall model is alias-driven. C1_Nets, C2_Nets, ALL_WEBS, ALL_DNS, S2_VEEAM, SITE1_VEEAM, and VEEAM_COPY_PORTS give the policy a vocabulary that matches the site design rather than a pile of unrelated IP entries [20], [21]. That matters operationally. When a rule allows SITE1_VEEAM -> S2_VEEAM on VEEAM_COPY_PORTS, both the traffic purpose and the expected endpoint pair are immediately visible.

SITE1_OVPN gives the MSP a controlled relationship between the two sites. Cross-site web rules allow C1_REMOTE to reach 172.30.65.170/32 on HTTP/HTTPS and C2_REMOTE to reach 172.30.65.162/32 on HTTP/HTTPS. The same tunnel also carries the static route to 192.168.64.20/32 for offsite backup copy traffic. S2Veeam

therefore sits in exactly the right place: on the MSP segment where it remains accessible as a recovery anchor even if a tenant workload or tenant LAN is unhealthy.

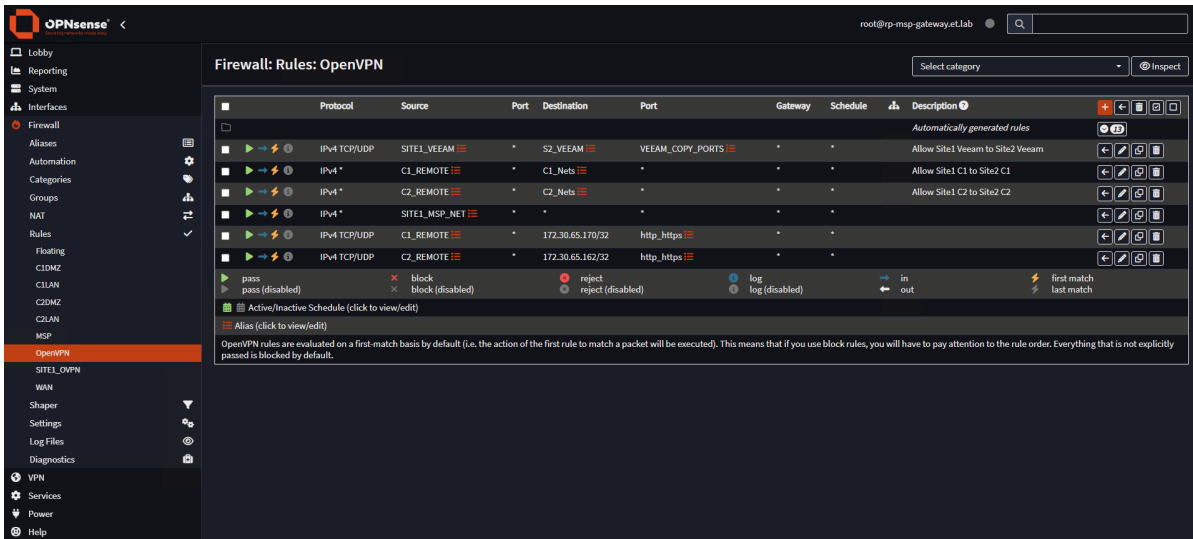


Figure 10. OPNsense OpenVPN inter-site tunnel and backup copy firewall rules

Table 12. MSP systems and platform summary

System	Address	Platform and Specs	How It Is Reached	Role	Operational Use
rp-msp-gateway	172.30.65.177	OPNsense 25.7, 4 vCPU, 2 GiB RAM, dual 16 GB disks, domain et.lab, timezone America/Toronto	Managed from MSPUbuntuJump across the MSP segment; management plane present and not anonymously accessible	Gateway, segmentation, policy control	Anchors interface zoning, alias-driven firewall policy, WAN publication, and the SITE1_OPENVPN path for the entire site.
Jump64	172.30.65.178	Windows Server 2022 Standard 10.0.20348 (Build 20348), 4 vCPU, 8 GiB RAM, C: 39.4 GB + D: 64 GB	RDP through WAN NAT 33464 -> 172.30.65.178:3389	Windows bastion	Primary Windows inspection platform for Company 1 and S2Veeam administration; workgroup-hosted.
MSPUbuntuJump	172.30.65.179	Ubuntu 22.04.5 LTS,	SSH through WAN NAT	Linux bastion	Primary Linux

		4 vCPU, 4 GiB RAM, 32 GB disk, 15 GB root LV, ens18 172.30.65.179/29	33564 -> 172.30.65.179:22		inspection platform for OPNsense reachability, Company 2 validation, and tenant port checks.
S2Veeam	172.30.65.180	Windows Server 2022 Standard 10.0.20348 (Build 20348), 4 vCPU, 8 GiB RAM, C: 63.3 GB + R: 499.9 GB + Z: 499.9 GB	Managed from Jump64 and reachable from MSPUbuntuJump on 445, 3389, 9392, 5985, 10005, and 10006	Backup and offsite-copy platform	Recovery anchor for the site, hosting the repository and copy workflows even when tenant workloads are impaired; workgroup-hosted.

That placement matters because it defines the order of operations during an outage. The support team can approach the site through the bastions, inspect rp-msp-gateway, and then decide whether the issue is tenant-specific or site-wide before touching backup or recovery. Operationally, that is the difference between a controlled recovery workflow and a scramble through unrelated systems.

3.2 Company 1 Services

Company 1 should be read as a complete tenant service stack rather than as a collection of disconnected hosts. The retained material below preserves the private-cloud-side identity and DHCP configuration references and the updated Site 2 evidence for observed service state, file presentation, web delivery, and client access.

Configuration reference retained from the private-cloud handover

Table 13. Company 1 domain controllers

Server	Role	IP Address	VLAN
C1-DC1	Primary Domain Controller	172.30.64.130	VLAN 20
C1-DC2	Secondary Domain Controller	172.30.64.131	VLAN 20

Table 14. Company 1 DHCP scope summary

Network	Address Range	Purpose
172.30.64.0/26	172.30.64.2 – 172.30.64.62	C1 Client Network

3. Discussion

```
C1-DC1 - 172.30.64.130 - 远程桌面连接
Administrator: Windows PowerShell
Windows PowerShell
Copyright (C) Microsoft Corporation. All rights reserved.

Install the latest PowerShell for new features and improvements! https://aka.ms/PSWindows

PS C:\Users\Administrator> Get-ADDomain

AllowedDNSSuffixes           : {}
ChildDomains                 : {}
ComputersContainer           : CN=Computers,DC=c1,DC=local
DeletedObjectsContainer      : CN=Deleted Objects,DC=c1,DC=local
DistinguishedName            : DC=c1,DC=local
DNSRoot                      : c1.local
DomainControllersContainer   : OU=Domain Controllers,DC=c1,DC=local
DomainMode                   : Windows2016Domain
DomainSID                    : S-1-5-21-235313996-2469821287-1438874053
ForeignSecurityPrincipalsContainer : CN=ForeignSecurityPrincipals,DC=c1,DC=local
Forest                      : c1.local
InfrastructureMaster         : C1-DC1.c1.local
LastLogonReplicationInterval :
LinkedGroupPolicyObjects     : {cn={BCA94FBF-5A83-4D39-99E0-C3EAAF7BAC9},cn=policies,cn=system,DC=c1,DC=local,
cn={B4A73C87-FA18-495F-B493-5C54A6F84A7F},cn=policies,cn=system,DC=c1,DC=local,
CN={31B2F340-016D-11D2-945F-00C04FB984F9},CN=Policies,CN=System,DC=c1,DC=local}
LostAndFoundContainer        : CN=LostAndFound,DC=c1,DC=local
ManagedBy                   :
Name                         : c1
NetBIOSName                  : C1
ObjectClass                   : domainDNS
ObjectGUID                   : b3370a87-a4ee-4e80-b9bc-22d4f948ce77
ParentDomain                  :
PDCEmulator                  : C1-DC1.c1.local
PublicKeyRequiredPasswordRolling : True
QuotasContainer              : CN=NTDS Quotas,DC=c1,DC=local
ReadOnlyReplicaDirectoryServers : {}
ReplicaDirectoryServers      : {C1-DC1.c1.local, C1-DC2.c1.local, C1DC1.c1.local, C1DC2.c1.local}
RIDMaster                     : C1-DC1.c1.local
SubordinateReferences        : {DC=ForestDnsZones,DC=c1,DC=local, DC=DomainDnsZones,DC=c1,DC=local, CN=Configuration,DC=c1,DC=local}
SystemsContainer              : CN=System,DC=c1,DC=local
UsersContainer                : CN=Users,DC=c1,DC=local
```

```
C1-DC1 - 172.30.64.130 - 远程桌面连接
Administrator: Windows PowerShell
PS C:\Users\Administrator> Get-DnsServerZone

ZoneName                ZoneType      IsAutoCreated  IsDnsIntegrated  IsReverseLookupZone  IsSigned
-----                -
_msdcs.c1.local         Primary       False          True              False                False
0.in-addr.arpa          Primary       True           False             True                 True
127.in-addr.arpa        Primary       True           False             True                 False
255.in-addr.arpa        Primary       True           False             True                 False
64.30.172.in-addr.arpa  Primary       False          True              True                 False
c1.local                Primary       False          True              False                False
c2.local                Primary       False          False             False                False
TrustAnchors            Primary       False          True              False                False

PS C:\Users\Administrator> Get-DnsServerForwarder

UseRootHint           : True
Timeout(s)            : 3
EnableReordering      : True
IPAddress              : 8.8.8.8
ReorderedIPAddress    : 8.8.8.8

PS C:\Users\Administrator> Get-DhcpServerv4Scope

ScopeId      SubnetMask      Name                State      StartRange      EndRange      LeaseDuration
-----
172.30.64.0  255.255.255.192 VLAN10_Clients Active    172.30.64.2    172.30.64.62  8.00:00:00

PS C:\Users\Administrator> Get-DhcpServerv4Reservation -ScopeId 172.30.64.0

IPAddress      ScopeId      ClientId            Name                Type              Description
-----
172.30.64.2    172.30.64.0  bc-24-11-84-a6-6c  client1.c1.local    Dhcp
172.30.64.3    172.30.64.0  bc-24-11-b6-e1-97  C1-Client2-Linux... Dhcp

PS C:\Users\Administrator>
```

Figure 11. C1-DC1 Windows PowerShell validation for Active Directory, DNS, and DHCP

Observed service state, client access, and platform evidence

Service Overview

Company 1 is the Windows-heavy tenant at Site 2. Its stack includes C1DC1 at 172.30.65.2 and C1DC2 at 172.30.65.3 as domain controllers, C1FS at 172.30.65.4 as the Windows file server and iSCSI consumer, C1WindowsClient at 172.30.65.11 as the Windows endpoint, C1UbuntuClient at 172.30.65.36 as the Linux endpoint, C1WebServer at 172.30.65.162 as the internal IIS server, and C1SAN at 172.30.65.186/29 as the isolated storage target. This is a complete service chain, not a partial tenant sample.

Table 15. Company 1 service and platform summary

System	Address	Platform and Specs	Function	Observed State	Operational Notes
C1DC1	172.30.65.2	Windows Server 2022 Standard 10.0.20348 (Build 20348), 4 vCPU, 4 GiB RAM, C: 39.3 GB	Primary domain controller	Inspection confirmed domain, forest, and directory service state through authenticated tenant sessions launched through Jump64; AD DS and DNS were installed.	One of two Company 1 controllers supplying identity and name services in c1.local.
C1DC2	172.30.65.3	Windows Server 2022 Standard 10.0.20348 (Build 20348), 4 vCPU, 4 GiB RAM, C: 39.3 GB	Secondary domain controller	Inspection confirmed domain, forest, and directory service state through authenticated tenant sessions launched through Jump64; AD DS and DNS were installed.	Reduces single-node identity risk for Company 1 while preserving local directory and DNS authority.
C1FS	172.30.65.4	Windows Server 2022 Standard 10.0.20348 (Build	File server and iSCSI consum	Windows SMB active; F: SharedData confirmed; named SMB	File publication is separated from the

		20348), 4 vCPU, 6 GiB RAM, C: 31.3 GB + F: 159.8 GB	er	shares present; active Get-IscsiSession confirmed.	directory tier and backed by isolated block storage.
C1WindowsClient	172.30.65.11	Windows 10 Education 10.0.19045 (Build 19045), 4 vCPU, 2 GiB RAM, C: 63.4 GB	Company 1 Windows endpoint	Domain membership confirmed through WMI and remote process execution; HTTP 200 to both web hostnames confirmed.	TCP 5985/WinRM not open; Jump64 used WMI instead.
C1UbuntuClient	172.30.65.36	Ubuntu 25.04, 4 vCPU, 2 GiB RAM, 32 GB disk, ens18 172.30.65.36/26	Company 1 Linux endpoint	Interactive shell validation on C1UbuntuClient confirmed C1.LOCAL realm state and HTTP 200 responses from both internal web hostnames.	Confirms that Company 1 service consumption does not depend on a Windows-only client path.
C1WebServer	172.30.65.162	Windows Server 2022 Standard 10.0.20348 (Build 20348), 4 vCPU, 2 GiB RAM, C: 31.3 GB	Internal web server	Workgroup-hosted, not domain-joined; binding restricted to c1-webserver.c1.local on TCP 443 only; raw IP returned HTTP 404.	The web tier is intentionally isolated from the Company 1 domain trust boundary.
C1SAN	172.30.65.186/29, gw 172.30.65.185	Windows Server 2022 storage target, 4 vCPU, 2 GiB RAM, 32 GB system disk + 160 GB data disk, SeaBIOS,	Company 1 iSCSI storage target	Direct bastion management access is intentionally blocked; health is inferred from the active consumer session on	The correct support path is through the consuming file server rather than direct

		VirtIO SCSI single, net0 on isolated bridge G264		C1FS.	administrative exposure.
--	--	---	--	-------	-----------------------------

Inspection from MSPUbuntuJump first confirmed that the Company 1 network was exposed exactly where the design expected it to be. C1DC1 and C1DC2 answered on TCP 53, 88, 389, 445, 3389, and 5985. C1FS answered on TCP 445 and 3389 while TCP 5985 timed out. C1WindowsClient answered on TCP 3389. C1UbuntuClient answered on TCP 22. C1WebServer answered on TCP 443, 3389, and 5985. That was a clean sign that Company 1 services were reachable through the MSP management plane without opening Company 1 directly to the WAN.

Supplemental hardware capture also filled the remaining Company 1 and MSP platform gaps. Guest-visible OS, vCPU, RAM, disk, and role data for the Company 1 Windows systems, the Company 1 Linux client, and the MSP platforms are now documented, so Tables 14 and 4 speak at roughly the same level of detail as the Company 2 tables without pretending the hypervisor layer is in scope.

```

admin@mSPubuntuJump:~$ echo "==== C1DC1 ===="; for p in 53 88 389 445 3389 5985; do nc -zvw2 172.30.65.2 $p; done
echo "==== C1DC2 ===="; for p in 53 88 389 445 3389 5985; do nc -zvw2 172.30.65.3 $p; done
echo "==== C1FS ===="; for p in 445 3389 5985; do nc -zvw2 172.30.65.4 $p; done
echo "==== C1WindowsClient ===="; nc -zvw2 172.30.65.11 3389
echo "==== C1UbuntuClient ===="; nc -zvw2 172.30.65.36 22
echo "==== C1WebServer ===="; for p in 443 3389 5985; do nc -zvw2 172.30.65.162 $p; done
==== C1DC1 ====
Connection to 172.30.65.2 53 port [tcp/domain] succeeded!
Connection to 172.30.65.2 88 port [tcp/kerberos] succeeded!
Connection to 172.30.65.2 389 port [tcp/ldap] succeeded!
Connection to 172.30.65.2 445 port [tcp/microsoft-ds] succeeded!
Connection to 172.30.65.2 3389 port [tcp/ms-wbt-server] succeeded!
Connection to 172.30.65.2 5985 port [tcp/*] succeeded!
==== C1DC2 ====
Connection to 172.30.65.3 53 port [tcp/domain] succeeded!
Connection to 172.30.65.3 88 port [tcp/kerberos] succeeded!
Connection to 172.30.65.3 389 port [tcp/ldap] succeeded!
Connection to 172.30.65.3 445 port [tcp/microsoft-ds] succeeded!
Connection to 172.30.65.3 3389 port [tcp/ms-wbt-server] succeeded!
Connection to 172.30.65.3 5985 port [tcp/*] succeeded!
==== C1FS ====
Connection to 172.30.65.4 445 port [tcp/microsoft-ds] succeeded!
Connection to 172.30.65.4 3389 port [tcp/ms-wbt-server] succeeded!
Connection to 172.30.65.4 5985 port [tcp/*] succeeded!
==== C1WindowsClient ====
Connection to 172.30.65.11 3389 port [tcp/ms-wbt-server] succeeded!
==== C1UbuntuClient ====
Connection to 172.30.65.36 22 port [tcp/ssh] succeeded!
==== C1WebServer ====
Connection to 172.30.65.162 443 port [tcp/https] succeeded!
Connection to 172.30.65.162 3389 port [tcp/ms-wbt-server] succeeded!
Connection to 172.30.65.162 5985 port [tcp/*] succeeded!
admin@mSPubuntuJump:~$

```

Figure 12. Company 1 services visible from the MSP management path (MSPUbuntuJump port checks)

Architectural Rationale

Two domain controllers were retained because Company 1 required a professionally supported identity tier that would not collapse onto a single Windows host. This improves resiliency, keeps DNS and authentication local to Company 1, and aligns with the client's expectation of controlled enterprise-style administration. A file share outage on C1FS

should not automatically imply a directory outage, and a directory issue on one controller should not erase the entire tenant namespace.

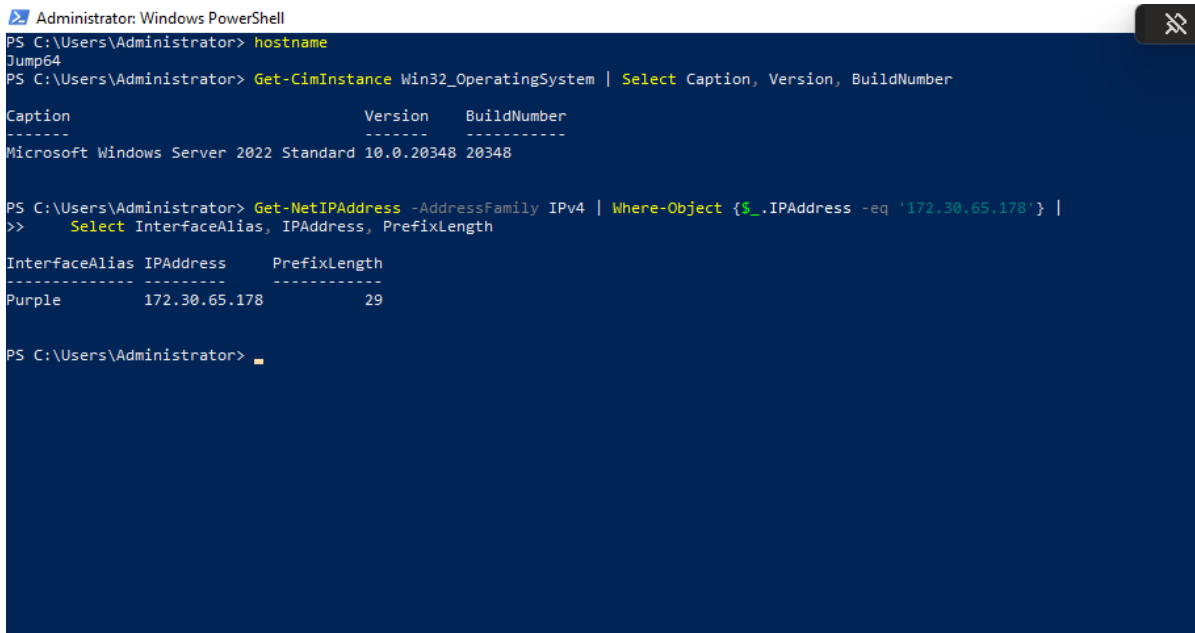
C1FS stands apart from the controllers because file I/O, share permissions, and block storage consumption are different operational problems than directory services [20]. Separating the role preserves predictable administration, allows the SMB tier to own the F: SharedData volume and iSCSI initiator session, and supports Company 1's preference for controlled data handling on a professionally managed platform.

C1WebServer is workgroup-hosted and not domain-joined on purpose. The Company 1 web tier is intended to consume the Company 1 namespace without inheriting the full Company 1 trust boundary. IIS is bound only to c1-webserver.c1.local on TCP 443, and raw IP access returns HTTP 404. That design keeps the site explicitly hostname-driven and reduces the blast radius of a compromise on the DMZ host [25].

C1SAN is isolated for the same reason that C2SAN is isolated: block storage should not live on the same trust path as user-facing traffic. Direct bastion management access to C1SAN is also intentionally blocked. That is not a missing feature. It is the design. The correct operational health signal is the active iSCSI consumer session on C1FS, not a habit of logging into storage targets from the management layer.

Observed Service State

Jump64 was the right place to perform the deep Company 1 inspection. Its own baseline as the Windows bastion at 172.30.65.178 was confirmed before it was used to reach the tenant systems. That mattered because Company 1 inspection needed PowerShell-native tooling and WMI access that the Linux bastion was not meant to replace.



```
Administrator: Windows PowerShell
PS C:\Users\Administrator> hostname
Jump64
PS C:\Users\Administrator> Get-CimInstance Win32_OperatingSystem | Select Caption, Version, BuildNumber

Caption                                Version      BuildNumber
-----
Microsoft Windows Server 2022 Standard 10.0.20348 20348

PS C:\Users\Administrator> Get-NetIPAddress -AddressFamily IPv4 | Where-Object {$_.IPAddress -eq '172.30.65.178'} |
>> Select InterfaceAlias, IPAddress, PrefixLength

InterfaceAlias  IPAddress      PrefixLength
-----
Purple          172.30.65.178 29

PS C:\Users\Administrator>
```

Figure 13. Jump64 Windows bastion baseline

Through authenticated tenant sessions launched from Jump64, inspection confirmed that C1DC1 and C1DC2 returned Company 1 domain, forest, controller inventory, and directory service state successfully. That is the strongest evidence needed for the core Company 1 identity layer: the tenant's own controllers answered from the approved MSP Windows path, and they did so as functioning domain controllers rather than as merely reachable hosts.

The image contains two screenshots of PowerShell terminal windows. The top screenshot is titled 'C1DC1 directory state' and shows the following commands and output:

```

Select Administrator: Windows PowerShell
PS C:\Users\Administrator.C1> hostname
C1DC1
PS C:\Users\Administrator.C1> Get-ADDomain | Select DNSRoot, NetBIOSName, DomainMode

DNSRoot  NetBIOSName      DomainMode
-----
c1.local C1                  Windows2016Domain

PS C:\Users\Administrator.C1> Get-ADForest | Select RootDomain, ForestMode

RootDomain      ForestMode
-----
c1.local        Windows2016Forest

PS C:\Users\Administrator.C1> Get-ADDomainController -Filter * | Select HostName, Domain, Forest, IsGlobalCatalog

HostName      Domain  Forest  IsGlobalCatalog
-----
C1-DC2.c1.local c1.local c1.local      True
C1-DC1.c1.local c1.local c1.local      True
C1DC1.c1.local  c1.local c1.local      True
C1DC2.c1.local  c1.local c1.local      True

PS C:\Users\Administrator.C1>

```

The bottom screenshot is titled 'C1DC2 directory state' and shows the following commands and output:

```

Administrator: Windows PowerShell
PS C:\Users\Administrator.C1> hostname
C1DC2
PS C:\Users\Administrator.C1> Get-ADDomain | Select DNSRoot, NetBIOSName, DomainMode

DNSRoot  NetBIOSName      DomainMode
-----
c1.local C1                  Windows2016Domain

PS C:\Users\Administrator.C1> Get-ADForest | Select RootDomain, ForestMode

RootDomain      ForestMode
-----
c1.local        Windows2016Forest

PS C:\Users\Administrator.C1> Get-ADDomainController -Filter * | Select HostName, Domain, Forest, IsGlobalCatalog

HostName      Domain  Forest  IsGlobalCatalog
-----
C1-DC2.c1.local c1.local c1.local      True
C1-DC1.c1.local c1.local c1.local      True
C1DC2.c1.local  c1.local c1.local      True
C1DC1.c1.local  c1.local c1.local      True

PS C:\Users\Administrator.C1>

```

Figure 14. C1DC1 and C1DC2 directory service state from authenticated tenant sessions

C1FS also validated cleanly through an authenticated Windows session reached from Jump64. Inspection confirmed that Windows SMB was active, the F: drive labeled SharedData was present, named SMB shares existed on F:, and an active iSCSI initiator session was visible through Get-IscsiSession. That chain matters because it proves more than one layer at once: Company 1 block storage reached the consumer, the consumer mounted it, and the file service published it.

```
Administrator: Windows PowerShell
PS C:\Users\Administrator.C1> hostname
C1FS
PS C:\Users\Administrator.C1> Get-Volume | Where-Object DriveLetter -eq 'F' |
>> Select DriveLetter, FileSystemLabel, Size, SizeRemaining

DriveLetter FileSystemLabel          Size SizeRemaining
-----
F SharedData      171621478400 170886627328

PS C:\Users\Administrator.C1> Get-SmbShare | Select Name, Path

Name      Path
----
ADMIN$    C:\Windows
C$        C:\
F$        F:\
IPC$
Priv_S2   F:\Priv_S2
PublicData F:\
Pub_S2    F:\Pub_S2

PS C:\Users\Administrator.C1> Get-IscsiSession | Select TargetNodeAddress, TargetNodeName

TargetNodeAddress          TargetNodeName
-----
iqn.1991-05.com.microsoft:win-5p43jh50pmm-c1fs-data-target

PS C:\Users\Administrator.C1> _
```

Figure 15. C1FS storage volume, SMB shares, and iSCSI session

C1WebServer behaved exactly like a hostname-first IIS tier should. It was workgroup-hosted rather than domain-joined, and its IIS binding was restricted to c1-webserver.c1.local on TCP 443 only. Raw IP requests returned HTTP 404. That state is considered healthy because the site is supposed to teach the client and the support team to trust names, certificates, and bindings rather than direct address access [25].

```
Administrator: Windows PowerShell
PS C:\Users\Administrator> hostname
C1WEBSERVER
PS C:\Users\Administrator> Get-CimInstance Win32_ComputerSystem | Select PartOfDomain, Domain, Workgroup

PartOfDomain Domain      Workgroup
-----
False WORKGROUP WORKGROUP

PS C:\Users\Administrator> Import-Module WebAdministration
PS C:\Users\Administrator> Get-WebBinding | Select protocol, bindingInformation

protocol bindingInformation
-----
https    *:443:c1-webserver.c1.local

PS C:\Users\Administrator> curl.exe -k -I https://c1-webserver.c1.local
curl: (6) Could not resolve host: c1-webserver.c1.local
PS C:\Users\Administrator> curl.exe -k -I https://172.30.65.162
HTTP/1.1 404 Not Found
Content-Length: 315
Content-Type: text/html; charset=us-ascii
Server: Microsoft-HTTPAPI/2.0
Date: Sun, 29 Mar 2026 02:44:55 GMT
Connection: close

PS C:\Users\Administrator> _
```

Figure 16. C1WebServer local workgroup status and IIS binding evidence

C1WindowsClient required a different inspection path because TCP 5985 was not open from the MSP side. WMI and remote process execution from Jump64 were therefore used during validation instead of pretending WinRM was available when it was not. For the final evidence package the same endpoint state was captured directly on the client console. Both paths confirmed c1.local domain membership, Company 1 DNS use, and HTTP 200 responses to both web hostnames [19].

```
Windows PowerShell
PS C:\Users\employee2> hostname
c1windowsclient
PS C:\Users\employee2> Get-CimInstance Win32_ComputerSystem | Select Domain, PartOfDomain

Domain    PartOfDomain
-----    -
c1.local      True

PS C:\Users\employee2> Get-DnsClientServerAddress -AddressFamily IPv4 | Select -ExpandProperty ServerAddresses
172.30.65.2
172.30.64.130
PS C:\Users\employee2> curl.exe -k -I https://c1-webserver.c1.local
HTTP/1.1 200 OK
Content-Length: 11022
Content-Type: text/html
Last-Modified: Mon, 09 Feb 2026 02:39:06 GMT
Accept-Ranges: bytes
ETag: "ac6edd426d99dc1:0"
Server: Microsoft-IIS/10.0
Date: Sun, 29 Mar 2026 02:52:26 GMT

PS C:\Users\employee2> curl.exe -k -I https://c2-webserver.c2.local
HTTP/1.1 200 OK
Server: nginx/1.18.0 (Ubuntu)
Date: Sat, 28 Mar 2026 23:52:34 GMT
Content-Type: text/html
Content-Length: 7311
Last-Modified: Fri, 20 Mar 2026 21:22:54 GMT
Connection: keep-alive
ETag: "69bdbaae-1c8f"
Accept-Ranges: bytes

PS C:\Users\employee2>
```

Figure 17. C1WindowsClient local domain membership, DNS, and dual-web access

C1UbuntuClient gave us the Linux-side counterpart. That validation was captured from an interactive client user session, so the final evidence panel below now shows the live C1UbuntuClient user context directly: the active C1.LOCAL realm and HTTP 200 responses to both web hostnames. That matters because it shows that Company 1 service consumption does not depend on a Windows-only client assumption. The hostname model holds from the Linux endpoint too.

```

s2-test@C1UbuntuClient:~$ hostname
echo
realm list
echo
curl -k -I https://c1-webserver.c1.local
echo
curl -k -I https://c2-webserver.c2.local
C1UbuntuClient

c1.local
type: kerberos
realm-name: C1.LOCAL
domain-name: c1.local
configured: kerberos-member
server-software: active-directory
client-software: sssd
required-package: sssd-tools
required-package: sssd
required-package: libnss-sss
required-package: libpam-sss
required-package: adcli
required-package: samba-common-bin
login-formats: %U
login-policy: allow-realm-logins

HTTP/2 200
content-length: 11022
content-type: text/html
last-modified: Mon, 09 Feb 2026 02:39:06 GMT
accept-ranges: bytes
etag: "ac6edd426d99dc1:0"
server: Microsoft-IIS/10.0
date: Sun, 29 Mar 2026 01:08:08 GMT

HTTP/1.1 200 OK
Server: nginx/1.18.0 (Ubuntu)
Date: Sun, 29 Mar 2026 01:08:08 GMT
Content-Type: text/html
Content-Length: 7311
Last-Modified: Fri, 20 Mar 2026 21:22:54 GMT
Connection: keep-alive
ETag: "69bdbaae-1c8f"
Accept-Ranges: bytes

s2-test@C1UbuntuClient:~$ █

```

Figure 18. C1UbuntuClient user-session realm state and dual-web access

Service Composition

Company 1's service chain is layered in a way that support staff can reason about quickly. C1DC1 and C1DC2 provide the identity and naming foundation. C1FS consumes isolated block storage from C1SAN and republishes it as SMB shares. C1WebServer publishes HTTPS only for c1-webserver.c1.local. C1WindowsClient and C1UbuntuClient consume those services through the Company 1 namespace. Nothing in that chain depends on exposing the tenant directly to the WAN.

The web behavior confirms the same design discipline. From MSPUbuntuJump on March 25, 2026, https://c1-webserver.c1.local returned HTTP/2 200 with Microsoft-IIS/10.0, while https://172.30.65.162 returned HTTP/2 404 with Microsoft-HTTPAPI/2.0. That exact split was the intended result. It tells the support team that a hostname failure and a web-service failure are not the same problem. If the raw IP works like the hostname, the binding is too loose. If the hostname works and the raw IP returns 404,

the binding is doing its job. In practice, that makes web triage faster and prevents the team from treating a DNS issue as an IIS outage.

3.3 Company 2 Services

Company 2 combines Samba AD, Linux-based client delivery, and a file-service design that differs materially from Company 1 while remaining part of the same managed environment. The retained blocks below preserve both the private-cloud-side reference points and the updated Site 2 evidence for the tenant's active service state.

Configuration reference retained from the private-cloud handover

Table 16. Company 2 domain controllers

Server	Role	IP Address	VLAN
C2-DC1	Primary Domain Controller	172.30.64.146	VLAN 120
C2-DC2	Secondary Domain Controller	172.30.64.147	VLAN 120

Table 17. Company 2 DHCP scope summary

Network	Address Range	Purpose
172.30.64.64/26	172.30.64.66 – 172.30.64.126	C2 Client Network

Table 18. Company 2 service stack summary

Function	Final Software / Service	Implementation Notes
DNS, directory, and authentication	Samba AD DC on C2-DC1 and C2-DC2	Provides the c2.local domain, AD-compatible authentication, and integrated DNS services. Validation confirmed the c2.local and _msdcs.c2.local zones plus successful internal and recursive lookups.
DHCP	ISC DHCP failover pair	C2-DC1 operates as primary and C2-DC2 as secondary for the 172.30.64.64/26 client scope, with MCLT 3600, split 128, and load balance max seconds 3.
SMB file sharing	Samba	Publishes C2_Public and C2_Private to internal users.
Replicated storage	GlusterFS volume gv0	Both domain controllers mount the replicated data

		set at /mnt/sync_disk so the same share data remains available on either node.
Client access workflow	admin was validated as the local Linux admin account, while employee1@c2.local and employee2@c2.local were validated to complete Company 2 domain-user shell/home-directory sessions and receive ~/C2_Public plus the correct per-user ~/C2_Private over SMB 3.1.1 from 172.30.64.146.	Confirms that the Linux client participates in the Company 2 identity and share workflow through successful domain-user session handling, home-directory landing, and per-user SMB mount presentation.
iSCSI initiators and SAN isolation	open-iscsi on C2-DC1 and C2-DC2	Sessions from 172.30.64.195 and 172.30.64.196 connect to Server2 SAN_C2 at 172.30.64.194, and the SAN interfaces remain separate from vlan120 tenant traffic.
Remote client access	OpenSSH on Company.c2.local	SSH access to 172.30.64.66 was validated from the management path, which satisfies the Linux client remote-access requirement.

```

ssh ▶ 172.30.64.146
[14:08:18] 98 admin@c2-dc1:~$ hostnamectl --static
[14:08:18] 99 systemctl is-active samba-ad-dc
[14:08:18] 100 systemctl is-active isc-dhcp-server
[14:08:18] 101 sudo grep -NE 'failover peer|primary;|subnet 172\.30\.64\.64' /etc/dhcp/dhcpd.conf
[14:08:18] 102 sudo samba-tool user list | egrep '^(Administrator|admin|employee1|employee2)$'
[14:08:18] 103 sudo samba-tool group list | grep -x 'c2_file_users'
[14:08:20] 104 c2-dc1.c2.local
[14:08:20] 105 active
[14:08:20] 106 active
[14:08:20] 107 6:failover peer "dhcp-failover" {
[14:08:20] 108 7: primary;
[14:08:20] 109 22:subnet 172.30.64.64 netmask 255.255.255.192 {
[14:08:20] 110 29: failover peer "dhcp-failover";
[14:08:21] 111 Administrator
[14:08:21] 112 employee1
[14:08:21] 113 admin
[14:08:21] 114 employee2
[14:08:21] 115 c2_file_users
[14:08:21] 116 admin@c2-dc1:~$

```

Figure 19. C2-DC1 primary-node DHCP failover, workflow accounts, and group evidence

```

ssh ▶ 172.30.64.146
136 admin@dc1:~$ host c2-webserver.c2.local 127.0.0.1
137 host microsoft.com 127.0.0.1
138 Using domain server:
139 Name: 127.0.0.1
140 Address: 127.0.0.1#53
141 Aliases:
142
143 c2-webserver.c2.local has address 172.30.64.170
144 c2-webserver.c2.local has address 172.30.65.170
145 Using domain server:
146 Name: 127.0.0.1
147 Address: 127.0.0.1#53
148 Aliases:
149
150 microsoft.com has address 13.107.253.51
151 microsoft.com has address 13.107.226.51
152 microsoft.com has IPv6 address 2620:1ec:29:1::51
153 microsoft.com has IPv6 address 2620:1ec:48:1::51
154 microsoft.com mail is handled by 10 microsoft-com.mail.protection.outlook.com.
155 admin@dc1:~$

```

Figure 20. C2-DC1 local DNS resolution and recursive lookup validation

```

32 admin@dc2:~$ hostnamed --static
33 systemctl is-active samba-ad-dc
34 systemctl is-active isc-dhcp-server
35 grep -nE 'failover peer|secondary;|subnet 172\.\.30\.\.64\.\.64' /etc/dhcp/dhcpd.conf
36 host c2-dc1.c2.local 127.0.0.1
37 host c2-dc2.c2.local 127.0.0.1
38 host c1-webserver.c1.local 127.0.0.1
39 host c2-webserver.c2.local 127.0.0.1
40 c2-dc2.c2.local
41 active
42 active
43 112:failover peer "dhcp-failover" {
44     secondary;
45     subnet 172.30.64.64 netmask 255.255.255.192 {
46         failover peer "dhcp-failover";
47 Using domain server:
48 Name: 127.0.0.1
49 Address: 127.0.0.1#53
50 Aliases:
51
52 c2-dc1.c2.local has address 172.30.64.146
53 Using domain server:
54 Name: 127.0.0.1
55 Address: 127.0.0.1#53
56 Aliases:
57
58 c2-dc2.c2.local has address 172.30.64.147
59 Using domain server:
60 Name: 127.0.0.1
61 Address: 127.0.0.1#53
62 Aliases:
63
64 c1-webserver.c1.local has address 172.30.64.162
65 c1-webserver.c1.local has address 172.30.65.162
66 Using domain server:
67 Name: 127.0.0.1
68 Address: 127.0.0.1#53
69 Aliases:
70
71 c2-webserver.c2.local has address 172.30.64.170
72 c2-webserver.c2.local has address 172.30.65.170
73 admin@dc2:~$

```

Figure 21. C2-DC2 secondary-node DHCP failover and DNS resolution validation

Identity, DNS, client delivery, and namespace evidence

Service Overview

Company 2 is the Linux-heavy tenant at Site 2. Its stack includes C2IdM1 at 172.30.65.66 and C2IdM2 at 172.30.65.67 as Samba AD DC, DNS, and DHCP nodes; C2F5 at 172.30.65.68 as the Samba file server and iSCSI consumer; C2LinuxClient at 172.30.65.70 as the endpoint; C2WebServer at 172.30.65.170 as the nginx HTTPS host; and C2SAN at 172.30.65.194/29 as the isolated storage target. Just like Company 1, this is a complete tenant stack rather than a single proof-of-concept node.

Table 19. Company 2 identity, DNS, and DHCP summary

Attribute	C2IdM1	C2IdM2
Address and platform	172.30.65.66, Ubuntu 22.04.5 LTS, 4 vCPU, 4 GiB RAM, 32 GB disk, 15 GB root LV, ens18	172.30.65.67, Ubuntu 22.04.5 LTS, 4 vCPU, 4 GiB RAM, 32 GB disk, 15 GB root LV, ens18
Directory state	samba-ad-dc active	samba-ad-dc active
DHCP state	isc-dhcp-server active, PRIMARY role	isc-dhcp-server active, SECONDARY role
Hosted zones	c2.local, c1.local, msdcs.c2.local	c2.local, c1.local, msdcs.c2.local
Key A records	c1-webserver.c1.local -> 172.30.64.162 and 172.30.65.162; c2-webserver.c2.local -> 172.30.64.170 and 172.30.65.170	Same records as C2IdM1; replication confirmed
Observed directory principals	Administrator, admin, employee1, employee2, krbtgt, Guest	Same directory principal inventory as C2IdM1
Observed groups	c2_file_users, DnsAdmins, Domain Controllers, Domain Computers, and standard AD groups	Same group inventory as C2IdM1

Table 20. Company 2 service summary

System	Address	Platform and Specs	Function	Observed State
C2IdM1	172.30.65.66	Ubuntu 22.04.5 LTS, 4 vCPU, 4 GiB RAM, 32 GB disk, 15 GB root LV, ens18	Primary Samba AD DC, DNS, DHCP	Directory, DNS, and DHCP services active; shared-zone visibility confirmed.
C2IdM2	172.30.65.67	Ubuntu 22.04.5 LTS, 4 vCPU, 4 GiB RAM, 32 GB disk, 15 GB root	Secondary Samba AD DC, DNS,	Replication and secondary DHCP failover role confirmed.

		LV, ens18	DHCP	
C2FS	172.30.65.68	Ubuntu 22.04.5 LTS, 4 vCPU, 6 GiB RAM, 16 GB system disk + 160 GB data disk at /mnt/c2_public, service NIC ens19 (172.30.65.68/26) , storage NIC ens18 (172.30.65.195/29)	Samba file server and iSCSI consumer	smbd active, /dev/sdb mounted, active iSCSI session to C2SAN, C2_Public and C2_Private shares defined.
C2LinuxClient	172.30.65.70	Ubuntu 25.04, 4 vCPU, 2 GiB RAM, 32 GB disk, ens18	Company 2 Linux endpoint	C2.LOCAL realm active; employee1@c2.local and employee2@c2.local resolved; SMB and HTTPS access validated.
C2WebServer	172.30.65.170	Ubuntu 22.04.5 LTS + nginx, 4 vCPU, 2 GiB RAM, 32 GB disk, 30 GB root LV, ens18	Internal HTTPS web server	nginx active; server_name c2-webserver.c2.local on port 443; raw IP returns HTTP 404.
C2SAN	172.30.65.194/29, gw 172.30.65.193	Ubuntu 22.04.5 LTS storage target, 4 vCPU, 2 GiB RAM, 32 GB system disk + 160 GB data disk, SeaBIOS, VirtIO SCSI single, net0 on isolated bridge M265	Company 2 iSCSI storage target	Consumed by C2FS across the dedicated storage interface rather than the tenant service interface.

The Company 2 identity nodes also expose the cross-domain DNS visibility between c1.local and c2.local. Inspection confirmed that both C2IdM1 and C2IdM2 hosted the zones c2.local, c1.local, and _msdcs.c2.local. The DNS A records for c1-webserver.c1.local returned both 172.30.64.162 and 172.30.65.162, while c2-webserver.c2.local returned both 172.30.64.170 and 172.30.65.170. That is operationally significant because Company 2 clients can resolve both Company 2 and Company 1 web identities from the same local identity tier.

Architectural Rationale

Samba AD on Ubuntu 22.04.5 LTS was selected for Company 2 because the client preference was to minimize licensing cost and vendor lock-in while still preserving Microsoft-compatible identity workflows [4]. Compared with Windows AD DS, Samba AD reduced platform cost and aligned more naturally with the tenant's Linux-heavy operating model, but it also required stronger Linux administration discipline around DNS, Kerberos, and client integration. That trade-off was acceptable here because the environment still needed domain-style authentication, policy-aware naming, and protocol compatibility for mixed clients rather than a reduced identity feature set.

Two identity nodes were also retained and DHCP was split into PRIMARY and SECONDARY roles because Company 2 should not lose authentication, naming, and address assignment all at once due to a single host failure. C2IdM1 and C2IdM2 mirror the same zones, the same A records, and the same directory principals. That gives support staff a clear redundancy story rather than a single Linux server doing too much.

C2F5 has a dual-NIC design because Company 2 file services need one path for users and another for block storage. The service NIC ens19 carries 172.30.65.68/26 on the tenant LAN. The storage NIC ens18 carries 172.30.65.195/29 toward C2SAN. That separation keeps iSCSI transport off the client-facing network and makes it much easier to tell whether a problem belongs to the LAN, the mount, or the block path [23], [24].

C2WebServer uses nginx because the Company 2 web tier does not need the rest of a large application platform. It needs clean HTTPS delivery, host-based publication, and easy inspection of the active site binding. nginx handles that cleanly with `server_name c2-webserver.c2.local` on port 443 and a default `server_name _ path` that returns HTTP 404 for raw IP access [18].

Observed Service State

C2IdM1 was healthy during direct inspection. `systemctl is-active` returned active for `samba-ad-dc` and active for `isc-dhcp-server`. Local `dig` queries against 127.0.0.1 returned both A records for `c1-webserver.c1.local` and `c2-webserver.c2.local`. `dhcpd.conf` showed the PRIMARY failover role, and `samba-tool dns zonelist` confirmed the hosted zones `c2.local`, `c1.local`, and `_msdcs.c2.local`. The directory user inventory included Administrator, admin, employee1, employee2, krbtgt, and Guest. Group output included `c2_file_users`, DnsAdmins, Domain Controllers, Domain Computers, and the standard AD groups.

```

admin@c2idm1:~$ hostname
systemctl is-active samba-ad-dc
systemctl is-active isc-dhcp-server
dig @127.0.0.1 c1-webserver.c1.local A +short
dig @127.0.0.1 c2-webserver.c2.local A +short
sudo grep -nE 'failover peer|primary;|secondary;' /etc/dhcp/dhcpd.conf
sudo samba-tool dns zonelist 127.0.0.1 -UAdministrator
c2idm1
active
active
172.30.64.162
172.30.65.162
172.30.64.170
172.30.65.170
10:failover peer "dhcp-failover" {
11:   primary;
27:   failover peer "dhcp-failover";
Password for [C2\Administrator]:
  3 zone(s) found

pszZoneName      : c2.local
Flags            : DNS_RPC_ZONE_DSINTEGRATED DNS_RPC_ZONE_UPDATE_SECURE
ZoneType        : DNS_ZONE_TYPE_PRIMARY
Version         : 50
dwDpFlags       : DNS_DP_AUTOCREATED DNS_DP_DOMAIN_DEFAULT DNS_DP_ENLISTED
pszDpFqdn       : DomainDnsZones.c2.local

pszZoneName      : c1.local
Flags            : DNS_RPC_ZONE_DSINTEGRATED DNS_RPC_ZONE_UPDATE_SECURE
ZoneType        : DNS_ZONE_TYPE_PRIMARY
Version         : 50
dwDpFlags       : DNS_DP_AUTOCREATED DNS_DP_DOMAIN_DEFAULT DNS_DP_ENLISTED
pszDpFqdn       : DomainDnsZones.c2.local

pszZoneName      : _msdcs.c2.local
Flags            : DNS_RPC_ZONE_DSINTEGRATED DNS_RPC_ZONE_UPDATE_SECURE
ZoneType        : DNS_ZONE_TYPE_PRIMARY
Version         : 50
dwDpFlags       : DNS_DP_AUTOCREATED DNS_DP_FOREST_DEFAULT DNS_DP_ENLISTED
pszDpFqdn       : ForestDnsZones.c2.local
admin@c2idm1:~$ █

```

Figure 22. C2IdM1 Active Directory, DNS, and DHCP evidence

C2IdM2 matched that reading. `systemctl is-active` again returned active for `samba-ad-dc` and active for `isc-dhcp-server`. Local `dig` queries returned the same A records as C2IdM1, `dhcpd.conf` showed the `SECONDARY` failover role, and `samba-tool dns zonelist` exposed the same hosted zones. That confirmed replication rather than a second server that merely existed on paper.

```

admin@c2idm2:~$ hostname
systemctl is-active samba-ad-dc
systemctl is-active isc-dhcp-server
dig @127.0.0.1 c1-webserver.c1.local A +short
dig @127.0.0.1 c2-webserver.c2.local A +short
sudo grep -nE 'failover peer|primary;|secondary;' /etc/dhcp/dhcpd.conf
sudo samba-tool dns zonelist 127.0.0.1 -UAdministrator
c2idm2
active
active
172.30.64.162
172.30.65.162
172.30.64.170
172.30.65.170
10:failover peer "dhcp-failover" {
11:     secondary;
25:     failover peer "dhcp-failover";
Password for [C2\Administrator]:
 3 zone(s) found

pszZoneName      : c2.local
Flags            : DNS_RPC_ZONE_DSINTEGRATED DNS_RPC_ZONE_UPDATE_SECURE
ZoneType        : DNS_ZONE_TYPE_PRIMARY
Version         : 50
dwDpFlags       : DNS_DP_AUTOCREATED DNS_DP_DOMAIN_DEFAULT DNS_DP_ENLISTED
pszDpFqdn       : DomainDnsZones.c2.local

pszZoneName      : c1.local
Flags            : DNS_RPC_ZONE_DSINTEGRATED DNS_RPC_ZONE_UPDATE_SECURE
ZoneType        : DNS_ZONE_TYPE_PRIMARY
Version         : 50
dwDpFlags       : DNS_DP_AUTOCREATED DNS_DP_DOMAIN_DEFAULT DNS_DP_ENLISTED
pszDpFqdn       : DomainDnsZones.c2.local

pszZoneName      : _msdcs.c2.local
Flags            : DNS_RPC_ZONE_DSINTEGRATED DNS_RPC_ZONE_UPDATE_SECURE
ZoneType        : DNS_ZONE_TYPE_PRIMARY
Version         : 50
dwDpFlags       : DNS_DP_AUTOCREATED DNS_DP_FOREST_DEFAULT DNS_DP_ENLISTED
pszDpFqdn       : ForestDnsZones.c2.local
admin@c2idm2:~$ █

```

Figure 23. C2IdM2 Active Directory, DNS, and DHCP evidence

The shared-zone view on both identity nodes is one of the most interesting outcomes in the whole site. Design evidence indicates that Company 2 was meant to carry a shared DNS visibility across both domain names, and inspection confirmed that intent live. Both c1.local and c2.local were present on the Company 2 identity nodes, and both internal web names resolved there. That gives the support team a local place to troubleshoot cross-domain naming rather than treating Company 1 names as something remote and mysterious.

```

admin@c2idm1:~$ hostname
systemctl is-active samba-ad-dc
systemctl is-active isc-dhcp-server
dig @127.0.0.1 c1-webserver.c1.local A +short
dig @127.0.0.1 c2-webserver.c2.local A +short
sudo grep -nE 'failover peer|primary;|secondary;' /etc/dhcp/dhcpd.conf
sudo samba-tool dns zonelist 127.0.0.1 -UAdministrator
c2idm1
active
active
172.30.64.162
172.30.65.162
172.30.64.170
172.30.65.170
10:failover peer "dhcp-failover" {
11:   primary;
27:   failover peer "dhcp-failover";
Password for [C2\Administrator]:
  3 zone(s) found

  pszZoneName      : c2.local
  Flags            : DNS_RPC_ZONE_DSINTEGRATED DNS_RPC_ZONE_UPDATE_SECURE
  ZoneType         : DNS_ZONE_TYPE_PRIMARY
  Version          : 50
  dwDpFlags        : DNS_DP_AUTOCREATED DNS_DP_DOMAIN_DEFAULT DNS_DP_ENLISTED
  pszDpFqdn        : DomainDnsZones.c2.local

  pszZoneName      : c1.local
  Flags            : DNS_RPC_ZONE_DSINTEGRATED DNS_RPC_ZONE_UPDATE_SECURE
  ZoneType         : DNS_ZONE_TYPE_PRIMARY
  Version          : 50
  dwDpFlags        : DNS_DP_AUTOCREATED DNS_DP_DOMAIN_DEFAULT DNS_DP_ENLISTED
  pszDpFqdn        : DomainDnsZones.c2.local

  pszZoneName      : _msdcs.c2.local
  Flags            : DNS_RPC_ZONE_DSINTEGRATED DNS_RPC_ZONE_UPDATE_SECURE
  ZoneType         : DNS_ZONE_TYPE_PRIMARY
  Version          : 50
  dwDpFlags        : DNS_DP_AUTOCREATED DNS_DP_FOREST_DEFAULT DNS_DP_ENLISTED
  pszDpFqdn        : ForestDnsZones.c2.local
admin@c2idm1:~$ █

```

Figure 24. Cross-domain DNS visibility without AD trust

C2FS validated cleanly as a layered service. systemctl showed smbd active. /dev/sdb was mounted at /mnt/c2_public as ext4 with 160 GB and 0% used. iscsiadm -m session showed tcp:[20] 172.30.65.194:3260,1 iqn.2024-03.org.clearroots:c2san (non-flash). testparm -s confirmed the [C2_Public] and [C2_Private] definitions exactly as designed, including valid users = @c2_file_users for the public share and valid users = %U for the private share.

```

admin@c2fs:~$ hostname
systemctl is-active smb
sudo iscsiadm -m session
findmnt /mnt/c2_public
testparm -s | sed -n '\[C2_Public\]/,/^$/p;\[C2_Private\]/,/^$/p'
c2fs
active
[sudo] password for admin:
tcp: [1] 172.30.65.194:3260,1 iqn.2024-03.org.clearroots:c2san (non-flash)
TARGET          SOURCE          FSTYPE OPTIONS
/mnt/c2_public /dev/sdb ext4    rw,relatime,stripe=8191
Load smb config files from /etc/samba/smb.conf
Loaded services file OK.
Weak crypto is allowed

Server role: ROLE_DOMAIN_MEMBER

[C2_Public]
create mask = 0770
directory mask = 0770
force group = c2_file_users
path = /mnt/c2_public/Public
read only = No
valid users = @c2_file_users

[C2_Private]
browseable = No
create mask = 0700
directory mask = 0700
force group = c2_file_users
path = /mnt/c2_public/Private/%U
read only = No
valid users = %U
admin@c2fs:~$ █

```

Figure 25. C2FS iSCSI-backed storage, mounted volume, and share definitions

The C2FS directory tree and sync state added one more useful operational layer. `/mnt/c2_public/Public` and `/mnt/c2_public/Private/{admin, admindc, employee1, employee2, user1, user2}` were present. The most recent sync log entry stated that the Site1->Site2 C2 sync completed successfully on 2026-03-25 02:00:31. That confirms that the file server is not only mounted and shared, but also participating in the site's data-movement design.

```

admin@c2fs:~$ sudo tail -n 20 /var/log/c2_site1_sync.log
sudo ls -lah /mnt/c2_public
sudo ls -lah /mnt/c2_public/Private
[sudo] password for admin:
[2026-03-27 02:00:03] Mirroring staged Public into /mnt/c2_public/Public
[2026-03-27 02:00:03] Mirroring staged Public into /mnt/c2_public/Public
[2026-03-27 02:00:03] Pulling Private from 172.30.64.146:/mnt/sync_disk/Private
[2026-03-27 02:00:03] Pulling Private from 172.30.64.146:/mnt/sync_disk/Private
[2026-03-27 02:00:25] Mirroring staged Private into /mnt/c2_public/Private
[2026-03-27 02:00:25] Mirroring staged Private into /mnt/c2_public/Private
[2026-03-27 02:00:25] Sync completed successfully
[2026-03-27 02:00:25] Sync completed successfully
[2026-03-28 02:00:02] Starting Site1 -> Site2 C2 sync
[2026-03-28 02:00:02] Starting Site1 -> Site2 C2 sync
[2026-03-28 02:00:02] Pulling Public from 172.30.64.146:/mnt/sync_disk/Public
[2026-03-28 02:00:02] Pulling Public from 172.30.64.146:/mnt/sync_disk/Public
[2026-03-28 02:00:03] Mirroring staged Public into /mnt/c2_public/Public
[2026-03-28 02:00:03] Mirroring staged Public into /mnt/c2_public/Public
[2026-03-28 02:00:03] Pulling Private from 172.30.64.146:/mnt/sync_disk/Private
[2026-03-28 02:00:03] Pulling Private from 172.30.64.146:/mnt/sync_disk/Private
[2026-03-28 02:00:22] Mirroring staged Private into /mnt/c2_public/Private
[2026-03-28 02:00:22] Mirroring staged Private into /mnt/c2_public/Private
[2026-03-28 02:00:22] Sync completed successfully
[2026-03-28 02:00:22] Sync completed successfully
total 40K
drwxrws--- 5 root      c2_file_users 4.0K Mar 28 06:06 .
drwxr-xr-x 3 root      root           4.0K Mar  5 15:43 ..
drwxrws--- 2 root      c2_file_users 16K Mar  5 15:43 lost+found
drwx--x--x 8 odengiz   c2_file_users 4.0K Mar 11 23:59 Private
drwxrws--- 2 odengiz   c2_file_users 4.0K Mar 28 22:20 Public
-rwxrws--- 1 root      c2_file_users 25 Mar 11 21:53 test-from-admin.txt
-rwxr--w--- 1 employee1 c2_file_users 21 Mar 11 21:57 test-from-employee1.txt
total 32K
drwx--x--x 8 odengiz   c2_file_users 4.0K Mar 11 23:59 .
drwxrws--- 5 root      c2_file_users 4.0K Mar 28 06:06 ..
drwx----- 2 admin     c2_file_users 4.0K Mar 25 23:32 admin
drwx--x--x 2 root      c2_file_users 4.0K Feb 28 20:38 admindc
drwx----- 2 employee1 c2_file_users 4.0K Mar 28 22:20 employee1
drwx----- 2 employee2 c2_file_users 4.0K Mar 28 22:20 employee2
-rwxrwxrwx 1 root      root           0 Mar  5 20:53 force_sync.txt
drwx--x--x 2 root      c2_file_users 4.0K Feb 28 20:38 user1
drwx--x--x 2 root      c2_file_users 4.0K Feb 28 20:38 user2
admin@c2fs:~$ █

```

Figure 26. C2FS synchronization log and directory structure

C2LinuxClient confirmed the client-side identity path. realm list showed C2.LOCAL with kerberos-member, server-software: active-directory, and client-software: sssd. resolvectl status showed DNS servers 172.30.65.66 and 172.30.65.67 and the DNS search domain c2.local during the March 25 pass. getent passwd returned employee1@c2.local with UID 952001106 and employee2@c2.local with UID 952001110, and both home directories - /home/employee1@c2.local and /home/employee2@c2.local - were present. SMB access to //c2fs.c2.local/C2_Public and //c2fs.c2.local/C2_Private was also validated from the client.

That March 25 pass also caught a real issue: c1-webserver.c1.local returned REFUSED because c1.local was not in the resolver search scope at the time. The issue was corrected in a later pass by adding c1.local and c2.local to the resolver search scope. The final validated state is that both hostnames resolve and return HTTP 200. The evidence package therefore includes both the client identity snapshot and a follow-up

dual-hostname check after the resolver correction. The failed observation was retained in the documentation because it is valuable troubleshooting history, not something to hide.

A second C2LinuxClient issue was identified and resolved on March 29, 2026 during the final validation pass. Domain user logins were succeeding but the C2_Public and C2_Private shares were not mounting automatically at session open. Diagnosis via journalctl confirmed that pam_mount was triggering the mount at login time but cifs.upcall was failing to obtain a Kerberos service ticket for c2fs.c2.local, returning error code -1765328230. Two root causes were identified: C2FS\$ was missing the required cifs/c2fs and cifs/c2fs.c2.local Service Principal Names, and C2LinuxClient could not reach a KDC for C2.LOCAL due to a DNS resolution gap. Resolution proceeded in three steps. On C2IdM1, the missing SPNs were added using samba-tool spn add cifs/c2fs C2FS\$ and samba-tool spn add cifs/c2fs.c2.local C2FS\$. On C2LinuxClient, the DNS resolver was corrected to use 172.30.65.66 and 172.30.65.67 as authoritative name servers, and /etc/krb5.conf was updated to pin the C2.LOCAL KDC to the same addresses. The pam_mount volume options were then simplified by removing the multiuser, uid=, and gid= parameters from the mount option string, leaving sec=krb5,vers=3.0,cuid=%(USERUID),noperm with appropriate file_mode and dir_mode values per share. After those changes, employee2@c2.local SSH login automatically mounted both //c2fs.c2.local/C2_Public and //c2fs.c2.local/C2_Private via CIFS/Kerberos, confirmed by mount | grep c2fs showing both shares active with sec=krb5. The same configuration applies to all domain users.

```

Identity, resolver, and account visibility

employee2@c2.local@c2linuxclient:~$ hostname
realm list
resolvectl status
getent passwd employee1@c2.local
getent passwd employee2@c2.local
curl -k -I https://c2-webserver.c2.local
c2linuxclient
c2.local
  type: kerberos
  realm-name: C2.LOCAL
  domain-name: c2.local
  configured: kerberos-member
  server-software: active-directory
  client-software: sssd
  required-package: sssd-tools
  required-package: sssd
  required-package: libnss-sss
  required-package: libpam-sss
  required-package: adcli
  required-package: samba-common-bin
  login-formats: %U@c2.local
  login-policy: allow-realm-logins
Global
  Protocols: -LLMNR -mDNS -DNSoverTLS DNSSEC=no/unsupported
  resolv.conf mode: stub

Link 2 (ens18)
  Current Scopes: DNS
  Protocols: +DefaultRoute -LLMNR -mDNS -DNSoverTLS DNSSEC=no/unsupported
Current DNS Server: 172.30.65.66
  DNS Servers: 172.30.65.66 172.30.65.67
  DNS Domain: c1.local c2.local
  Default Route: yes
employee1@c2.local:*:952001106:952000513:employee1:/home/employee1@c2.local:/bin/bash
employee2@c2.local:*:952001110:952000513:employee2:/home/employee2@c2.local:/bin/bash
HTTP/1.1 200 OK
Server: nginx/1.26.3 (Ubuntu)
Date: Sun, 29 Mar 2026 00:06:28 GMT
Content-Type: text/html
Content-Length: 7311
Last-Modified: Thu, 19 Mar 2026 14:04:52 GMT
Connection: keep-alive
ETag: "69bc0284-1c8f"
Accept-Ranges: bytes

employee2@c2.local@c2linuxclient:~$ █

```

```

Corrected dual-hostname validation after resolver update

employee2@c2.local@c2linuxclient:~$ whoami
hostname
realm list
resolvectl status
curl -k -I https://c1-webserver.c1.local
curl -k -I https://c2-webserver.c2.local
employee2@c2.local
c2linuxclient
c2.local
  type: kerberos
  realm-name: C2.LOCAL
  domain-name: c2.local
  configured: kerberos-member
  server-software: active-directory
  client-software: sssd
  required-package: sssd-tools
  required-package: sssd
  required-package: libnss-sss
  required-package: libpam-sss
  required-package: adcli
  required-package: samba-common-bin
  login-formats: %U@c2.local
  login-policy: allow-realm-logins
Global
  Protocols: -LLMNR -mDNS -DNSoverTLS DNSSEC=no/unsupported
  resolv.conf mode: stub

Link 2 (ens18)
  Current Scopes: DNS
  Protocols: +DefaultRoute -LLMNR -mDNS -DNSoverTLS DNSSEC=no/unsupported
Current DNS Server: 172.30.65.66
  DNS Servers: 172.30.65.66 172.30.65.67

```

Figure 27. C2LinuxClient domain identity, resolver state, and corrected service access

C2WebServer completed the Company 2 picture. nginx was active. The active configuration used server_name c2-webserver.c2.local on HTTPS port 443, and the default server_name _ caught raw IP access and returned HTTP 404. With the correct Host header, the server returned HTTP 200 OK and Content-Length: 7311. Without the Host header, raw IP access returned HTTP 404 Not Found. On the March 25 MSPUbuntuJump path the hostname returned HTTP/1.1 200 OK with nginx/1.18.0 while the raw IP returned HTTP/1.1 404 Not Found with nginx/1.18.0. On the C2LinuxClient path the response identified nginx/1.26.3. Both readings were preserved because they came from live validation paths and both still support the same hostname-first design conclusion.

```

admin@c2-webserver:~$ hostname
systemctl is-active nginx
grep -RIn 'server_name\\|listen 443' /etc/nginx/sites-enabled
curl -k -I -H 'Host: c2-webserver.c2.local' https://127.0.0.1
curl -k -I https://172.30.65.170
c2-webserver
active
/etc/nginx/sites-enabled/c2web:2: listen 443 ssl default_server;
/etc/nginx/sites-enabled/c2web:3: server_name _;
/etc/nginx/sites-enabled/c2web:12: listen 443 ssl;
/etc/nginx/sites-enabled/c2web:13: server_name c2-webserver.c2.local;
HTTP/1.1 200 OK
Server: nginx/1.18.0 (Ubuntu)
Date: Sun, 29 Mar 2026 00:07:08 GMT
Content-Type: text/html
Content-Length: 7311
Last-Modified: Fri, 20 Mar 2026 21:22:54 GMT
Connection: keep-alive
ETag: "69bdbaae-1c8f"
Accept-Ranges: bytes

HTTP/1.1 404 Not Found
Server: nginx/1.18.0 (Ubuntu)
Date: Sun, 29 Mar 2026 00:07:08 GMT
Content-Type: text/html
Content-Length: 162
Connection: keep-alive

admin@c2-webserver:~$ █

```

Figure 28. C2WebServer nginx hostname binding and raw-IP behavior

Service Composition

Company 2's service chain begins with identity. C2IdM1 and C2IdM2 provide authentication, DNS, and DHCP. C2FS consumes isolated storage from C2SAN and

republishes it through Samba as C2_Public and C2_Private. C2WebServer publishes HTTPS only for c2-webserver.c2.local. C2LinuxClient consumes identity, web, and file services through the Company 2 namespace. Like Company 1, Company 2 is internally layered rather than flat.

The extra twist is the shared DNS visibility across both domain names. Because C2IdM1 and C2IdM2 host c1.local as well as c2.local, Company 2 users can resolve Company 1 web identities through the local identity tier once the resolver search scope includes both domains. That is why the March 25 resolver correction mattered. It was not just a client tweak. It made the intended multi-domain service composition visible to the endpoint. For the client taking over the site, that makes Company 2 name resolution a first-class dependency for cross-domain browsing rather than an optional convenience.

3.4 Shared Storage, SAN, and File Services

Storage and file presentation are intentionally grouped together here because both source documents describe them as multi-layer service chains: block transport first, mounted storage second, share publication third, and user consumption last.

Private-cloud storage and SAN baseline

Server2 is the storage backbone for the environment. It provides tenant-separated iSCSI volumes for Company 1 and Company 2 and dedicated Veeam repository volumes. SAN traffic stays on dedicated VLANs and interfaces rather than the routed tenant networks. In the Company 1 implementation, C1-DC1, C1-DC2, C1-Client1, and C1-Client2 use VLAN 40 to reach Server2 SAN_C1 at 172.30.64.186 from initiator addresses 172.30.64.187 through 172.30.64.190. Company 1 file services are then published through DFS namespace and DFS replication.

The final Company 1 file-service model uses Access-Based Enumeration on the underlying target shares and per-user private folders for admin, employee1, and employee2. Public collaboration data remains visible through the shared namespace path. Private content is restricted to the matching user folder. This preserves the intended collaboration-versus-private-data split while keeping the Windows-backed DFS design clean and supportable.

Table 21. Company 1 file service access control summary

Item	Final Configuration	Project Relevance
Namespace presentation	\\c1.local\namespace with public and private logical folders	Preserves the documented DFS namespace access model for Company 1 clients.
Back-end share targets	Pub1 and Priv1 on C1-DC1, Pub2 and Priv2 on C1-DC2, with C1FS as the secondary-site target	Provides multiple file-service nodes behind the namespace.
Share-layer permissions	Authenticated Users = Change, BUILTIN\Administrators =	Allows internal users to work in the shares while preserving administrative

	Full	control.
Private NTFS model	Per-user folders for admin, employee1, and employee2 with owner-only modify access plus SYSTEM/Administrators full control	Meets the requirement that private data is visible only to the owning user.
Fault tolerance	DFSR replicated folders for public and private; C1-DC2 766 GB SAN disk read-only state corrected before replication validation	Restores replicated file-service behavior across both Company 1 domain controllers.
Windows client validation	C1-Client1 validation focused on employee1 and employee2 workflows; Public was writable by both users and each user could see only the matching private folder	Provides direct client evidence for the intended employee-facing access-control outcome.
Linux client validation	C1-Client2-Linux now keeps admin as a local Linux admin account while employee1@c1.local and employee2@c1.local sign in as Company 1 domain users and automatically receive ~/C1_Public plus ~/C1_Private through per-user CIFS sessions; Public is presented from the user-facing shared-content directory and Private maps directly to the matching user path	Confirms that the Linux client also satisfies the mounted-share requirement and mirrors the Windows access-control outcome.

Company 1 Linux client access was normalized so that admin remains a local Linux account, while employee1@c1.local and employee2@c1.local sign in through the c1.local domain path and automatically receive ~/C1_Public and ~/C1_Private through per-user CIFS sessions. The public mount is bound to the shared-content directory, and the private mount points directly to the matching per-user private path. This gives the Ubuntu client the same effective access model as the Windows client without exposing DFSR system folders.

Company 2 uses a simpler Linux file-service model. C2-DC1 (172.30.64.146) and C2-DC2 (172.30.64.147) publish identical Samba shares from the GlusterFS replicated volume gv0 mounted at /mnt/sync_disk. C2_Public is the shared collaboration location, while C2_Private maps each user to /mnt/sync_disk/Private/%U so every account receives its own private path.

Table 22. Company 2 replicated file service summary

Item	Final Configuration	Project Relevance
File-service nodes	C2-DC1 172.30.64.146 and C2-DC2 172.30.64.147	Provides two Company 2 domain controllers that can both publish the same file shares
Shared data layer	GlusterFS replicated volume gv0 mounted at /mnt/sync_disk on both servers	Implements the replicated storage layer used for Company 2 file-service fault tolerance
Public share	C2_Public -> /mnt/sync_disk/Public	Allows Company 2 internal users to read and modify the shared collaboration folder
Private share	C2_Private -> /mnt/sync_disk/Private/%U	Redirects each authenticated user to a private per-user folder
Authorized users	Members of c2_file_users; share-access validation centered on employee1 and employee2 workflows while admin remains the local Linux administration account	Demonstrates controlled share access rather than anonymous or guest access
Client validation	Company.c2.local validated with employee1@c2.local and employee2@c2.local user workflows; successful domain-user sessions landed in the expected Linux home directories, live SMB 3.1.1 sessions mounted ~/C2_Public and ~/C2_Private from 172.30.64.146, and both workflows could write to Public while remaining limited to the matching Private content	Confirms that the Linux client can reach the shares, that the mounts persist in the final workflow, and that the private-folder model is enforced per user.

For Company 2, fault tolerance is implemented at the shared-storage and service layer. GlusterFS volume gv0 operates as a two-brick replicate volume across C2-DC1 and C2-DC2. Linux client validation confirmed that admin remains a local Linux account, while employee1@c2.local and employee2@c2.local complete successful domain-user sessions, land in their own home directories, mount ~/C2_Public and ~/C2_Private, and remain limited to their matching private content.

Company 2 storage validation also confirms tenant-separated iSCSI presentation from Server2. The GlusterFS brick disks are delivered over the dedicated SAN interfaces, while user and server traffic continues over the routed tenant interfaces. This preserves the required separation between storage traffic and general user traffic.

Table 23. Server2 volume layout

Drive Letter	Volume Name	File System	Purpose
C:	System Volume	NTFS	Windows Server operating system
S:	C1_ISCSI	NTFS	iSCSI storage for Company 1
T:	C2_ISCSI	NTFS	iSCSI storage for Company 2
V:	Backups	ReFS	Primary Veeam backup repository
Site2 R:	Site1OffsiteFromServer2\Repo	NTFS via SMB shared folder	Remote offsite shared-folder repository on 172.30.65.180 used by the Site1-to-Site2 Backup Copy job

Table 24. Company 1 iSCSI targets

Target Name	Initiator IP	Storage Location	Purpose
c1-dc1	172.30.64.187	S:\iSCSIVirtualDisks	DC1 Data Storage
c1-dc2	172.30.64.188	S:\iSCSIVirtualDisks	DC2 Data Storage
target-win10	172.30.64.189	S:\iSCSIVirtualDisks	Client1 SAN Disk
target-ubuntu	172.30.64.190	S:\iSCSIVirtualDisks	Client2 SAN Disk

Table 25. Company 2 iSCSI targets

Target Name	Initiator IP	Storage Location	Purpose
c2-dc1	172.30.64.195	T:\iSCSIVirtualDisks	DC1 Data Storage
c2-dc2	172.30.64.196	T:\iSCSIVirtualDisks	DC2 Data Storage

Table 26. SAN VLAN design

VLAN	Network	Purpose
VLAN 40	172.30.64.184/29	Company 1 SAN Network
VLAN 140	172.30.64.192/29	Company 2 SAN Network

Table 27. Storage server SAN interfaces

Interface	IP Address	Purpose
-----------	------------	---------

SAN_C1	172.30.64.186	Company 1 iSCSI Network
SAN_C2	172.30.64.194	Company 2 iSCSI Network

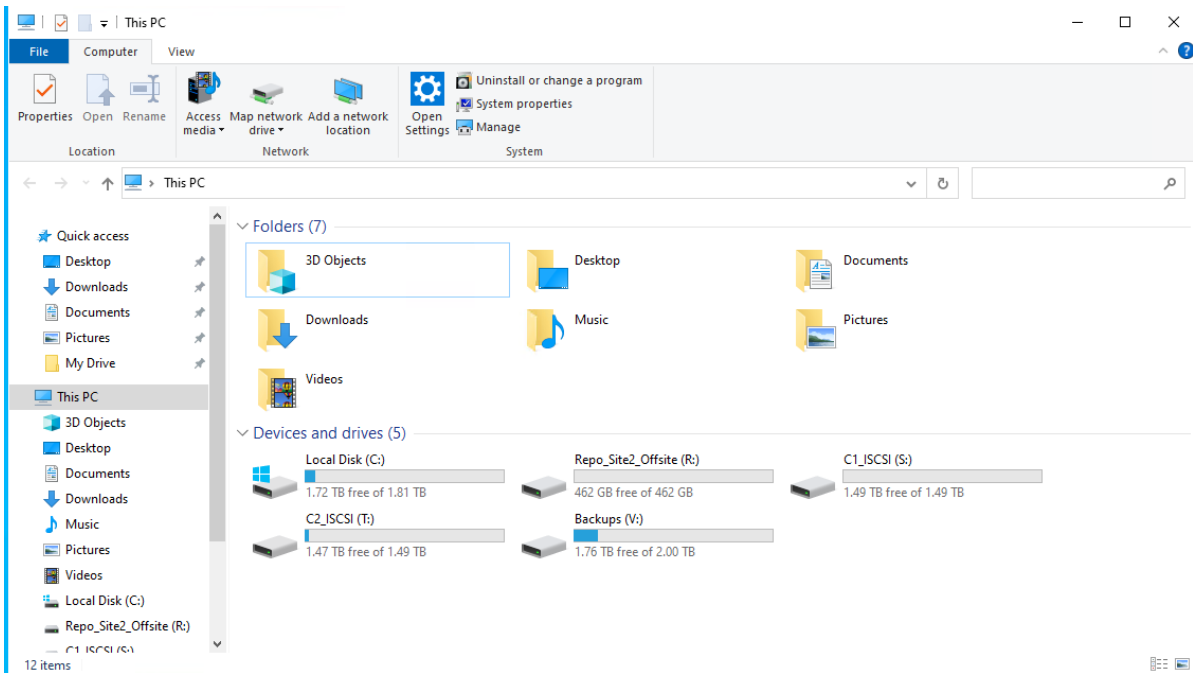


Figure 29. Server2 storage volume presentation

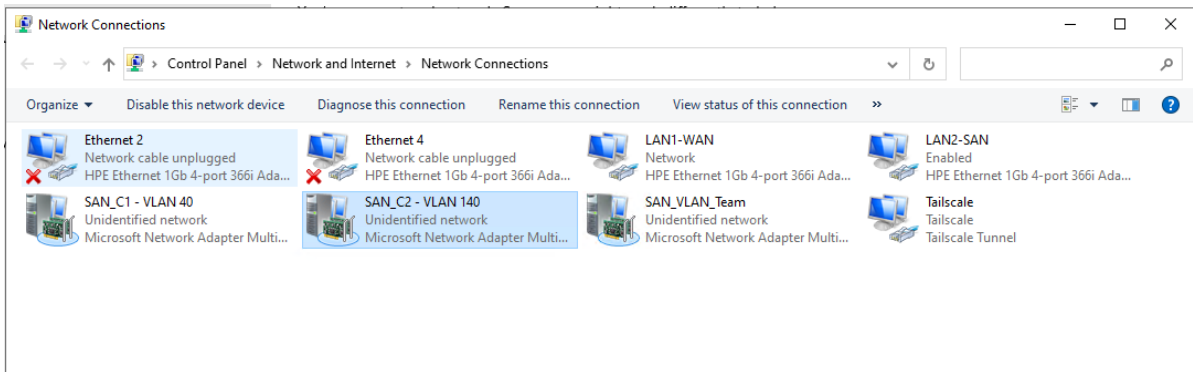


Figure 30. Server2 network interface layout

3. Discussion

```
ssh > 172.30.64.3
admin@c1-Client2-Linux:~$ hostnamectl --static
18
19 realm list
20 getent passwd 'employee@c1.local'
21 getent passwd 'employee@c2c1.local'
22 su - 'employee@c1.local' -c 'whoami; pwd; ls -ld /home/employee@c1.local/C1_Public /home/employee@c1.local/C1_Private'
23 su - 'employee@c2c1.local' -c 'whoami; pwd; ls -ld /home/employee@c2c1.local/C1_Public /home/employee@c2c1.local/C1_Private'
24 mount | grep -F 'C1_Public|C1_Private'
25 C1-Client2-Linux
26 c1.local
27 type: kerberos
28 realm-name: C1.LOCAL
29 domain-name: c1.local
30 configured: kerberos-member
31 server-software: active-directory
32 client-software: sssd
33 required-package: sssd-tools
34 required-package: sssd
35 required-package: libsss-sss
36 required-package: libsss-css
37 required-package: adcli
38 required-package: samba-common-bin
39 login-formats: %u@c1.local
40 login-policy: allow-realm-logins
41 employee@c1.local:~$ sshroot@10:82880013:employee1:/home/employee@c1.local:/bin/bash
42 employee@c2c1.local:~$ sshroot@104:82880013:employee2:/home/employee@c2c1.local:/bin/bash
43 Password:
44 employee@c1.local
45 /home/employee@c1.local
46 drwxr-xr-x 2 employee@c1.local domain users@c1.local 0 Mar 26 13:42 /home/employee@c1.local/C1_Public
47 drwxr-xr-x 2 employee@c1.local domain users@c1.local 4096 Mar 26 13:42 /home/employee@c1.local/C1_Private
48 Password:
49 employee@c2c1.local
50 /home/employee@c2c1.local
51 drwxr-xr-x 2 employee@c2c1.local domain users@c2c1.local 0 Mar 26 13:43 /home/employee@c2c1.local/C1_Public
52 drwxr-xr-x 2 employee@c2c1.local domain users@c2c1.local 4096 Mar 26 13:42 /home/employee@c2c1.local/C1_Private
53 //172.30.64.130/Priv1 on /home/employee@c1.local/C1_Private type cifs (rw,relatime,vers=3.0,cache=strict,upcall_target_app,username=employee,domain=c1.local,uid=828801110,forceuid,gid=828800513,forcegid,addr=172.30.64.130,file_mode=0755,dir_mode=0755,iocharset=utf8,soft,nounix,mapposix,reparse=fs,nativeosocket,symlink-native,rsize=4194384,wsize=4194384,bsize=1048576,retrans=1,echo_interval=60,actimeo=1,closetimeo=1)
54 //172.30.64.130/Pub1 on /home/employee@c1.local/C1_Public type cifs (rw,relatime,vers=3.0,cache=strict,upcall_target_app,username=employee,domain=c1.local,uid=828801110,forceuid,gid=828800513,forcegid,addr=172.30.64.130,file_mode=0755,dir_mode=0755,iocharset=utf8,soft,nounix,mapposix,reparse=fs,nativeosocket,symlink-native,rsize=4194384,wsize=4194384,bsize=1048576,retrans=1,echo_interval=60,actimeo=1,closetimeo=1)
55 //172.30.64.130/Priv1 on /home/admin/C1_Private type cifs (rw,relatime,vers=3.0,cache=strict,upcall_target_app,username=admin,domain=c1.local,uid=1001,forceuid,gid=1001,forcegid,addr=172.30.64.130,file_mode=0755,dir_mode=0755,iocharset=utf8,soft,nounix,mapposix,reparse=fs,nativeosocket,symlink-native,rsize=4194384,wsize=4194384,bsize=1048576,retrans=1,echo_interval=60,actimeo=1,closetimeo=1)
56 //172.30.64.130/Pub1 on /home/admin/C1_Public type cifs (rw,relatime,vers=3.0,cache=strict,upcall_target_app,username=admin,domain=c1.local,uid=1001,forceuid,gid=1001,forcegid,addr=172.30.64.130,file_mode=0755,dir_mode=0755,iocharset=utf8,soft,nounix,mapposix,reparse=fs,nativeosocket,symlink-native,rsize=4194384,wsize=4194384,bsize=1048576,retrans=1,echo_interval=60,actimeo=1,closetimeo=1)
57 //172.30.64.130/Priv1 on /home/employee@c2c1.local/C1_Private type cifs (rw,relatime,vers=3.0,cache=strict,upcall_target_app,username=employee2,domain=c1.local,uid=828802104,forceuid,gid=828800513,forcegid,addr=172.30.64.130,file_mode=0755,dir_mode=0755,iocharset=utf8,soft,nounix,mapposix,reparse=fs,nativeosocket,symlink-native,rsize=4194384,wsize=4194384,bsize=1048576,retrans=1,echo_interval=60,actimeo=1,closetimeo=1)
58 //172.30.64.130/Pub1 on /home/employee@c2c1.local/C1_Public type cifs (rw,relatime,vers=3.0,cache=strict,upcall_target_app,username=employee2,domain=c1.local,uid=828802104,forceuid,gid=828800513,forcegid,addr=172.30.64.130,file_mode=0755,dir_mode=0755,iocharset=utf8,soft,nounix,mapposix,reparse=fs,nativeosocket,symlink-native,rsize=4194384,wsize=4194384,bsize=1048576,retrans=1,echo_interval=60,actimeo=1,closetimeo=1)
59 admin@c1-Client2-Linux:~$
```

Figure 31. C1-Client2 domain-user session and Company 1 share mount validation

```
ssh > 172.30.64.66
Web console: https://Company.c2.local:9090/ or https://172.30.64.66:9090/
15
16 Last login: Thu Mar 26 11:22:01 2026 from 172.30.64.179
17 admin@Company:~$ hostnamectl --static
18
19 realm list
20 getent passwd 'employee@c2c1.local'
21 getent passwd 'employee@c2c2.local'
22 su - 'employee@c2c1.local' -c 'whoami; pwd; ls -ld /home/employee@c2c1.local/C2_Public /home/employee@c2c1.local/C2_Private'
23 su - 'employee@c2c2.local' -c 'whoami; pwd; ls -ld /home/employee@c2c2.local/C2_Public /home/employee@c2c2.local/C2_Private'
24 mount | grep -F 'C2_Public|C2_Private'
25 Company
26 c2.local
27 type: kerberos
28 realm-name: C2.LOCAL
29 domain-name: c2.local
30 configured: kerberos-member
31 server-software: active-directory
32 client-software: sssd
33 required-package: sssd-tools
34 required-package: sssd
35 required-package: libsss-sss
36 required-package: libsss-css
37 required-package: adcli
38 required-package: samba-common-bin
39 login-formats: %u@c2.local
40 login-policy: allow-realm-logins
41 employee@c2c1.local:~$ sshroot@106:95200013:employee1:/home/employee@c2c1.local:/bin/bash
42 employee@c2c2.local:~$ sshroot@110:95200013:employee2:/home/employee@c2c2.local:/bin/bash
43 Password:
44 employee@c2c1.local
45 /home/employee@c2c1.local
46 drwxr-xr-x 2 employee@c2c1.local domain users@c2c1.local 0 Mar 26 13:42 /home/employee@c2c1.local/C2_Public
47 drwxr-xr-x 2 employee@c2c1.local domain users@c2c1.local 0 Mar 26 13:42 /home/employee@c2c1.local/C2_Private
48 Password:
49 employee@c2c2.local
50 /home/employee@c2c2.local
51 drwxr-xr-x 2 employee@c2c2.local domain users@c2c2.local 0 Mar 26 13:42 /home/employee@c2c2.local/C2_Public
52 drwxr-xr-x 2 employee@c2c2.local domain users@c2c2.local 0 Mar 26 13:42 /home/employee@c2c2.local/C2_Private
53 //172.30.64.146/C2_Private on /home/admin/C2_Private type cifs (rw,relatime,vers=3.1.1,cache=strict,upcall_target_app,username=admin,domain=c2.local,uid=1001,forceuid,gid=1001,forcegid,addr=172.30.64.146,file_mode=0660,dir_mode=0700,soft,nounix,mapposix,reparse=fs,nativeosocket,symlink-native,rsize=4194384,wsize=4194384,bsize=1048576,retrans=1,echo_interval=60,actimeo=1,netdev)
54 //172.30.64.146/C2_Public on /home/admin/C2_Public type cifs (rw,relatime,vers=3.1.1,cache=strict,upcall_target_app,username=admin,domain=c2.local,uid=1001,forceuid,gid=1001,forcegid,addr=172.30.64.146,file_mode=0660,dir_mode=0700,soft,nounix,mapposix,reparse=fs,nativeosocket,symlink-native,rsize=4194384,wsize=4194384,bsize=1048576,retrans=1,echo_interval=60,actimeo=1,netdev)
55 //172.30.64.146/C2_Private on /home/employee@c2c2.local/C2_Private type cifs (rw,relatime,vers=3.1.1,cache=strict,upcall_target_app,username=employee2,domain=c2.local,uid=952001110,forceuid,gid=952000513,forcegid,addr=172.30.64.146,file_mode=0664,dir_mode=0775,soft,nounix,mapposix,reparse=fs,nativeosocket,symlink-native,rsize=4194384,wsize=4194384,bsize=1048576,retrans=1,echo_interval=60,actimeo=1,netdev)
56 //172.30.64.146/C2_Public on /home/employee@c2c2.local/C2_Public type cifs (rw,relatime,vers=3.1.1,cache=strict,upcall_target_app,username=employee2,domain=c2.local,uid=952001110,forceuid,gid=952000513,forcegid,addr=172.30.64.146,file_mode=0660,dir_mode=0700,soft,nounix,mapposix,reparse=fs,nativeosocket,symlink-native,rsize=4194384,wsize=4194384,bsize=1048576,retrans=1,echo_interval=60,actimeo=1,netdev)
57 //172.30.64.146/C2_Private on /home/employee@c2c1.local/C2_Private type cifs (rw,relatime,vers=3.1.1,cache=strict,upcall_target_app,username=employee1,domain=c2.local,uid=952001106,forceuid,gid=952000513,forcegid,addr=172.30.64.146,file_mode=0664,dir_mode=0775,soft,nounix,mapposix,reparse=fs,nativeosocket,symlink-native,rsize=4194384,wsize=4194384,bsize=1048576,retrans=1,echo_interval=60,actimeo=1,netdev)
58 //172.30.64.146/C2_Public on /home/employee@c2c1.local/C2_Public type cifs (rw,relatime,vers=3.1.1,cache=strict,upcall_target_app,username=employee1,domain=c2.local,uid=952001106,forceuid,gid=952000513,forcegid,addr=172.30.64.146,file_mode=0660,dir_mode=0700,soft,nounix,mapposix,reparse=fs,nativeosocket,symlink-native,rsize=4194384,wsize=4194384,bsize=1048576,retrans=1,echo_interval=60,actimeo=1,netdev)
59 admin@Company:~$
```

Figure 32. C2-Client1 domain-user session and Company 2 share mount validation

3. Discussion

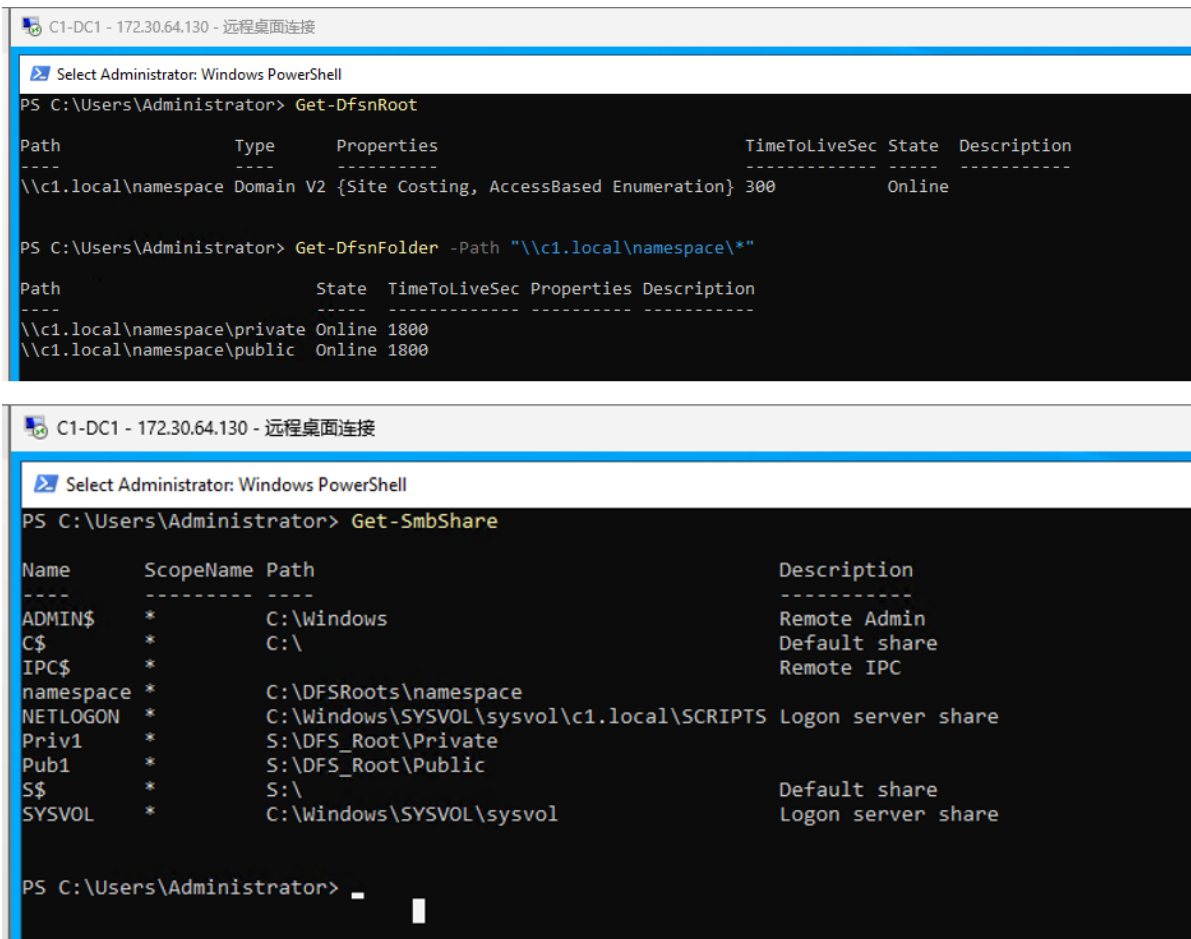


Figure 33. C1-DC1 Windows PowerShell validation for DFS namespace and SMB share publication

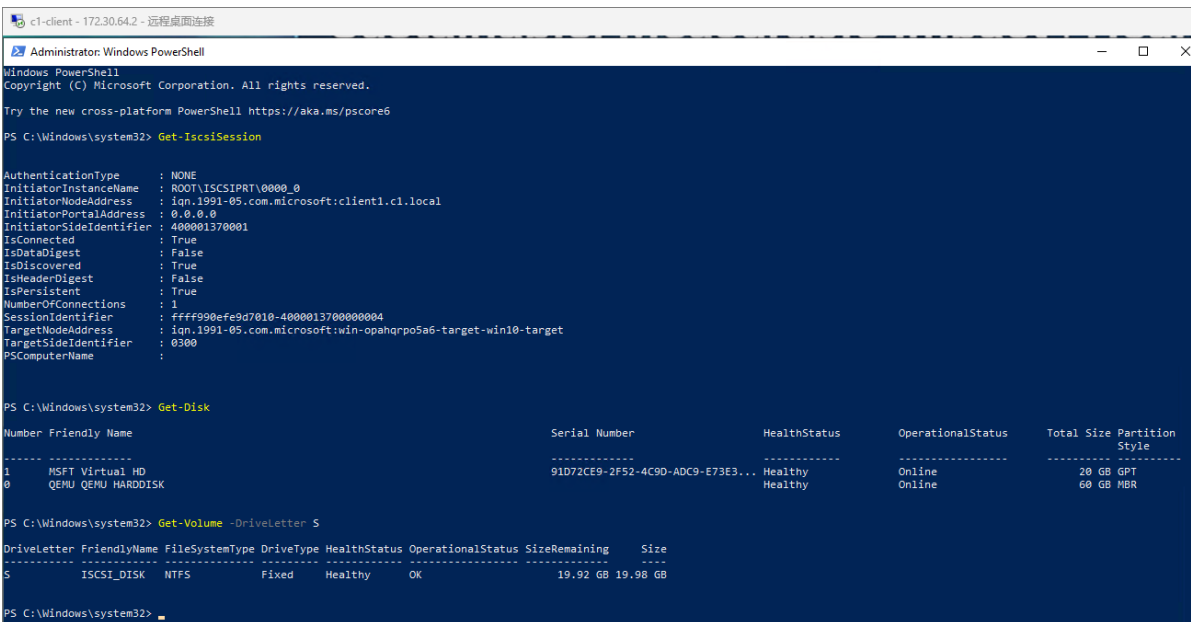


Figure 34. C1-Client1 Windows PowerShell validation for iSCSI session and mounted SAN volume

Public-cloud storage and isolated SAN evidence

Both tenants use iSCSI-backed isolated storage, and this is one of the strongest design choices in Site 2. Storage is not just data capacity. It is a fault domain. By separating block transport from the tenant LANs and keeping iSCSI off the OPNsense routing path, avoidable latency was reduced, the firewall was kept out of the storage hot path, and shares and mounted volumes were able to fail independently of tenant routing. That separation helps support staff find the right problem faster.

Company 1 expresses the pattern through C1SAN and C1FS. C1SAN sits at 172.30.65.186/29 with gateway 172.30.65.185, and it is not routed through OPNsense. C1FS is the iSCSI consumer. Inspection confirmed the active Get-IscsiSession on C1FS, the presence of F: SharedData, and the published SMB shares. In other words, block storage stayed hidden while the file server translated it into a service users can consume.

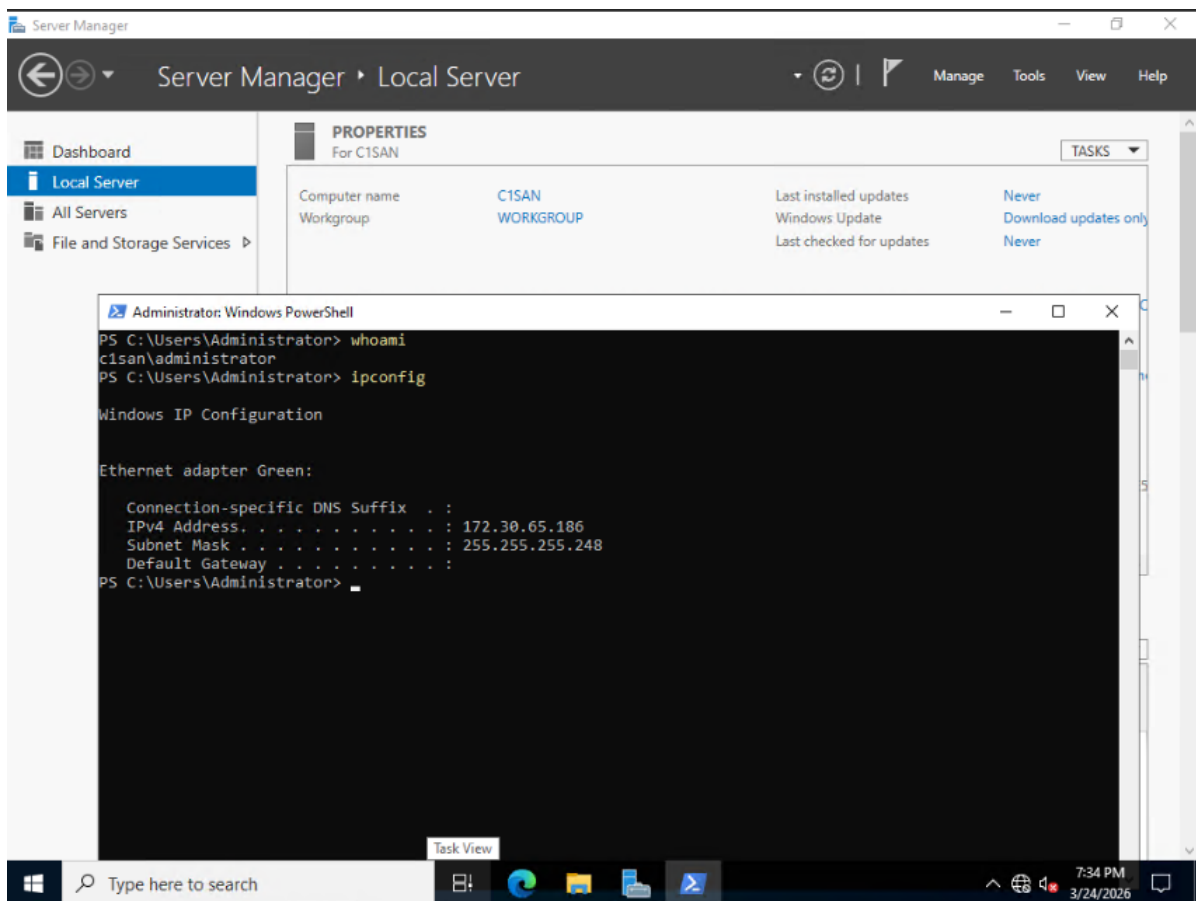


Figure 35. C1SAN isolated storage bridge interface evidence

Company 2 expresses the same idea even more visibly because C2FS has a dedicated storage NIC. C2SAN sits at 172.30.65.194/29 with gateway 172.30.65.193. C2FS reaches it through ens18 172.30.65.195/29 while serving users through ens19 172.30.65.68/26. The active session tcp:[20] 172.30.65.194:3260,1 iqn.2024-

03.org.clearroots:c2san (non-flash) and the mounted /dev/sdb volume at /mnt/c2_public prove that the isolated storage path is live.

```

bdengiz@c2san:~$ ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: ens18: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000
    link/ether bc:24:11:c7:92:f2 brd ff:ff:ff:ff:ff:ff
    altname enp0s18
    inet 172.30.65.194/29 brd 172.30.65.199 scope global ens18
        valid_lft forever preferred_lft forever
    inet6 fe80::be24:11ff:fec7:92f2/64 scope link
        valid_lft forever preferred_lft forever
bdengiz@c2san:~$
    
```

Figure 36. C2SAN isolated storage bridge interface evidence

Table 28. Storage architecture summary

Tenant	SAN Addressing	File Server Consumer	Mount or Data Volume	iSCSI Evidence	Published Shares and Isolation Method
Company 1	C1SAN 172.30.65.186 /29, gw 172.30.65.185	C1FS 172.30.65.4	F: SharedData	Active iSCSI initiator session confirmed by Get-IscsiSession on C1FS	Named SMB shares on F:; block transport stays off the routed tenant LAN because C1SAN is not routed through OPNsense.
Company 2	C2SAN 172.30.65.194 /29, gw	C2FS 172.30.65.68 with storage	/mnt/c2_public on /dev/sdb	iscsiadm -m session showed tcp:[1]	C2_Public and C2_Priva

	172.30.65.193	NIC ens18 172.30.65.195/29		172.30.65.194:3260,1 iqn.2024-03.org.clearroots:c2san (non-flash)	te; block transport is isolated on the dedicated storage NIC instead of the service NIC ens19.
--	---------------	-------------------------------	--	---	--

This is more than a neat network trick. It is an operational safety margin. If the tenant LAN is congested or misrouted, block storage is not automatically implicated. If the iSCSI path drops, the dedicated storage side should be checked before user credentials or firewall rules are blamed. That is why the storage design is explained in its own section instead of letting it disappear under the file-service headings. Operationally, isolated storage turns a vague share outage into a smaller and more testable fault domain.

3.5 Backup, Recovery, and Cross-Site Protection

Backup and recovery remain platform-level concerns because they depend on shared repositories, approved management paths, and the cross-site protection workflow defined for the environment as a whole.

Private-cloud backup baseline

Veeam Backup & Replication is hosted on Server2 as the main backup platform for Site 1. The implemented design uses three protection classes: agentless system-level VM backup, file-based client backup, and offsite backup copy to Site 2 over the site-to-site OpenVPN tunnel. This platform was retained because it provides a professional, supportable recovery workflow for Company 1 and a centralized low-operator-burden protection model for the shared MSP estate.

Table 29. Backup infrastructure components

Component	Role	System
Veeam Backup Server	Central backup management and restore orchestration	Server2
Backup Proxy	Data processing and transport for Proxmox VE backup jobs managed from Server2	Proxmox Host

Primary Backup Repository	Primary server-side storage for Proxmox VE backup chains and the C1-Client1 Veeam agent backup	Server2
Veeam Agent for Microsoft Windows	Agent-based file-level protection	C1-Client1
Offsite Backup Repository	Remote backup copy storage provided through a Veeam-managed SMB shared folder over the Site 1 to Site 2 OpenVPN tunnel	Site2-Offsite-SharedRepo (\\172.30.65.180\Site1OffsiteFromServer2\$\Repo)
Backup Copy Job	Immediate-copy duplication of seven selected Proxmox VE backup jobs to Site2-Offsite-SharedRepo	Server2

Table 30. Backup repository layout

Drive	Volume Name	File System	Purpose
V:	Backups	ReFS	Primary backup repository on Server2
Site2 R:	Site1OffsiteFromServer2\Repo	NTFS via SMB shared folder	Remote shared-folder repository on 172.30.65.180 for Site1 backup copy
Site2 Local Repository	Not documented in this handover section	NTFS	Reserved for Site2 local backup operations on 172.30.65.180 and kept separate from the Site1 offsite copy path

Table 31. Backup scope by class

Backup Class	Implemented Scope
Agentless VM backup - Company 1	C1-DC1, C1-DC2, C1-WebServer, C1-Client1, and C1-Client2 protected through Proxmox VE backup jobs
Agentless VM backup - Company 2	C2-DC1, C2-DC2, and C2-Client1 protected through Proxmox VE backup jobs
Agentless VM backup - Shared infrastructure	Jumpbox Win10 and OPNsense protected through Proxmox VE backup jobs
File-based backup	C1-Client1 designated folder path protected by Veeam Agent for Microsoft Windows in the Default Backup Repository on Server2
Offsite VPN backup copy	Seven selected Proxmox VE backup jobs copied through the Site 1 to Site 2 OpenVPN tunnel to Site2-Offsite-SharedRepo at \\172.30.65.180\Site1OffsiteFromServer2\$\Repo

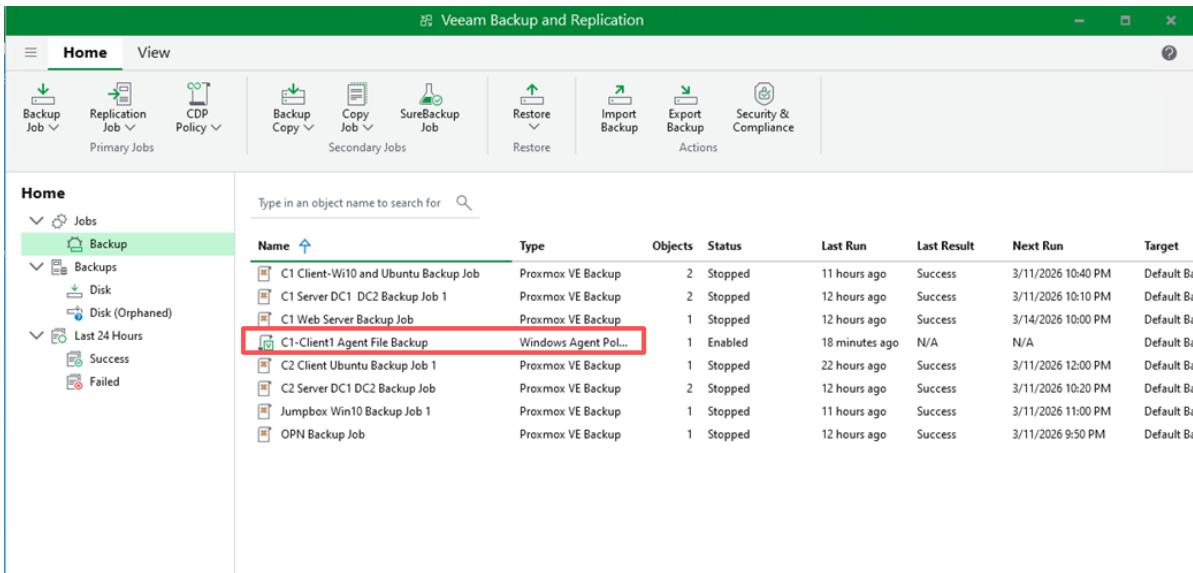


Figure 37. Veeam backup overview

Agentless System-Level VM Backup

Agentless / system-level backup class: drive V: on Server2 formatted with ReFS is the primary repository for the Proxmox VE backup chains. These jobs protect the C1 servers, C1 clients, C1-WebServer, C2 servers, C2 client, Jumpbox Win10, and OPNsense at the virtual machine level.

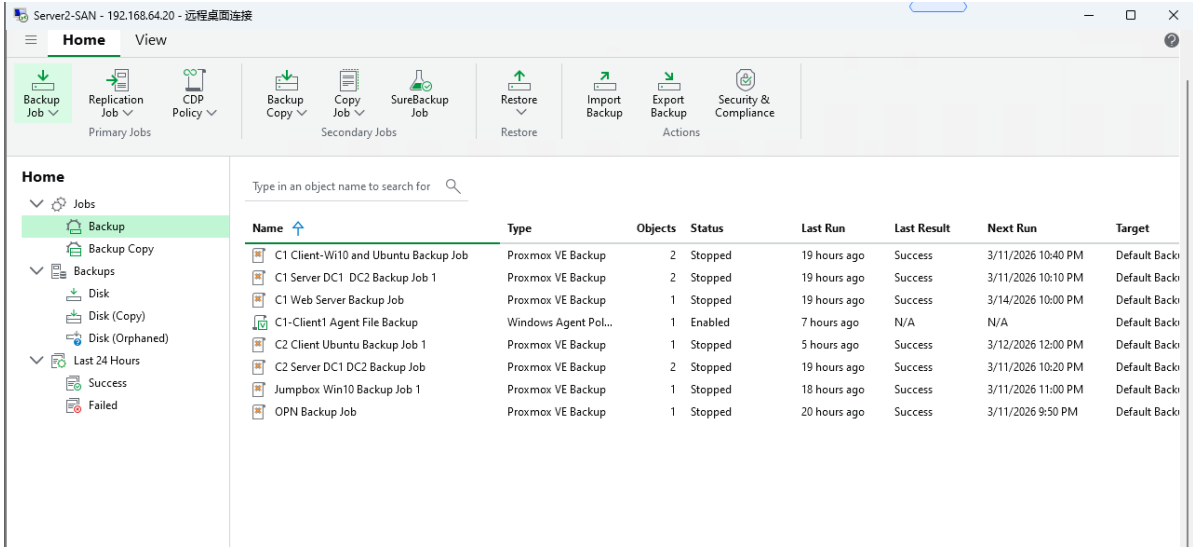


Figure 38. Primary Veeam job status

File-Based Client Backup

File-based backup class: C1-Client1 uses Veeam Agent for Microsoft Windows to back up the designated client folder path to the Default Backup Repository on Server2, which allows direct file-level restore from the Veeam console without restoring the full VM.

Offsite Backup Copy to Site 2

Selected local backup chains are copied to the Site2-Offsite-SharedRepo SMB shared-folder repository at \\172.30.65.180\Site1OffsiteFromServer2\$\Repo on the R: volume through the Site-to-Site OpenVPN tunnel, providing offsite backup protection for Site 1. For operational planning, the working recovery target is an RPO of less than 24 hours and an RTO of less than 4 hours for standard VM or file-level restores under normal lab conditions. The offsite backup copy therefore provides both an additional recovery location and a practical recovery baseline for the receiving support team.

3. Discussion

The screenshot displays the Veeam Backup and Replication interface. The main window shows a list of backup copy jobs under the 'Backup Copy' tab. The 'Site1-to-Site2 Backup Copy' job is selected, and its details are shown in a pop-up window.

Summary:

- Jobs: 7
- Objects: 10
- Processed: 10

Data:

- Processed: 32.9 GB
- Read: 32.9 GB
- Transferred: 17 GB (1.9x)

Status:

- Success: 7
- Warnings: 0
- Errors: 0

Throughput (last 24 hours):

Speed: 19 MB/s

Job Details Table:

Name	Status	Action	Duration
C1 Web Server Backu...	Success	Backup copy for C2-DC2-Secondary Backup started at 3/21/2026 4:01:07 AM	
C1 Server DC1 DC2...	Success	C2-DC2-Secondary Backup (424 MB) processing finished at 3/21/2026 4:08:50 AM: 138...	0:07:42
C1 Client-Win10 and...	Success	Backup copy for C2-DC1-Primary Backup started at 3/21/2026 4:01:07 AM	
OPN Backup Job	Success	C2-DC1-Primary Backup (469 MB) processing finished at 3/21/2026 4:09:24 AM: 157.2...	0:08:16
Jumpbox Win10 Bac...	Success	Backup copy for C1-Client1-Win10 Backup started at 3/21/2026 4:01:06 AM	
C2 Server DC1 DC2 B...	Success	C1-Client1-Win10 Backup (4.8 GB) processing finished at 3/21/2026 4:19:25 AM: 2.5 G...	0:18:18
C2 Client Ubuntu Ba...	Success	Backup copy for jumpbox-win10 Backup started at 3/21/2026 4:01:07 AM	
		jumpbox-win10 Backup (5.2 GB) processing finished at 3/21/2026 4:09:24 AM: 2.9 GB t...	0:08:17
		Backup copy for C2-Client1-Linux Backup started at 3/21/2026 4:01:07 AM	
		C2-Client1-Linux Backup (745 MB) processing finished at 3/21/2026 4:06:24 AM: 271.3...	0:05:17
		Backup copy for OPNsense Backup started at 3/21/2026 4:01:07 AM	
		OPNsense Backup (338 MB) processing finished at 3/21/2026 4:03:47 AM: 55.5 MB tra...	0:02:39

At the bottom of the pop-up window, there are buttons for 'All', 'Errors', 'Warnings', and 'Success', along with an 'OK' button.

Figure 39. Site 1 to Site 2 backup copy job

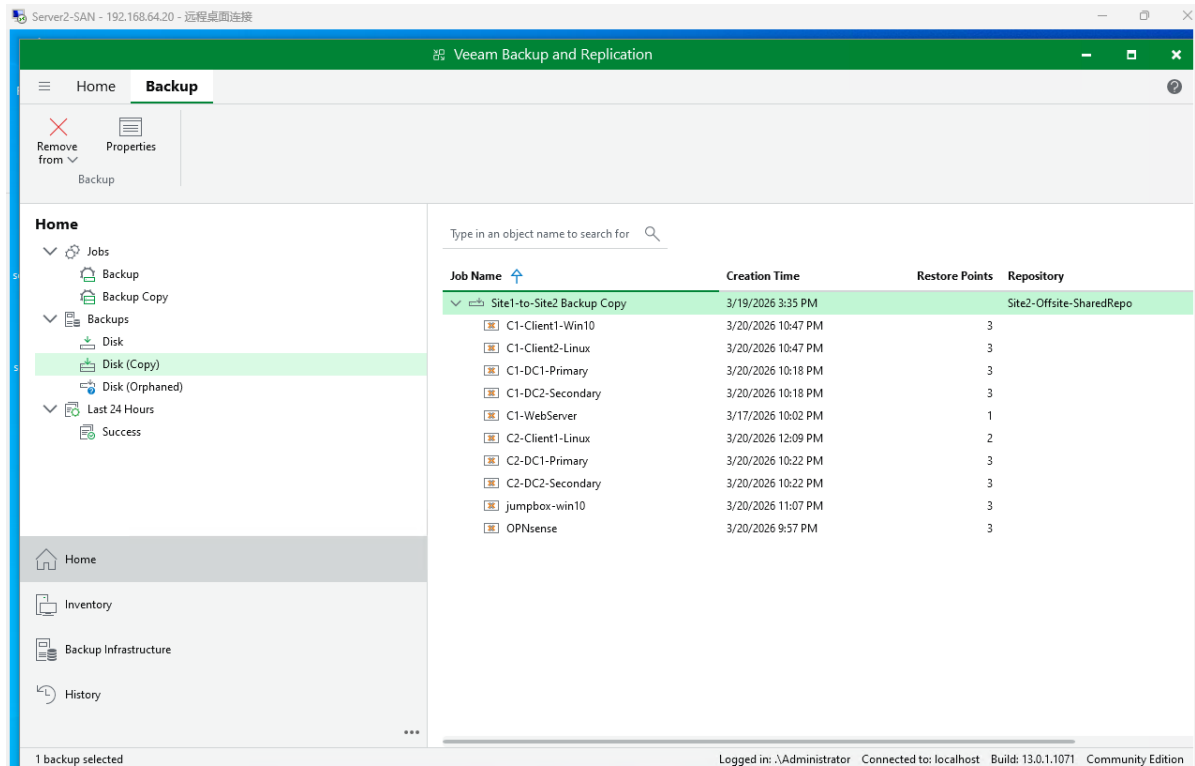


Figure 40. Remote backup copy inventory

Public-cloud offsite and repository evidence

S2Veeam at 172.30.65.180 is deliberately placed on the MSP segment, not on a tenant segment. That keeps the protection platform adjacent to the bastions and the gateway rather than burying it inside a tenant trust boundary. From MSPUbuntuJump, TCP 445, 3389, 9392, 5985, 10005, and 10006 were confirmed reachable. From Jump64, an administrative session and the active Veeam services VeeamBackupSvc, VeeamBrokerSvc, VeeamDeploySvc, VeeamExplorersRecoverySvc, VeeamFilesysVssSvc, VeeamMountSvc, and VeeamNFSSvc were confirmed.

The repository design is clear and supportable. The local repository is Site2Veeam on Z:\Site2AgentBackups. The offsite SMB target is Site1OffsiteSmbShare at \\192.168.64.20\Site2OffsiteFromSite2. The job inventory is organized into four families: Ubuntu_Servers, Windows_Servers, C1_FileShare, and C2_FileShare. Offsite copy jobs exist for both the file-share and server-backup families. The protected workload count is 10 machines. That is a coherent protection plan rather than a one-off backup proof.

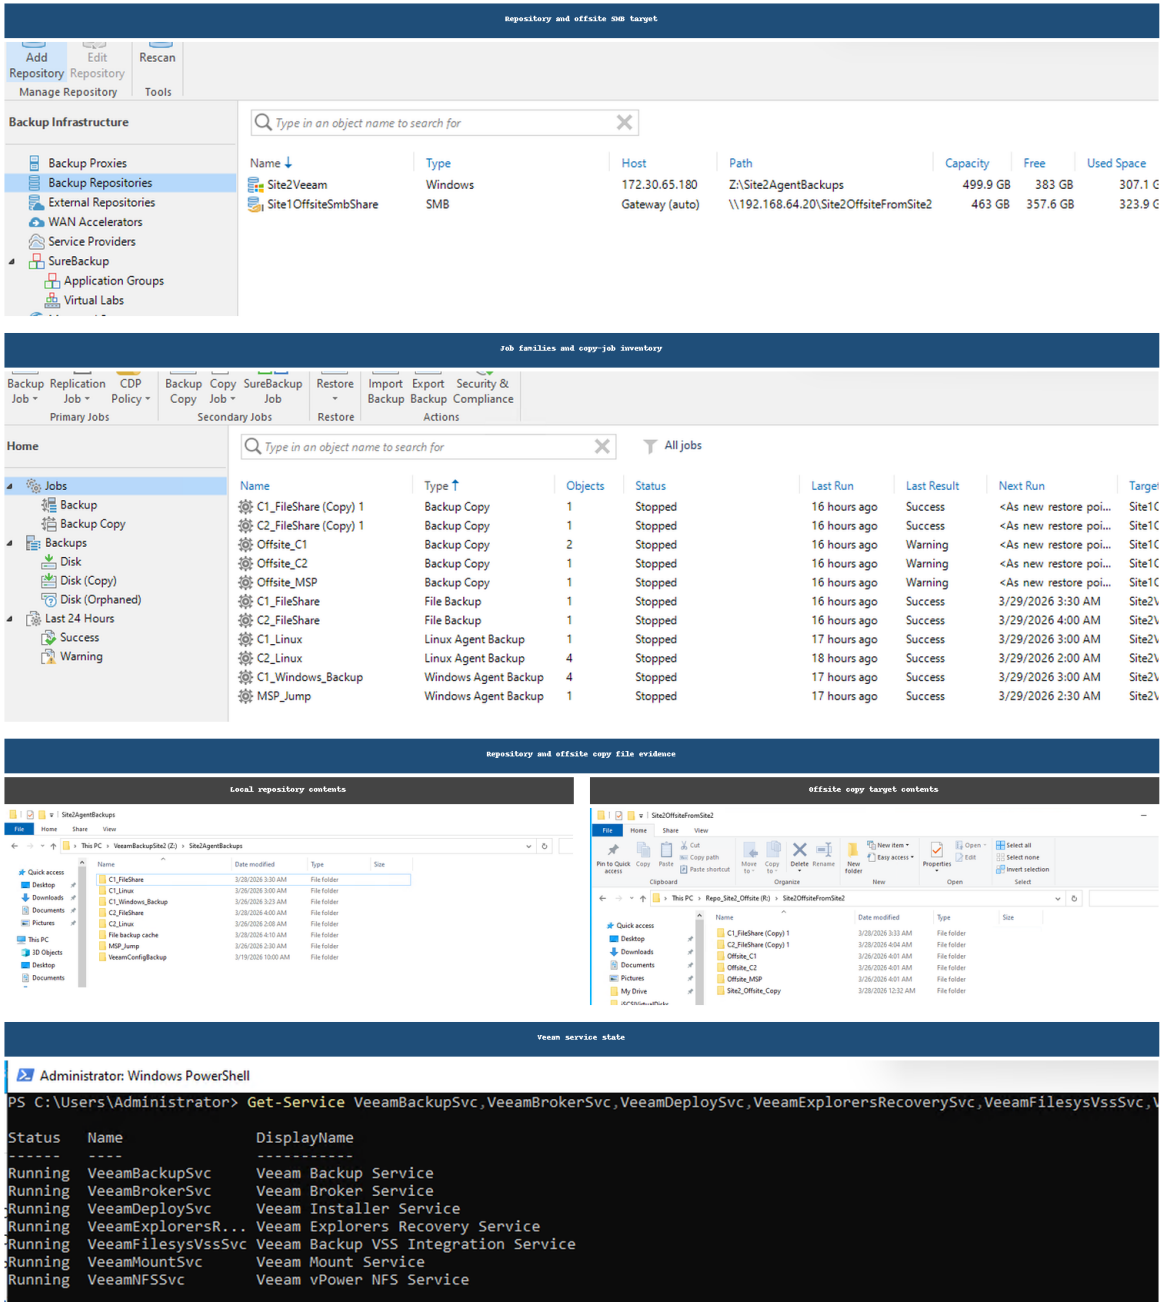


Figure 41. S2Veeam repository, job inventory, and offsite-copy configuration

Table 32. Backup and protection summary

Protection Element	Configuration	Observed State	Operational Meaning
Backup platform	S2Veeam at 172.30.65.180 on the MSP segment	Ports 445, 3389, 9392, 5985, 10005, and 10006 reachable from	Backup administration remains reachable from the

		MSPUbuntuJump	management layer even if a tenant service stack is degraded.
Active services	VeeamBackupSvc, VeeamBrokerSvc, VeeamDeploySvc, VeeamExplorersRecoverySvc, VeeamFilesysVssSvc, VeeamMountSvc, VeeamNFSSvc	Confirmed from Jump64	Core backup, broker, mounting, VSS, and recovery services are running.
Local repository	Site2Veeam on Z:\Site2AgentBackups	Repository present in the live console state	Primary local restore points stay on a known path under MSP control.
Offsite target	Site1OffsiteSmbShare at \\192.168.64.20\Site2OffsiteFromSite2	Offsite SMB target defined and reachable through the inter-site route	A second recovery location exists beyond Site 2 itself.
Job families	Ubuntu_Servers, Windows_Servers, C1_FileShare, C2_FileShare	All four families present	Protection is organized by workload class and tenant data type rather than a single monolithic job.
Copy jobs	Offsite copy jobs for both file-share and server-backup families	Present	Local backup is paired with offsite retention instead of stopping at one repository.
Protected workload count	10 machines	Confirmed in the final inspection	The platform is protecting the whole site set, not a partial sample.

It is also important to separate synchronization from backup. C2FS showed a successful Site1->Site2 C2 sync completed successfully on 2026-03-25 02:00:31. That is useful, but it is not a backup. Sync moves the current state of selected data. Backup creates

recoverable restore points with repository control, job history, and offsite retention. Both were documented because both exist in Site 2, but sync is not a substitute for backup [17]. That distinction matters because confusing sync with backup creates the wrong recovery expectations at exactly the moment the site needs discipline.

Inter-site VPN protection path




Figure 42. Site 1 OpenVPN client status




Figure 43. Site 2 OpenVPN server status

Site-to-site connectivity between Site 1 and Site 2 is implemented with OpenVPN. The current configuration state shows Site 1 operating as the OpenVPN client and Site 2 operating as the OpenVPN server. The tunnel uses TCP port 33664, tun mode, subnet topology, certificate validation under the OpenVPN-S2S-CA, and a shared static TLS key named openvpn-s2s-key.

On Site 1, the client instance is described as Site1-to-Site2 and uses the certificate site1-openvpn-client to connect to the remote endpoint 10.50.17.31. On Site 2, the server instance is described as Site2-S2S-OpenVPN and uses the certificate site2-openvpn-server with a tunnel network of 192.168.65.0/24. The routed intent carried across the tunnel remains 172.30.64.0/24 for Site 1, 172.30.65.0/24 for Site 2, and a host-specific path for Server2 at 192.168.64.20/32 so that the Site 1 Veeam server can reach the Site 2 shared-folder repository hosted on 172.30.65.180. In the current firewall state shown in the OPNsense screenshots, the OpenVPN interface now includes two explicit IPv4 block rules between Alias_C1_GLOBAL and Alias_C2_GLOBAL in both directions, while a broader IPv4 any-to-any pass rule labeled "Allow OpenVPN traffic for testing" still remains underneath them. The effective inter-site backup path is narrowed further on the WAN interface by host-specific ICMP and TCP 445 rules between 192.168.64.20 and 172.30.65.180. Tenant separation across the VPN is therefore reinforced by both the bidirectional OpenVPN block rules and the routed VLAN interface rules.

Table 33. Inter-site VPN routing and firewall control summary

Control	Site 1 / Source	Site 2 / Destination	Security Intent
---------	-----------------	----------------------	-----------------

Area			
Shared VPN transport	Site1-to-Site2 OpenVPN client	Site2-S2S-OpenVPN server	Single inter-site tunnel; Site 1 currently applies explicit IPv4 block rules between Alias_C1_GLOBAL and Alias_C2_GLOBAL in both directions before a remaining any-to-any OpenVPN testing pass rule.
OpenVPN tenant isolation	Alias_C1_GLOBAL and Alias_C2_GLOBAL	Alias_C2_GLOBAL and Alias_C1_GLOBAL	Bidirectional OpenVPN block rules now deny cross-tenant C1-to-C2 and C2-to-C1 traffic across the VPN before the broader testing pass can match.
WAN backup exception	192.168.64.20/32	172.30.65.180/32	WAN rules permit ICMP and TCP 445 so Server2 can reach the Site 2 SMB shared-folder repository used for backup copy.
Company 1 client controls	C1_ClientsVLAN10 net / 172.30.64.2	Alias_C1_GLOBAL, C1_DMZVLAN30, !RFC1918_Private, 192.168.64.20	Same-tenant access is allowed, Alias_C2_GLOBAL is blocked, TCP 80-443 is allowed to the C1 DMZ, non-RFC1918 internet access is allowed, and C1-Client1 has host-specific backup access to Server2. Company 1 SAN/iSCSI traffic

			uses the separate VLAN 40 direct-storage path between the initiator SAN NICs and Server2 rather than the routed client firewall path.
Company 1 server controls	C1_ServersVLAN20 net	Alias_C1_GLOBAL, SCHOOL_DNS, !RFC1918_Private	C1 servers can reach C1 resources, school DNS is allowed, non-RFC1918 internet access is allowed, and Alias_C2_GLOBAL is blocked.
Company 1 DMZ controls	C1_DMZVLAN30 net / any	SCHOOL_DNS, !RFC1918_Private, C1_DMZVLAN30 net	The C1 DMZ can use school DNS and non-RFC1918 internet access, and outside access to the published C1 web server is permitted.
Company 2 client controls	C2_ClientsVLAN110 net	Alias_C2_GLOBAL, C2_DMZVLAN130, !RFC1918_Private	Same-tenant access is allowed, Alias_C1_GLOBAL is blocked, TCP 80-443 is allowed to the C2 DMZ, and internet access is limited to non-RFC1918 destinations.
Company 2 server controls	C2_ServersVLAN120 net	Alias_C2_GLOBAL, SCHOOL_DNS, !RFC1918_Private	C2 servers can reach C2 resources, school DNS, and non-RFC1918 internet destinations while

			Alias_C1_GLOBAL is blocked.
Company 2 DMZ controls	C2_DMZVLAN130 net / any	SCHOOL_DNS, IRFC1918_Private, C2_DMZVLAN130 net	The C2 DMZ can use school DNS and non-RFC1918 internet access, and outside access to the published c2-webserver is permitted.

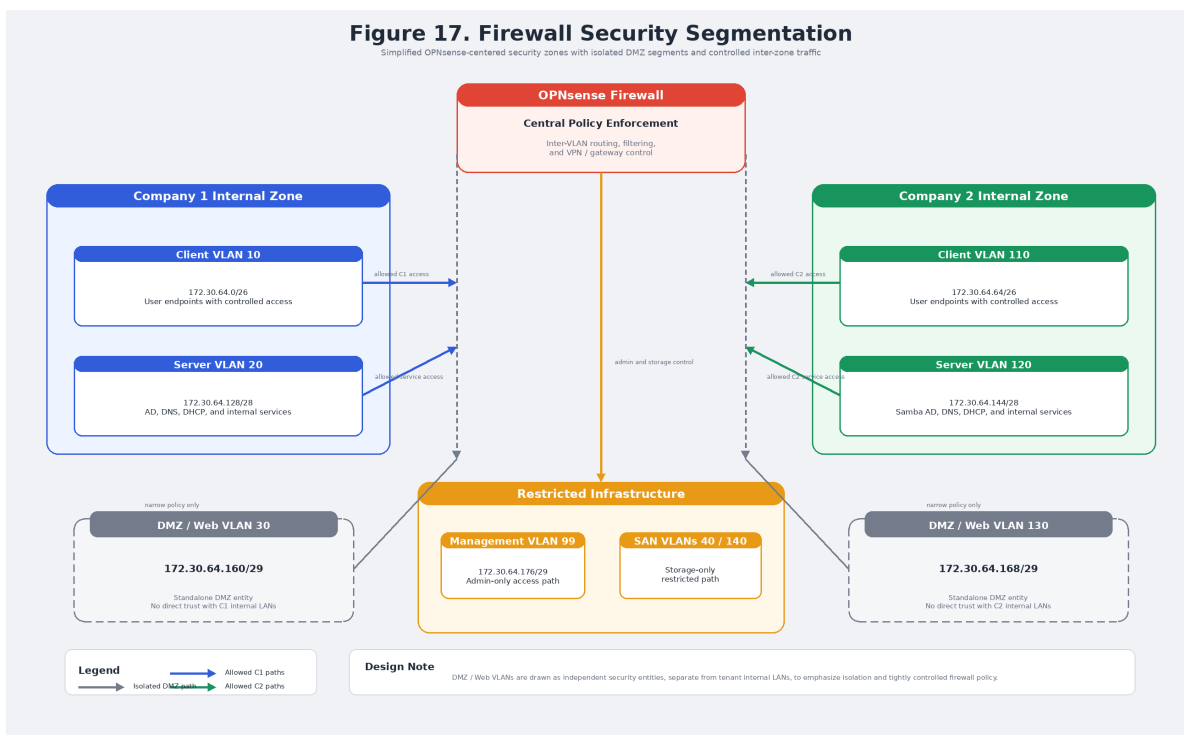


Figure 44. Firewall-centered security segmentation

SAN VLANs should be limited to iSCSI traffic between Server2 and authorized initiators. Cross-tenant connectivity should remain blocked unless explicitly justified and approved. Firewall changes should be documented with pre-change backup and post-change validation steps. In the current implementation, the OpenVPN interface blocks Alias_C1_GLOBAL to Alias_C2_GLOBAL and Alias_C2_GLOBAL to Alias_C1_GLOBAL before a remaining broad testing pass rule, C1 and C2 client/server interfaces use alias-based allow and block rules to preserve tenant separation, the DMZ interfaces allow SCHOOL_DNS and non-RFC1918 internet access while permitting published outside access to the tenant web servers, and the WAN interface uses host-specific ICMP and

TCP 445 rules to permit Server2 (192.168.64.20) to reach the Site 2 SMB repository host (172.30.65.180).

3.6 Administrative Access, Dependencies, Maintenance, and Triage

The final core section preserves the parts of both handovers that are most useful once the environment enters steady-state support: approved administrative access, dependency interpretation, maintenance discipline, and fast triage.

Administrative access, risks, and support reading from the private-cloud handover

Routine administrative access should originate from Jumpbox Windows (172.30.64.179) or Jumpbox Ubuntu (172.30.64.180) on VLAN 99. Browser-based tools such as Grafana, Cockpit, and Windows Admin Center sit behind that same management model, while direct tenant-side administration should remain the exception rather than the default. This keeps support traffic legible and reduces the number of systems exposed for day-to-day control.

Maintenance, Change Workflow, and Failure Domains

The following guidance is written for the receiving support team and should be treated as the minimum operating rhythm for Site 1. These checks are not a substitute for monitoring and alerting, but they do establish the recurring actions most likely to preserve service continuity.

Table 34. Recommended operational checks

Cadence	Required Checks	Evidence / Outcome
Daily	Review Proxmox host health, verify Veeam backup and backup copy job success, verify OpenVPN site-to-site availability when intersite copy is required, confirm firewall/gateway availability, check Grafana for abnormal CPU, memory, or I/O trends.	Dashboard screenshots or ticket note
Weekly	Review local and offsite repository free space, validate jumpbox accessibility, confirm tenant DNS/DHCP health, verify reachability to the Site2 repository host at 172.30.65.180, inspect SAN interface state on Server2.	Ops checklist or change record

Monthly	Apply approved patch cycle, export configuration backups where available, review firewall rule exceptions, test one representative restore path from the local or offsite repository.	Change record and restore evidence
After Any Major Change	Re-test authentication, DNS resolution, VM state, backup success, backup copy reachability, monitoring visibility, and OpenVPN site-to-site reachability.	Post-change validation record

Standard Change Workflow

Confirm an approved maintenance window and rollback owner before making infrastructure changes.

Verify the latest successful backup or snapshot for affected systems.

Capture current-state evidence for firewall rules, network settings, VM configuration, or storage mappings before modification.

Validate service health after the change by checking authentication, routing, storage access, and backup job status.

Record the change outcome, operator, timestamp, and any follow-up action in the client change log.

Failure Domains and Recovery Priority

Table 35. Primary failure domains

Component	Operational Impact	Support Priority
OPNsense Firewall	Loss of routing between VLANs and potential loss of internet egress.	Critical
Proxmox VE Host	Loss of hosted tenant workloads and management access to virtual machines.	Critical
Server2	Loss of SAN presentation and backup repository availability.	Critical
Tenant Domain Controllers	Authentication, DNS, and DHCP degradation for the affected tenant.	High
Jumpbox Systems	Reduced remote administration capability	Medium

	from the management network.	
--	------------------------------	--

Risks, Deferred Improvements, and Support Assumptions

The architecture is operationally usable, but several dependencies and deferred improvements should be acknowledged during handover. The most significant are the single Proxmox host, the shared dependence on Server2 for SAN and backup functions, and the need to formalize operational recovery targets and alert-routing standards. Table 36 summarizes these items in a form suitable for service review and future roadmap planning. A next-phase expansion would add a second Proxmox node, shared cluster-aware storage, and quorum design so that live migration and more resilient maintenance windows could be introduced without redesigning the tenant network model.

Table 36. Risks, deferred improvements, and support assumptions

Item	Current State	Support Reading / Required Action
Single Proxmox host dependency	All Site 1 workloads depend on one Proxmox VE host.	Introduce a second Proxmox node and cluster design to reduce business-interruption risk, or document downtime expectations for the current single-host design.
Shared dependence on Server2	Server2 remains the shared storage and backup backbone for both tenants.	Document contingency planning for SAN and backup service interruption, or add service redundancy in a future phase.
Backup retention / RPO / RTO	Working targets of RPO < 24 hours and RTO < 4 hours are documented in Section 3.5, but formal client-approved retention policy and service-level recovery objectives remain to be finalized.	Confirm retention periods, offsite policy, and client-approved service-level recovery objectives in Veeam documentation for steady-state support.
Monitoring alert routing	Dashboard visibility exists, but alert thresholds, recipients, and escalation ownership are still provisional.	Formalize alert thresholds, recipients, and escalation path so monitoring operates as a managed service rather than dashboard-only visibility.

Service dependency, routine checks, and troubleshooting from the Site 2 handover

Support teams do not need a flat inventory nearly as much as they need a dependency map. Site 2 is a layered environment, so failure impact is layered too. If rp-msp-gateway goes down, routing, policy, remote entry, and the inter-site path fail together. If both C2IdM1 and C2IdM2 fail, Company 2 loses authentication, DNS, and DHCP together. If C2FS loses its iSCSI session, the shares disappear even if users can still authenticate. If S2Veeam becomes unreachable, new backups stop even if the tenant workloads themselves are still online.

Table 37. Service dependency map

Service	Depends On	Downstream Impact If It Fails	First Check
rp-msp-gateway	WAN connectivity, internal interfaces, firewall policy, NAT, SITE1_OVPN	Remote entry, tenant routing, cross-site access, and backup copy paths stop together	Check WAN status, interface state, NAT publication, and tunnel health first.
Jump64	rp-msp-gateway MSP segment and WAN NAT 33464	Windows-side administration of Company 1 and S2Veeam becomes difficult	Confirm 33464 publication and 172.30.65.178 reachability.
MSPUbuntuJump	rp-msp-gateway MSP segment and WAN NAT 33564	Linux-side inspection of Company 2 and port validation is lost	Confirm 33564 publication and 172.30.65.179 reachability.
Company 1 identity	C1DC1, C1DC2, C1LAN, Company 1 DNS	Company 1 authentication, policy, and name resolution degrade or fail	Start with C1DC1 and C1DC2 service state and reachability.
Company 1 file services	C1FS, C1SAN path, SMB service, iSCSI session	Shares disappear or become stale even if the Company 1 domain is healthy	Inspect SMB state, F: SharedData, and Get-IscsiSession on C1FS.
Company 1 web services	C1WebServer HTTPS binding, Company 1 DNS, C1DMZ path	Users receive resolution errors or HTTP 404 instead of the site	Validate hostname resolution and the c1-webserver.c1.local binding on TCP 443.
Company 2 identity	C2IdM1 or C2IdM2, Samba AD, DNS, DHCP	Authentication, name resolution, and address assignment fail for Company 2 if both	Check samba-ad-dc and isc-dhcp-server on C2IdM1 first, then C2IdM2.

		nodes are lost	
Company 2 file services	C2FS, C2SAN path, /mnt/c2_public, smbd	C2_Public and C2_Private vanish or become inaccessible	Check iscsiadm -m session, findmnt /mnt/c2_public, and systemctl is-active smbd.
Company 2 web services	C2WebServer nginx binding, Company 2 DNS, C2DMZ path	Users receive resolution errors or HTTP 404 instead of the site	Validate server_name c2-webserver.c2.local and curl behavior with and without the Host header.
S2Veeam	MSP segment reachability, repository path, inter-site route to 192.168.64.20	No new backups run and offsite copy jobs stall	Check Veeam services, repository visibility, and the SITE1_VEEAM route.

The practical lesson is to start with the component that owns the symptom. A share problem is usually not a domain controller problem first. A web 404 is usually not an OPNsense problem first. A remote-access outage is almost always an MSP edge problem before it is a tenant problem. The dependency table turns those instincts into a repeatable operating model. For day-two support, this dependency map is what shortens outages when multiple stacks are healthy enough to distract the operator.

Maintenance, Daily Duties, and Operational Checks

Routine support for Site 2 should begin on the bastions, not directly on tenant servers. The bastions are the supported entry points, so daily checks should preserve that discipline. Start by confirming that Jump64 and MSPUbuntuJump are reachable through the NAT publications. Then verify that SITE1_OVPN is up, because the inter-site path matters for both cross-site web access and the Veeam copy path.

Table 38. Routine operational checks

Check	How To Verify	Expected Result	Why It Matters
Jump host reachability	Confirm WAN NAT 33464 and 33564 reach Jump64 and MSPUbuntuJump	At least one bastion is reachable and both are preferred	Every support action begins on the bastion layer.
OPNsense tunnel state	Inspect SITE1_OVPN state and the route to 192.168.64.20/32	Inter-site path is up and the Site 1 backup endpoint is routable	Cross-site web access and offsite copy depend on the tunnel.
Company 1 directory health	From Jump64, rerun Get-ADDomainController or the approved WinRM-based domain check against C1DC1 and	Both controllers return domain and service state successfully	Company 1 logons, DNS, and policy checks depend on this

	C1DC2		foundation.
Company 1 file service path	From Jump64, verify Get-SmbShare, Get-IscsiSession, and the F: SharedData volume on C1FS	SMB shares are present and the iSCSI session is active	This proves that Company 1 file publication and its hidden storage path are both healthy.
Company 2 identity services	systemctl is-active samba-ad-dc and systemctl is-active isc-dhcp-server on C2IdM1 and C2IdM2	Both services active on both nodes; PRIMARY on C2IdM1 and SECONDARY on C2IdM2	Company 2 authentication, DNS, and DHCP are concentrated here.
Company 2 file service mount	findmnt /mnt/c2_public and iscsiadm -m session on C2FS	/dev/sdb mounted and the session to 172.30.65.194:3260,1 active	If the block layer is gone, the shares will fail next.
Synchronization health	Review /var/log/c2_site1_sync.log	Most recent Site1->Site2 C2 sync completed successfully	Sync issues do not break backup, but they do affect replicated data freshness.
Internal web delivery	curl -kI https://c1-webserver.c1.local and curl -kI https://c2-webserver.c2.local	Hostname requests return HTTP 200, while raw IP checks still return HTTP 404	The site intentionally serves by hostname rather than by address.
Backup status	Review Veeam job history and service state on S2Veeam	Four job families and copy jobs complete without new failures	Protection is only useful if it continues to run after deployment day.

From there the daily work becomes specific and fast: confirm samba-ad-dc and isc-dhcp-server on C2IdM1 and C2IdM2, confirm the C2FS mount and iSCSI session, review the sync log, confirm that both web hostnames return HTTP 200 while raw IP access still returns HTTP 404, and verify Veeam job state. For Company 1, the concrete daily checks on Jump64 are the controller inventory command against C1DC1 and C1DC2 plus Get-SmbShare, Get-IscsiSession, and the F: SharedData check on C1FS. That is what turns routine maintenance from a checklist into a stable support habit. When each tenant has a clear and repeatable inspection path, daily checks stay fast, comparable, and useful under pressure.

Troubleshooting Guide

The fastest triage in Site 2 starts with the symptom and the owning layer. The guide below was written for the failures a support team is most likely to see first: no remote

access, broken name resolution, wrong web response, unavailable shares, failed authentication, and failed backups. In each case the goal is the same: inspect the first system that can answer the question cleanly instead of wandering through unrelated hosts.

Table 39. Fast triage guide

Symptom	First System To Check	Quick Command or Check	Healthy Expectation	Likely Fault Domain If Unhealthy
Cannot reach Site 2 remotely at all	rp-msp-gateway, then Jump64 and MSPUbuntu Jump	Verify WAN publication 33464 -> 172.30.65.178:3389 and 33564 -> 172.30.65.179:22; confirm bastion reachability	At least one bastion responds and both NAT rules are present	WAN edge, NAT publication, provider path, or bastion outage
c1-webserver.c1.local or c2-webserver.c2.local does not resolve	Tenant DNS authorities: C1DC1 or C1DC2 for Company 1 names; C2IdM1 or C2IdM2 for Company 2 names	Run nslookup or dig from the affected client; on Company 2 identity nodes confirm the c1.local and c2.local zones and both A records	DNS returns 172.30.64.162 and 172.30.65.162 for c1-webserver.c1.local, and 172.30.64.170 and 172.30.65.170 for c2-webserver.c2.local	Tenant DNS outage, missing resolver search scope, or bad zone data
Web hostname returns 404 instead of 200	C1WebServer or C2WebServer	Check the hostname-based binding: IIS on c1-webserver.c1.local TCP 443, or nginx server_name c2-webserver.c2.local on port 443	Hostname request returns 200 while raw IP still returns 404	Binding drift, wrong host header, or DNS pointing at the wrong server
Company 1 shares appear unavailable	C1FS (via Jump64)	From Jump64 run Get-SmbShare, Get-IscsiSession, and confirm the F: SharedData volume	Named SMB shares remain present, the iSCSI initiator session is active, and F: stays mounted	Windows file service fault, storage-session loss, or hidden SAN path issue
Company 2 shares appear unavailable	C2FS	Run systemctl is-active smbd, findmnt /mnt/c2_public, and iscsiadm -m session	smbd is active, /dev/sdb is mounted, and the	Samba service fault, mount loss, or

			172.30.65.194: 3260, 1 session is present	broken iSCSI path
Company 1 user cannot authenticate	C1DC1, then C1DC2	From Jump64 rerun Get-ADDomainController or the approved WinRM directory check; from the client confirm c1.local membership and DNS	Both controllers return healthy directory state and the client still points at Company 1 DNS	Active Directory, DNS, or client trust-path fault
Company 2 user cannot authenticate	C2IdM1, then C2IdM2	Run systemctl is-active samba-ad-dc; from C2LinuxClient confirm realm list and getent passwd employee1@c2.local	Samba AD is active and domain users still resolve	Directory, DNS, or client realm integration fault
Backup job shows as failed	S2Veeam	Check VeeamBackupSvc and companion services, confirm Site2Veeam on Z:\Site2AgentBackups, and verify access to \\192.168.64.20\Site2OffsiteFromSite2	Core Veeam services are running, repository is present, and the offsite target is reachable	Backup service failure, repository issue, or inter-site copy path problem
Domain user login succeeds but C2_Public and C2_Private are not mounted on C2LinuxClient	C2LinuxClient, then C2IdM1	On C2LinuxClient: sudo journalctl -b grep -Ei 'pam_mount cifs' to identify mount failure; kvno cifs/c2fs.c2.local to test Kerberos ticket acquisition. On C2IdM1: sudo samba-tool spn list C2FS\$ to verify cifs/c2fs and cifs/c2fs.c2.local SPNs are present.	kvno returns kvno = 1 for cifs/c2fs.c2.local; samba-tool spn list shows both cifs/c2fs and cifs/c2fs.c2.local; after login mount grep c2fs shows both shares mounted via sec=krb5.	Missing CIFS SPN on C2FS\$; KDC unreachable from C2LinuxClient due to DNS misconfiguration; pam_mount volume options incompatible with the kernel CIFS/Kerberos path.

That approach matters in this environment because Site 2 mixes Windows and Linux platforms on purpose. The same outward symptom can live in IIS, nginx, Samba AD, Windows AD, SMB, WMI, WinRM, or iSCSI. Good triage is therefore not about memorizing one favorite command set. It is about knowing which layer owns the failure

and then using the right bastion to read it. For the reader inheriting the site, this section turns the document from an inventory into an action sequence.

3.7 Value-Added Operational Tooling and Public Service Extension

The project includes value-added components beyond the minimum tenant-service baseline. They remain in the main body because they were implemented, validated, and should still be understood by the receiving support team.

Operational tooling and management support from the private-cloud handover.

In addition to the core services required to make the environment functional, several operational enhancements were retained because they improve manageability, visibility, and day-to-day support efficiency. These items are worth documenting separately because they change how the environment is operated even though they are not the primary service backbone.

3.7.1 Standardized Endpoint Policy Control (Company 1)

A Company 1 Group Policy Object was used to demonstrate centralized endpoint policy delivery on managed workstations. In this environment the most visible control is branded wallpaper, but the operational significance is the same as any other domain-delivered desktop policy: the organization can target client configuration from AD without touching each workstation manually.

Table 40. Group Policy branding configuration summary

Setting	Value
Domain	c1.local
Domain Controller	C1-DC1
GPO Name	GPO_Wallpaper_Client1
Linked OU	Computers
Target System	CLIENT1
Policy Type	Desktop Wallpaper Policy

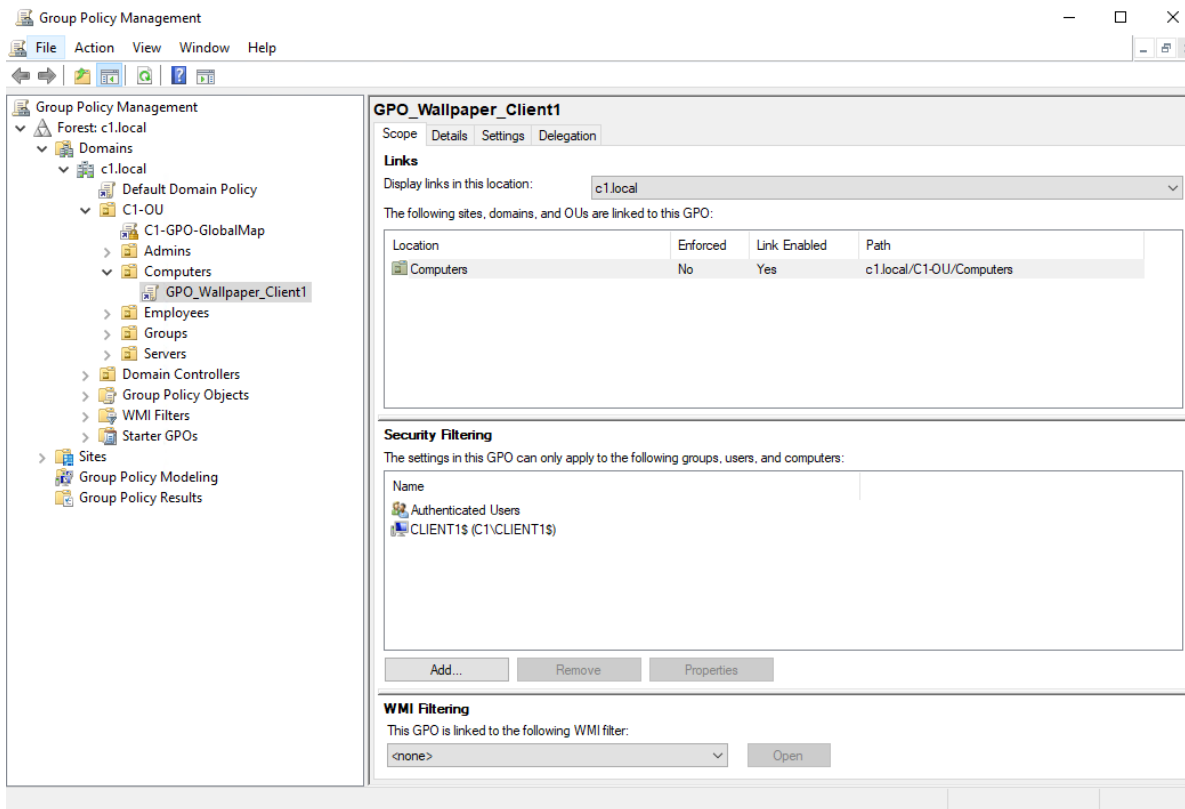


Figure 45. Group Policy desktop branding evidence

3.7.2 Infrastructure Monitoring Dashboard (Grafana)

A containerized Grafana and InfluxDB stack was deployed on the Ubuntu jumpbox (172.30.64.180) to provide a lightweight infrastructure visibility layer for Site 1. Proxmox VE exports host and workload metrics to InfluxDB, and Grafana presents them through a browser-based dashboard that is suitable for daily health checks and trend review from the management path.

Table 41. Grafana dashboard metrics

Metric	Description
Server CPU Usage	Real-time CPU utilization of the Proxmox host
Load Average	Current system load average
I/O Wait	Disk I/O waiting time
Memory Usage	Total memory consumption of the server
Running Virtual Machines	List of active virtual machines
Running LXC Containers	Status of running Linux containers
Network Interfaces	Monitoring of NIC traffic and performance

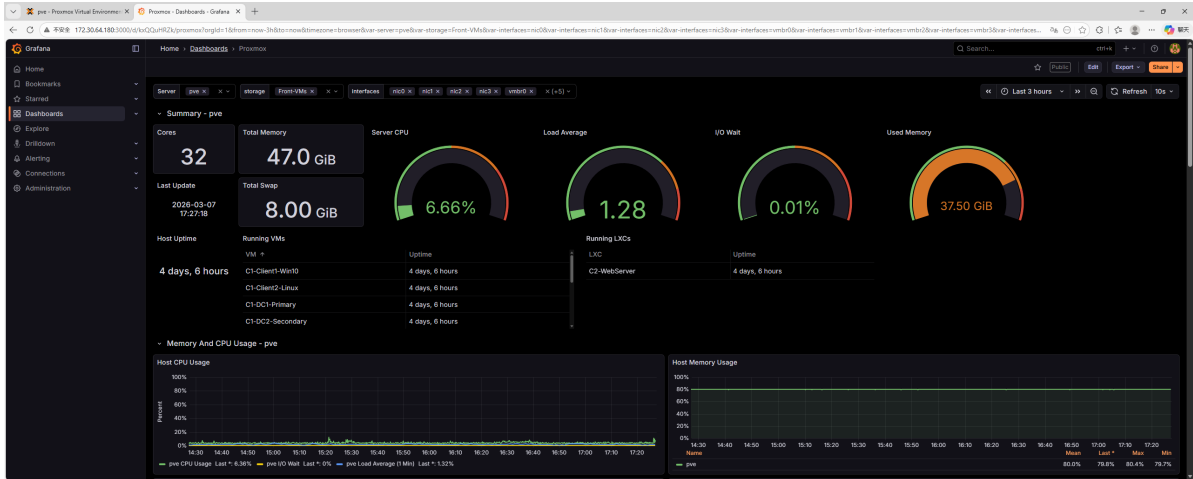


Figure 46. Grafana infrastructure dashboard

```

9  14 updates can be applied immediately.
10 3 of these updates are standard security updates.
11 To see these additional updates run: apt list --upgradable
12
13 1 additional security update can be applied with ESM Apps.
14 Learn more about enabling ESM Apps service at https://ubuntu.com/esm
15
16 Last login: Tue Mar 24 16:09:29 2026 from 172.30.64.179
17 admin@adminin-Standard-PC-i440FX-PIIX-1996:~$ hostnamectl --static
18 systemctl is-active grafana-server
19 systemctl is-active influxdb
20 ss -lntp | egrep '3000|8086'
21 curl -I http://127.0.0.1:3000
22 curl -I http://127.0.0.1:8086
23 admin-Standard-PC-i440FX-PIIX-1996
24 inactive
25 inactive
26 LISTEN 0      4096          0.0.0.0:3000    0.0.0.0:*
27 LISTEN 0      4096          0.0.0.0:8086    0.0.0.0:*
28 LISTEN 0      4096          :::3000         ::::*
29 LISTEN 0      4096          :::8086         ::::*
30 HTTP/1.1 302 Found
31 Cache-Control: no-store
32 Content-Type: text/html; charset=utf-8
33 Location: /login
34 X-Content-Type-Options: nosniff
35 X-Frame-Options: deny
36 X-Xss-Protection: 1; mode=block
37 Date: Thu, 26 Mar 2026 17:56:24 GMT
38
39 HTTP/1.1 404 Not Found
40 Content-Type: text/plain; charset=utf-8
41 X-Content-Type-Options: nosniff
42 X-Influxdb-Build: OSS
43 X-Influxdb-Version: 1.8.10
44 Date: Thu, 26 Mar 2026 17:56:24 GMT
45 Content-Length: 19
46
47 admin@adminin-Standard-PC-i440FX-PIIX-1996:~$

```

Figure 47. Ubuntu jumpbox CLI validation for Grafana and InfluxDB endpoints

3.7.3 Browser-Based Linux Administration (Cockpit)

Cockpit was retained as a browser-based administration layer for the Ubuntu-based Company 2 systems. It reduces dependence on direct SSH for routine checks by exposing service state, logs, basic health indicators, network details, and a web terminal through a single support-friendly interface.

Table 42. Cockpit management features

Feature	Description
System Health Monitoring	Displays system alerts and service status
CPU and Memory Monitoring	Shows real-time resource usage
System Logs	Allows administrators to view and filter system logs
Service Management	Start, stop, and manage system services
Terminal Access	Provides a browser-based terminal for command-line administration
Network Configuration	Displays network interface information

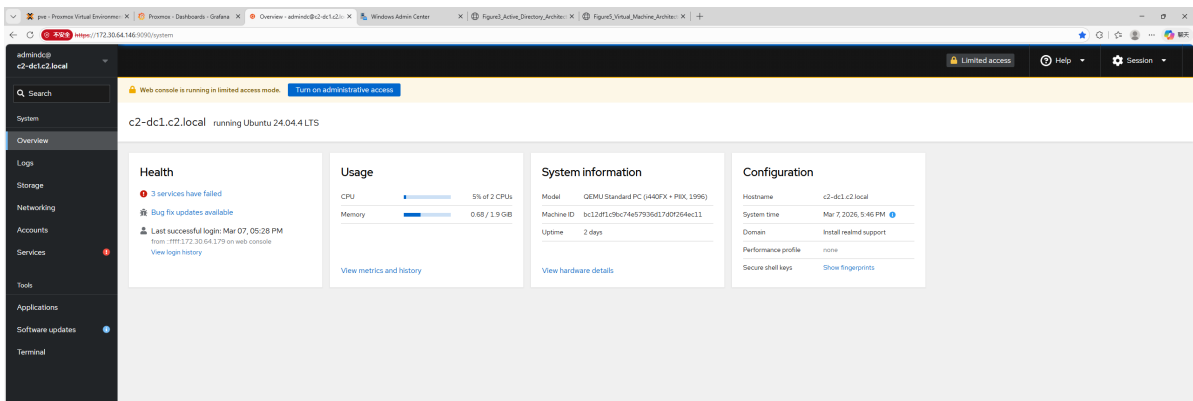


Figure 48. Cockpit administration overview

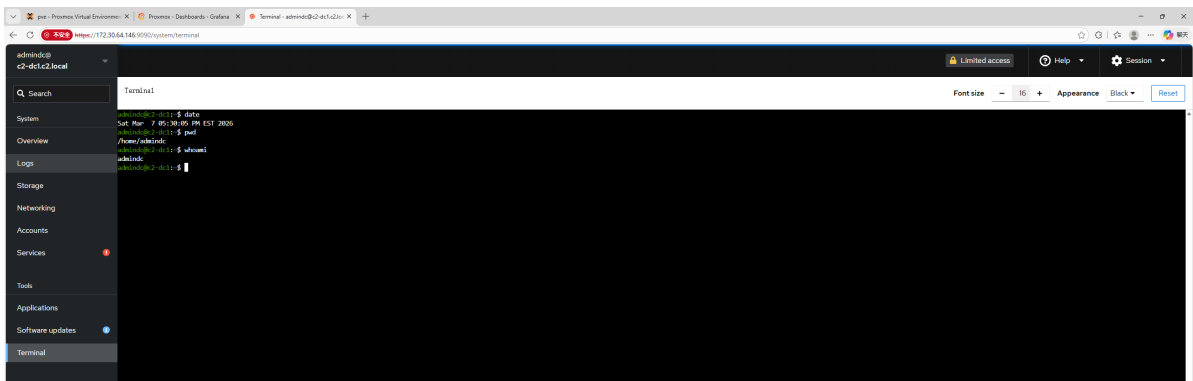


Figure 49. Cockpit web terminal

3.7.4 Centralized Windows Administration (Windows Admin Center)

Windows Admin Center was deployed on the Windows jumpbox to consolidate management of Windows-based infrastructure without requiring direct RDP to every

system for routine checks. It provides a single management surface for server overview, event review, PowerShell access, service management, and storage inspection.

Table 43. Windows Admin Center managed systems

System	IP Address	Role
C1-DC1	172.30.64.130	Company 1 Primary Domain Controller
C1-DC2	172.30.64.131	Company 1 Secondary Domain Controller
Client1	172.30.64.2	Company 1 Windows Client
Storage Server	192.168.64.20	Backup and SAN Storage Server
Jumpbox	172.30.64.179	Administrative Access System

Table 44. Windows Admin Center capabilities

Feature	Description
Server Monitoring	Displays CPU, memory, and storage usage
Event Viewer	Provides centralized access to system event logs
Remote PowerShell	Allows administrators to run PowerShell commands remotely
Service Management	Start, stop, and manage system services
Storage Management	Monitor disks and storage volumes
System Updates	Manage Windows updates for servers and clients




Figure 50. Windows Admin Center connection inventory

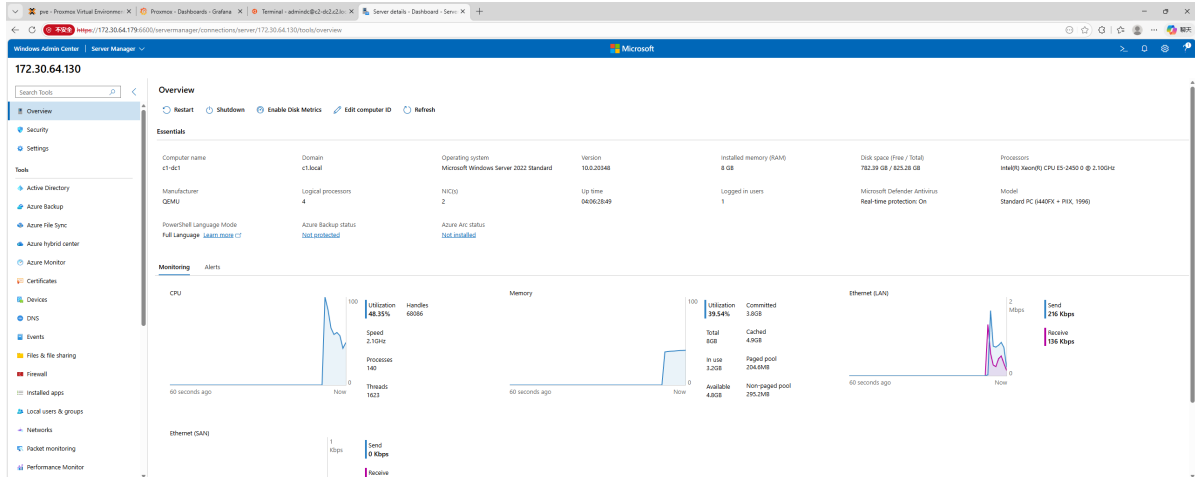


Figure 51. Windows Admin Center server overview

3.7.5 Cloud-hosted public service extension from the Site 2 handover

3.7.5.1 Value-Add Overview

The team extended the internal Company 2 web-delivery concept into a publicly reachable AWS deployment as the Site 2 Value-Add component. Internally, Company 2 already published an nginx-based HTTPS service from C2WebServer at 172.30.65.170 in C2DMZ. Externally, that same model is expressed through a containerized image, a two-node Kubernetes cluster on EC2, and automated HTTPS delivery through Caddy and Let's Encrypt at <https://clearroots.omerdengiz.com>. The public endpoint is therefore not a replacement for the internal Company 2 path. It is a parallel public-delivery implementation that demonstrates how the same service idea can be expressed beyond the Site 2 edge while still supporting ClearRoots' preference for scalable, low-vendor-lock-in service delivery.

This component extends the environment into orchestration, containerization, and public cloud delivery. Terraform provides the infrastructure-as-code layer, AWS EC2 provides compute, AWS S3 carries the remote Terraform state, AWS Route 53 publishes the public DNS record, Docker packages the site, and Kubernetes orchestrates the workload [21]-[25].

3.7.5.2 Architectural Rationale

Terraform was chosen because the cloud portion of the project needed to be reproducible, reviewable, and safe to rebuild. A manual sequence of console actions would have been harder to audit and harder for another team member to repeat consistently. Terraform kept the AWS provider configuration, EC2 definitions, security group, Elastic IP association, Route 53 record, and outputs in version-controlled files. The remote state backend in the manually created S3 bucket `clearroots.omerdengiz`, keyed at `clearroots/terraform.tfstate` in `us-east-1`, gives the team a shared state source without placing state data in the repository. The bucket was created outside Terraform on purpose so that Terraform did not have to manage its own backend before it could initialize [21].

The compute layer uses two Ubuntu 22.04 LTS EC2 instances in us-east-1, named clearroots-kube-master and clearroots-kube-worker, both built from AMI ami-0c7217cdde317cfec and both sized as t3.micro. That decision separates the control plane from the workload-serving node and therefore reflects a real Kubernetes architecture rather than a single-node demonstration. It also keeps the deployment within a low-cost footprint that fits a student project more appropriately than a managed control plane or larger instances would.

The worker alone carries the Elastic IP because it is the traffic-serving node. The master does not need a stable public address for browser traffic, and publishing both nodes would add exposure without improving the delivery path. By fixing the public address on the worker, the Route 53 A record remains valid across worker reboots and routine lifecycle events. That keeps the public endpoint tied to the node that actually accepts HTTPS requests [25].

The IAM role ec2connectcli-clearroots with EC2 Instance Connect permissions was chosen to reduce credential sprawl during bootstrap. The worker needed to retrieve the Kubernetes join token and CA hash from the master during provisioning, but distributing a shared SSH key for that purpose would have created a weaker and less tidy operating model. EC2 Instance Connect was delivered through a dedicated aws_iam_role_policy resource named ec2connect-inline attached to the ec2connectcli-clearroots role, which allowed the worker bootstrap to use mssh to reach the master and obtain the join material without introducing another long-lived shared credential [21].

Caddy was selected as the public HTTPS reverse proxy on the worker because the project needed reliable TLS with minimal administrative overhead. A single Caddyfile entry for clearroots.omerdengiz.com was enough to request and renew a Let's Encrypt certificate automatically while proxying traffic to 127.0.0.1:30080. In this architecture Caddy does exactly what the public edge needs to do and does not require a separate certificate-management subsystem [23].

The Kubernetes Service is published as a NodePort on 30080 because that is the simplest way to expose the clearroots-web Deployment to Caddy without introducing a cloud load balancer. The service abstraction still matters. It keeps the reverse proxy pointed at a stable Kubernetes endpoint rather than at individual pods, while preserving the ability for Kubernetes to manage two replicas behind that service [22].

The application image uses Docker with nginx:alpine as the base image because the site content is static and does not require a larger application stack. The Dockerfile copies the entire site/ directory into /usr/share/nginx/html and exposes port 80, keeping the image minimal while including all site assets in a single build step. That separation is operationally useful because the site can be rebuilt or restyled without redesigning the EC2 or Kubernetes foundation [24].

Route 53 ties the public name clearroots.omerdengiz.com to the worker's Elastic IP through an A record with TTL 300 in hosted zone Z05041043TC6QHEHQGZUG. That decision matters because the public endpoint needs a stable, human-readable name rather than a raw AWS address. The hosted zone was delegated from the parent

omerdengiz.com account, so the public site remains under a controlled DNS authority rather than an improvised hostname [25].

The cloud deployment is fully independent of the internal Site 2 environment. It does not traverse the OPNsense edge, the SITE1_OVPN tunnel, the MSP segment, or any Company 1 or Company 2 LAN resource. That separation is deliberate because the public cloud path is intended to stand on its own rather than tunneling internal Site 2 services onto the internet.

3.7.5.3 Infrastructure Components and Deployment Flow

Before Terraform was initialized, the team first created the S3 backend bucket clearroots.omerdengiz manually in AWS us-east-1. That step happened before Terraform initialization because the state backend had to exist before Terraform could use it safely.

Once the bucket existed, the working directory was connected to that remote backend with `terraform init -backend-config=backend.hcl`. That initialization step bound the configuration to the S3 state object at `clearroots/terraform.tfstate` and established the shared state path for future apply operations [21].

`terraform apply` then provisioned the cloud foundation described in `main.tf`, `variables.tf`, `route53.tf`, `s3.tf`, and `outputs.tf`. It created the IAM role and instance profile `ec2connectcli-clearroots`, the security group `clearroots-k8s-sec-gr`, the master instance `clearroots-kube-master`, the worker instance `clearroots-kube-worker` with an explicit dependency on the master, the Elastic IP attached to the worker, and the Route 53 A record that pointed `clearroots.omerdengiz.com` at that Elastic IP. The Terraform outputs then exposed `master_public_ip`, `master_private_ip`, `worker_public_ip`, and `website_url`.

On first boot, `master.sh` ran as user-data on the master node. It set the hostname to `kube-master`, installed Docker and Kubernetes 1.29.0, configured containerd with `SystemdCgroup=true`, ran `kubeadm init` with `--pod-network-cidr 10.244.0.0/16`, installed Flannel CNI and `local-path-provisioner`, and wrote `deployment.yaml` and `service.yaml` into `/home/ubuntu/clearroots/` [22].

On first boot, `worker.sh` ran as user-data on the worker node. It set the hostname to `kube-worker`, installed Docker and Kubernetes 1.29.0, used `mssh` through EC2 Instance Connect to retrieve the join token and CA hash from the master, ran `kubeadm join`, installed Caddy, and wrote a Caddyfile that defined `clearroots.omerdengiz.com { encode gzip zstd; reverse_proxy 127.0.0.1:30080 }` [22], [23].

After the cluster nodes were ready, the operator connected to the master and ran `kubectl apply -f /home/ubuntu/clearroots/`. That command applied the `clearroots-web` Deployment and the `clearroots-web-service` Service from the files that `master.sh` had already written. The Deployment created two replicas in the default namespace using `docker.io/ofdengiz/clearroots-web:latest`, exposed containerPort 80, and applied the defined resource requests and limits. The Service published port 80 to targetPort 80 on nodePort 30080.

After DNS propagation completed, Caddy requested a Let's Encrypt certificate automatically on the first HTTPS request to `clearroots.omerdengiz.com`. No separate certificate-installation workflow was required because TLS issuance and renewal were handled directly by Caddy's standard reverse-proxy behaviour [23], [25].

Once those steps had completed, the request path was straightforward. A browser request for `https://clearroots.omerdengiz.com` was resolved by Route 53 to the worker Elastic IP. Caddy on the worker accepted the TLS connection on port 443 and proxied the decrypted request to the Kubernetes NodePort service on `127.0.0.1:30080`. The NodePort service then forwarded the request to one of the two `clearroots-web` pods running `nginx:alpine` on containerPort 80, and the pod returned the ClearRoots Foundation HTML page.

3.7.5.4 Observed Deployment State

Terraform apply completed without errors, provisioning 9 resources in total. The output block confirmed `master_public_ip = 100.53.70.90`, `master_private_ip = 172.31.12.199`, `worker_public_ip = 52.22.77.66` as the Elastic IP, and `website_url = https://clearroots.omerdengiz.com`. The individual resource creation log confirmed the IAM role `ec2connectcli-clearroots`, security group `sg-0d0edc108381de577`, master instance `i-02985d4465a3170fb`, worker instance `i-0cf76aa23aa18062a`, Elastic IP `eipalloc-03ed028af6b538b67`, and Route 53 record `Z05041043TC6QHEHQZUG_clearroots.omerdengiz.com_A` all created successfully.

The AWS EC2 console confirmed that both provisioned instances were in Running state with 3/3 status checks passed. `clearroots-kube-master` (`i-02985d4465a3170fb`) carried public IPv4 address `100.53.70.90` in `us-east-1a`. `clearroots-kube-worker` (`i-0cf76aa23aa18062a`) carried the Elastic IP `52.22.77.66` in `us-east-1a`, which is the address referenced by the Route 53 A record. That console view confirmed that the Terraform resource definitions had translated into live, healthy AWS compute resources.

The S3 backend bucket `clearroots.omerdengiz` contained the object `clearroots/terraform.tfstate` after apply completed, confirmed from both the AWS S3 console and a direct `aws s3 ls` command run from `C2WebServer`. The console showed the `tfstate` object at 22.7 KB, last modified March 29, 2026 at 16:36:57 UTC. The terminal command confirmed the same object at 23,224 bytes timestamped 2026-03-29 20:36:57. Both paths confirmed that remote state had been written successfully and that future Terraform operations would have a consistent shared state source.

`kubectl apply -f /home/ubuntu/clearroots/` created both the Deployment and the Service in the default namespace. `kubectl get pods` showed two `clearroots-web` pods - `clearroots-web-9858b77ff-j6rr2` and `clearroots-web-9858b77ff-xfr9l` - both in Running state within 11 seconds of apply. `kubectl get svc` showed `clearroots-web-service` as a NodePort service with cluster IP `10.100.49.211` and port mapping `80:30080/TCP`. `kubectl get nodes` confirmed that `kube-master` carried the control-plane role and `kube-worker` was Ready, both running Kubernetes v1.29.0. That cluster state was the most direct confirmation that the worker-side reverse proxy had a live Kubernetes service target behind it.

A direct header check from MSPUbuntuJump confirmed the full delivery chain. `curl -I https://clearroots.omerdengiz.com` returned HTTP/2 200 with server: nginx/1.29.7 and via: 1.1 Caddy, confirming that Caddy terminated TLS and proxied the request to the nginx workload inside Kubernetes. The response was timestamped March 29, 2026 at 20:55:32 GMT with content-length 7326. Browser access to the same endpoint rendered the complete ClearRoots Foundation landing page over public HTTPS, confirming that public DNS, TLS termination via Caddy and Let's Encrypt, Kubernetes NodePort service forwarding on port 30080, and the nginx pod workload were all functioning end to end.

The Route 53 console confirmed the A record for `clearroots.omerdengiz.com` resolving to `52.22.77.66` with TTL 300 in hosted zone `Z05041043TC6QHEHQZUG`. The NS and SOA records were also present, confirming that the hosted zone was active and authoritative for the `clearroots.omerdengiz.com` subdomain. The A record value `52.22.77.66` matches the worker Elastic IP confirmed in both the Terraform outputs and the EC2 console.

3.7.5.5 Dependency Notes and Operational Considerations

The public cloud stack is independent of the internal Site 2 environment. It does not require the OPNsense edge, the `SITE1_OVPN` tunnel, the MSP segment, or any tenant LAN resource to serve `clearroots.omerdengiz.com`. That separation matters operationally because a Site 2 outage does not automatically take down the public AWS endpoint, and a cloud-side failure does not automatically imply an internal tenant problem.

Route 53 delegation must remain intact for the public endpoint to resolve. If the NS delegation in the parent `omerdengiz.com` account is removed, changed, or pointed elsewhere, `clearroots.omerdengiz.com` will fail publicly even if both EC2 instances and the Kubernetes workload remain healthy. Hosted zone `Z05041043TC6QHEHQZUG` must remain the authoritative zone for that subdomain path [25].

The Elastic IP is the DNS anchor for the worker node. If the worker is terminated and recreated without re-associating the same Elastic IP, the Route 53 A record will point to the wrong target until the association is restored. In normal operation Terraform manages that relationship, so a fresh terraform apply after reprovisioning is the intended way to restore the correct address binding [21], [25]. The confirmed worker Elastic IP at the time of final validation was `52.22.77.66`, consistent across the Terraform outputs, EC2 console, and Route 53 A record.

The S3 backend bucket `clearroots.omerdengiz` remains a manual prerequisite rather than a Terraform-managed resource. It must not be deleted between deployments. If it is lost, the Terraform state path disappears with it, and the live environment must be reconciled and rebuilt into state before additional apply operations are safe. That makes the bucket a small object with a large operational consequence [21].

The security group `clearroots-k8s-sec-gr` must continue to allow inbound TCP 80 and 443 from `0.0.0.0/0` so that public web traffic and Let's Encrypt HTTP-01 validation can succeed. If those rules are removed, certificate issuance or renewal will fail and public site delivery will fail with it. TCP 6443 must also remain open for Kubernetes control-

plane communication during any future cluster rebuild or kubeadm join workflow [22], [23].

4. Conclusion

Read together, the retained configuration and evidence blocks show one coherent managed environment rather than two unrelated site narratives. Company 1 and Company 2 remain distinct at the identity, naming, file-service, and client-delivery layers, but they are still supported through a shared architecture of bounded MSP entry, shared compute, isolated storage transport, and integrated recovery handling.

What was learned across the merged design is consistent with the strongest points in both source handovers: stable multi-tenant operation depends less on any single service build than on preserving segmentation, management-path discipline, and the relationships between identity, storage, client access, and backup. The practical challenge was keeping unlike service stacks interoperable without allowing DNS, routing, or storage policy to bleed across tenant boundaries.

Those lessons reinforce the final operating position of the environment. The estate is supportable, well-structured, and ready for operational handover within the documented boundaries, dependencies, and recovery model already described in the report.

5. Appendices

Appendix A. Addressing, Inventory, and Endpoint Reference

This appendix gathers the addressing, gateway, inventory, and endpoint tables most likely to be used during onboarding and early support work.

Table 45. Management endpoints

System	IP / URL	Access Method	Operational Use
Proxmox VE	https://192.168.64.10:8006	Web console	Hypervisor management
Grafana	http://172.30.64.180:3000	Web console	Monitoring dashboard
Cockpit	https://172.30.64.146:9090	Web console	Linux administration for Company 2
Windows Admin Center	https://172.30.64.179:6600	Web console	Windows infrastructure management
Jumpbox Windows	172.30.64.179	RDP / local admin tooling	Primary administrative entry point
Jumpbox Ubuntu	172.30.64.180	SSH / browser-based access	Linux administration plus Grafana and InfluxDB monitoring services
Proxmox Host iLO	https://192.168.64.11	Web console	Out-of-band hardware management for Server1
Server2 iLO	https://192.168.64.21	Web console	Out-of-band management endpoint for Server2
Lab Switch	ssh admin@192.168.64.2	SSH	Physical switching, VLAN trunk control, and upstream lab-network handoff

Table 46. Full VLAN addressing matrix

VLAN	Network	Gateway	Purpose	Traffic Type
VLAN 10	172.30.64.0/26	172.30.64.1	Company 1 Client Network	Routed
VLAN 20	172.30.64.128/28	172.30.64.129	Company 1 Server Network	Routed
VLAN 30	172.30.64.160/29	172.30.64.161	Company 1	Routed

			Web / DMZ	
VLAN 40	172.30.64.184/29	None (direct SAN)	Company 1 SAN Network	Direct storage
VLAN 99	172.30.64.176/29	172.30.64.177	Management Network	Restricted admin
VLAN 110	172.30.64.64/26	172.30.64.65	Company 2 Client Network	Routed
VLAN 120	172.30.64.144/28	172.30.64.145	Company 2 Server Network	Routed
VLAN 130	172.30.64.168/29	172.30.64.169	Company 2 Web / DMZ	Routed
VLAN 140	172.30.64.192/29	None (direct SAN)	Company 2 SAN Network	Direct storage
Infrastructure / WAN	192.168.64.0/24	192.168.64.1	Shared infrastructure and WAN-side management segment for OPNsense WAN, Proxmox VE, Server2 management, and iLO access	Shared infrastructure

Table 47. Appendix A addressing reference

Scope	Host or Interface	Address or Port	Role or Note
MSP Edge	rp-msp-gateway WAN	172.20.64.1/16	Provider-facing WAN interface on the OPNsense gateway VM; 2 GiB RAM; domain et.lab; timezone America/Toronto.
MSP Edge	rp-msp-gateway MSP	172.30.65.177/29	MSP segment gateway and policy control point.
MSP Edge	rp-msp-gateway C1LAN	172.30.65.1/26	Gateway for Company 1 LAN systems.
MSP Edge	rp-msp-gateway C1DMZ	172.30.65.161/29	Gateway for Company 1 DMZ systems.
MSP Edge	rp-msp-gateway C2LAN	172.30.65.65/26	Gateway for Company 2 LAN systems.
MSP Edge	rp-msp-gateway	172.30.65.169/29	Gateway for

	C2DMZ		Company 2 DMZ systems.
MSP Edge	SITE1_OVPN	OpenVPN inter-site tunnel interface	Carries cross-site web and backup copy traffic to Site 1.
MSP Edge	WAN NAT publication	33464 -> 172.30.65.178:3389	Jump64 RDP publication.
MSP Edge	WAN NAT publication	33564 -> 172.30.65.179:22	MSPUbuntuJump SSH publication.
MSP Edge	Static route	192.168.64.20/32 via Site 1 OpenVPN gateway	Route to SITE1_VEEAM across the tunnel.
MSP Systems	Jump64	172.30.65.178	Windows Server 2022 bastion; 4 vCPU, 8 GiB RAM, dual-path access via internal IP and Tailscale.
MSP Systems	MSPUbuntuJump	172.30.65.179	Ubuntu 22.04.5 LTS bastion; 4 vCPU, 4 GiB RAM, ens18 172.30.65.179/29, primary Linux inspection path.
MSP Systems	S2Veeam	172.30.65.180	Windows Server 2022 backup and offsite-copy platform; 4 vCPU, 8 GiB RAM, repositories on R: and Z:.
Company 1	C1DC1	172.30.65.2	Windows Server 2022 primary domain controller; 4 vCPU, 4 GiB RAM, AD DS and DNS confirmed.
Company 1	C1DC2	172.30.65.3	Windows Server 2022 secondary domain controller; 4 vCPU, 4 GiB RAM, AD DS and DNS confirmed.
Company 1	C1FS	172.30.65.4	Windows Server 2022 file server; 4 vCPU, 6 GiB RAM, SMB shares on F:, active iSCSI

			consumer.
Company 1	C1WindowsClient	172.30.65.11	Windows 10 endpoint; 4 vCPU, 2 GiB RAM, c1.local membership confirmed.
Company 1	C1UbuntuClient	172.30.65.36	Ubuntu 25.04 endpoint; 4 vCPU, 2 GiB RAM, C1.LOCAL realm active.
Company 1	C1WebServer	172.30.65.162	Windows Server 2022 IIS host in C1DMZ; 4 vCPU, 2 GiB RAM, hostname-bound HTTPS service.
Company 1	C1SAN	172.30.65.186/29, gw 172.30.65.185	Windows Server 2022 storage target; 4 vCPU, 2 GiB RAM, 32 GB system disk + 160 GB data disk, SeaBIOS, VirtIO SCSI single, net0 on isolated bridge G264.
Company 2	C2IdM1	172.30.65.66	Primary Samba AD DC, DNS, DHCP; Ubuntu 22.04.5 LTS, 4 vCPU, 4 GiB RAM.
Company 2	C2IdM2	172.30.65.67	Secondary Samba AD DC, DNS, DHCP; Ubuntu 22.04.5 LTS, 4 vCPU, 4 GiB RAM.
Company 2	C2FS service NIC	ens19 172.30.65.68/26	Tenant-facing file service interface for C2FS; Ubuntu 22.04.5 LTS, 4 vCPU, 6 GiB RAM.
Company 2	C2FS storage NIC	ens18 172.30.65.195/29	Dedicated storage path to C2SAN.
Company 2	C2LinuxClient	172.30.65.70	Ubuntu 25.04 endpoint; 4 vCPU, 2 GiB RAM.
Company 2	C2WebServer	172.30.65.170	Internal HTTPS web server in

			C2DMZ; Ubuntu 22.04.5 LTS + nginx, 4 vCPU, 2 GiB RAM.
Company 2	C2SAN	172.30.65.194/29, gw 172.30.65.193	Ubuntu 22.04.5 LTS storage target; 4 vCPU, 2 GiB RAM, 32 GB system disk + 160 GB data disk, SeaBIOS, VirtIO SCSI single, net0 on isolated bridge M265.
DNS Records	c1- webserver.c1.local	172.30.64.162 and 172.30.65.162	Observed on C2IdM1 and C2IdM2.
DNS Records	c2- webserver.c2.local	172.30.64.170 and 172.30.65.170	Observed on C2IdM1 and C2IdM2.

Appendix B. DHCP, Storage, and Backup Reference

Table 48. Documented DHCP scope

Scope Name	Network	Range
VLAN10_Clients	172.30.64.0/26	172.30.64.2 – 172.30.64.62
VLAN110_Clients	172.30.64.64/26	172.30.64.66 – 172.30.64.126

Table 49. Storage server interfaces

Interface	IP Address	Purpose
LAN1-WAN	192.168.64.20	Management and backup network
SAN_C1	172.30.64.186	Company 1 storage network
SAN_C2	172.30.64.194	Company 2 storage network

Table 50. iSCSI target mappings

Target Name	Initiator IP	Storage Location
c1-dc1	172.30.64.187	S:\iSCSIVirtualDisks
c1-dc2	172.30.64.188	S:\iSCSIVirtualDisks
c2-dc1	172.30.64.195	T:\iSCSIVirtualDisks
c2-dc2	172.30.64.196	T:\iSCSIVirtualDisks
target-win10	172.30.64.189	S:\iSCSIVirtualDisks
target-ubuntu	172.30.64.190	S:\iSCSIVirtualDisks

Table 51. Backup repository volumes

Drive	File System	Purpose
V:	ReFS	Primary backup storage on Server2
Site2 R:	NTFS via SMB shared folder	Remote offsite backup copy repository on 172.30.65.180 (R:\Site1OffsiteFromServer2\Repo)
Site2 Local Repository	NTFS	Reserved for Site2 local backup operations on 172.30.65.180 and intentionally separated from the Site1 offsite copy path

Table 52. Veeam backup file types

File Type	Extension	Description
Full backup file	.vbk	Primary full-backup file used as the anchor of a

		Veeam backup chain.
Incremental backup file	.vib	Incremental restore-point file that records changes after the full backup.
Backup metadata file	.vbm	Metadata and chain-index file used by Veeam to track restore points and job state.

Appendix C. Operational Handover, Verification, and Triage

This appendix gathers the most reusable handover, verification, and early-triage material so support staff can work from one quick-reference location during onboarding and first-response troubleshooting.

This appendix reorganizes the most important service commitments, administrative paths, and validation outcomes into a quick operational reference for handover and early support use.

Network and Shared Infrastructure

Physical and out-of-band management: Figures 29 and 3B document the rack installation and cable layout. Server1 iLO is identified at 192.168.64.11, and the Server2 iLO endpoint is identified at 192.168.64.21 for later hardware administration.

Type 1 hypervisor: Proxmox VE runs on Server1 as the Site 1 compute platform and hosts the tenant servers, clients, jump systems, and shared services described in the virtual machine inventory.

Storage server and SAN presentation: Server2 provides tenant-separated SAN presentation. Company 1 uses SAN_C1 172.30.64.186 with initiators 172.30.64.187 through 172.30.64.190, while Company 2 uses SAN_C2 172.30.64.194 with initiators 172.30.64.195 and 172.30.64.196.

Company 1

Identity, DNS, and DHCP: c1.local is hosted on C1-DC1 and C1-DC2 with Windows Server AD DS. Company 1 also provides recursive DNS through the tenant domain controllers and uses Windows DHCP failover for the client scope.

File services and namespace: \\c1.local\namespace publishes the public and private logical paths backed by Pub1, Priv1, Pub2, and Priv2 across the Company 1 domain controllers and the secondary-site file-service target.

Access model and client workflow: Public remains the shared collaboration area for internal users, while Private is structured as per-user folders. C1-Client1 uses the published H: and P: mappings for employee1 and employee2 workflow validation, and C1-Client2-Linux keeps admin as a local Linux administration account while employee1@c1.local and employee2@c1.local use domain-backed logins that automatically receive ~/C1_Public and the correct per-user ~/C1_Private.

Replication and permission state: Pub1, Priv1, Pub2, and Priv2 were aligned to the same Access-Based Enumeration and share-permission model, and DFS Replication resumed after the read-only state on the 766 GB SAN disk attached to C1-DC2 was corrected. At the firewall layer, C1_ClientsVLAN10 allows same-tenant access to Alias_C1_GLOBAL, blocks Alias_C2_GLOBAL, permits TCP 80-443 to C1_DMZVLAN30, and allows non-RFC1918 internet access. A host-specific rule also allows C1-Client1 (172.30.64.2) to reach Server2 (192.168.64.20) on VEEAM_CLIENT1_BACKUP_PORTS. Company 1 SAN/iSCSI traffic uses the separate VLAN 40 direct-storage path between the initiator SAN NICs and Server2, so that

storage access remains off the routed client firewall path. C1_ServersVLAN20 allows Alias_C1_GLOBAL, school DNS, and non-RFC1918 internet access while blocking Alias_C2_GLOBAL.

Company 2

Identity, DNS, and DHCP: Company 2 uses Samba AD DC on C2-DC1 and C2-DC2. The c2.local and _msdcs.c2.local zones are AD-integrated primary zones with secure updates, recursive lookup was verified for external names, and ISC DHCP runs as a failover pair with C2-DC1 as primary and C2-DC2 as secondary for the 172.30.64.64/26 client scope.

File services: C2_Public and C2_Private are published from the replicated GlusterFS volume gv0 mounted at /mnt/sync_disk. Public maps to /mnt/sync_disk/Public for shared collaboration, while Private maps to /mnt/sync_disk/Private/%U for per-user storage limited to members of c2_file_users.

Fault tolerance and client workflow: GlusterFS gv0 runs as a two-brick replicate volume across 172.30.64.146 and 172.30.64.147, and C2-Client1 keeps admin as a local Linux administration account while employee1@c2.local and employee2@c2.local complete successful Company 2 domain-user shell and home-directory sessions, then receive ~/C2_Public plus ~/C2_Private over SMB 3.1.1. Each user remains limited to the matching private content.

iSCSI and SAN isolation: open-iscsi sessions from 172.30.64.195 and 172.30.64.196 connect to Server2 SAN_C2 at 172.30.64.194, while the tenant server addresses remain on vlan120 at 172.30.64.146 and 172.30.64.147. This keeps storage traffic separate from the routed user and server path. At the firewall layer, C2_ClientsVLAN110 allows Alias_C2_GLOBAL, blocks Alias_C1_GLOBAL, permits TCP 80-443 to C2_DMZVLAN130, and allows internet access only to non-RFC1918 destinations. C2_ServersVLAN120 allows Alias_C2_GLOBAL, school DNS, and non-RFC1918 internet access while blocking Alias_C1_GLOBAL. The C2_DMZVLAN130 interface is permitted to use SCHOOL_DNS, reach non-RFC1918 internet destinations, and accept the published outside access shown for the c2-webserver.

Cross-Site and Remote Access

Site-to-site connectivity: Site 1 and Site 2 are linked by the documented OpenVPN tunnel. In the current Site 1 firewall state, the OpenVPN interface now includes two explicit block rules that deny Alias_C1_GLOBAL to Alias_C2_GLOBAL and Alias_C2_GLOBAL to Alias_C1_GLOBAL before a remaining broad "Allow OpenVPN traffic for testing" IPv4 pass rule. The inter-site backup path in active use is reinforced on the WAN interface by host-specific ICMP and TCP 445 rules from Server2 (192.168.64.20) to the Site2-Offsite-SharedRepo host at 172.30.65.180.

Administrative access: Jumpbox Windows and Jumpbox Ubuntu remain the approved management entry points. The Ubuntu jumpbox also provides the Grafana dashboard at <http://172.30.64.180:3000> and hosts the local InfluxDB service used for Proxmox metrics

collection. RDP is enabled on the Windows systems documented in the management section, and SSH is active on C1-Client2, C2-DC1, C2-DC2, and C2-Client1.

Backup and offsite copy: Server2 hosts Veeam Backup and Replication for local VM backup, the C1-Client1 file-based backup workflow, and the Site1-to-Site2 backup copy path to the Site2-Offsite-SharedRepo SMB repository on 172.30.65.180 at \\172.30.65.180\Site1OffsiteFromServer2\$\Repo on the R: volume. In the active Site 1 firewall view, this path is represented on WAN by ICMP and TCP 445 rules from 192.168.64.20 to 172.30.65.180.

Operational Handover Use

Operational use: this appendix is intended as a quick support-reference matrix. The body of the report explains the design and rationale in detail, while this section condenses the most important service obligations, access paths, and validation outcomes into a form suitable for onboarding, change review, or first-line triage.

Table 53. Routine operational checks

Check	How To Verify	Expected Result	Why It Matters
Jump host reachability	Confirm WAN NAT 33464 and 33564 reach Jump64 and MSPUbuntuJump	At least one bastion is reachable and both are preferred	Every support action begins on the bastion layer.
OPNsense tunnel state	Inspect SITE1_OVPN state and the route to 192.168.64.20/32	Inter-site path is up and the Site 1 backup endpoint is routable	Cross-site web access and offsite copy depend on the tunnel.
Company 1 directory health	From Jump64, rerun Get-ADDomainController or the approved WinRM-based domain check against C1DC1 and C1DC2	Both controllers return domain and service state successfully	Company 1 logons, DNS, and policy checks depend on this foundation.
Company 1 file service path	From Jump64, verify Get-SmbShare, Get-IscsiSession, and the F: SharedData volume on C1FS	SMB shares are present and the iSCSI session is active	This proves that Company 1 file publication and its hidden storage path are both healthy.
Company 2 identity services	systemctl is-active samba-ad-dc and systemctl is-active isc-dhcp-server on C2IdM1 and C2IdM2	Both services active on both nodes; PRIMARY on C2IdM1 and SECONDARY on C2IdM2	Company 2 authentication, DNS, and DHCP are concentrated here.
Company 2 file service mount	findmnt /mnt/c2_public and iscsiadm -m session on C2FS	/dev/sdb mounted and the session to 172.30.65.194:3260,1 active	If the block layer is gone, the shares will fail next.

Synchronization health	Review /var/log/c2_site1_sync.log	Most recent Site1->Site2 C2 sync completed successfully	Sync issues do not break backup, but they do affect replicated data freshness.
Internal web delivery	curl -kl https://c1-webserver.c1.local and curl -kl https://c2-webserver.c2.local	Hostname requests return HTTP 200, while raw IP checks still return HTTP 404	The site intentionally serves by hostname rather than by address.
Backup status	Review Veeam job history and service state on S2Veeam	Four job families and copy jobs complete without new failures	Protection is only useful if it continues to run after deployment day.

Table 54. Fast triage guide

Symptom	First System To Check	Quick Command or Check	Healthy Expectation	Likely Fault Domain If Unhealthy
Cannot reach Site 2 remotely at all	rp-msp-gateway, then Jump64 and MSPUbuntu Jump	Verify WAN publication 33464 -> 172.30.65.178:3389 and 33564 -> 172.30.65.179:22; confirm bastion reachability	At least one bastion responds and both NAT rules are present	WAN edge, NAT publication, provider path, or bastion outage
c1-webserver.c1.local or c2-webserver.c2.local does not resolve	Tenant DNS authorities: C1DC1 or C1DC2 for Company 1 names; C2IdM1 or C2IdM2 for Company 2 names	Run nslookup or dig from the affected client; on Company 2 identity nodes confirm the c1.local and c2.local zones and both A records	DNS returns 172.30.64.162 and 172.30.65.162 for c1-webserver.c1.local, and 172.30.64.170 and 172.30.65.170 for c2-webserver.c2.local	Tenant DNS outage, missing resolver search scope, or bad zone data
Web hostname returns 404 instead of 200	C1WebServer or C2WebServer	Check the hostname-based binding: IIS on c1-webserver.c1.local TCP 443, or nginx server_name c2-webserver.c2.local on port 443	Hostname request returns 200 while raw IP still returns 404	Binding drift, wrong host header, or DNS pointing at the wrong

				server
Company 1 shares appear unavailable	C1FS (via Jump64)	From Jump64 run Get-SmbShare, Get-IscsiSession, and confirm the F: SharedData volume	Named SMB shares remain present, the iSCSI initiator session is active, and F: stays mounted	Windows file service fault, storage-session loss, or hidden SAN path issue
Company 2 shares appear unavailable	C2FS	Run systemctl is-active smbd, findmnt /mnt/c2_public, and iscsiadm -m session	smbd is active, /dev/sdb is mounted, and the 172.30.65.194:3260,1 session is present	Samba service fault, mount loss, or broken iSCSI path
Company 1 user cannot authenticate	C1DC1, then C1DC2	From Jump64 rerun Get-ADDomainController or the approved WinRM directory check; from the client confirm c1.local membership and DNS	Both controllers return healthy directory state and the client still points at Company 1 DNS	Active Directory, DNS, or client trust-path fault
Company 2 user cannot authenticate	C2IdM1, then C2IdM2	Run systemctl is-active samba-ad-dc; from C2LinuxClient confirm realm list and getent passwd employee1@c2.local	Samba AD is active and domain users still resolve	Directory, DNS, or client realm integration fault
Backup job shows as failed	S2Veeam	Check VeeamBackupSvc and companion services, confirm Site2Veeam on Z:\Site2AgentBackups, and verify access to \\192.168.64.20\Site2OffsiteFromSite2	Core Veeam services are running, repository is present, and the offsite target is reachable	Backup service failure, repository issue, or inter-site copy path problem
Domain user login succeeds but C2_Public and C2_Private are not mounted on C2LinuxClient	C2LinuxClient, then C2IdM1	On C2LinuxClient: sudo journalctl -b grep -Ei 'pam_mount cifs' to identify mount failure; kvno cifs/c2fs.c2.local to test Kerberos ticket acquisition. On C2IdM1: sudo samba-tool spn list C2FS\$ to verify cifs/c2fs and cifs/c2fs.c2.local SPNs are present.	kvno returns kvno = 1 for cifs/c2fs.c2.local; samba-tool spn list shows both cifs/c2fs and cifs/c2fs.c2.local; after login mount grep c2fs shows	Missing CIFS SPN on C2FS\$; KDC unreachable from C2LinuxClient due to DNS misconfiguration;

			both shares mounted via sec=krb5.	pam_mount volume options incompatible with the kernel CIFS/Kerberos path.
--	--	--	-----------------------------------	---

Appendix D. C2FS Samba Configuration (sanitized excerpt)

The following sanitized excerpt is included verbatim because it captures the logic of the Company 2 file-service publication path in one place.

[global]

```
workgroup = C2
realm = C2.LOCAL
security = ADS
server role = member server
kerberos method = secrets and keytab
winbind use default domain = yes
winbind refresh tickets = yes
idmap config * : backend = tdb
idmap config * : range = 3000-7999
template shell = /bin/bash
template homedir = /home/%U
obey pam restrictions = no
log file = /var/log/samba/log.%m
max log size = 1000
logging = file
```

[C2_Public]

```
path = /mnt/c2_public/Public
browseable = yes
read only = no
valid users = @c2_file_users
force group = c2_file_users
create mask = 0770
directory mask = 0770
```

[C2_Private]

```
path = /mnt/c2_public/Private/%U
browseable = no
read only = no
valid users = %U
force group = c2_file_users
create mask = 0700
directory mask = 0700
```

Appendix E. Resolved Gaps and Troubleshooting History

OPNsense authenticated GUI walkthrough. Resolved. The OPNsense management surface was confirmed reachable through the MSP Linux bastion. HTTP 403 on port 80 was expected behaviour and confirmed that the management interface was present but not anonymously accessible. TCP 53 on rp-msp-gateway was also confirmed reachable from MSPUbuntuJump. The alias inventory, NAT publication rules, and OpenVPN firewall policy were then confirmed through GUI screenshot evidence that was captured and included in the report. No documentation gap remained.

Veeam GUI screenshots. Resolved. Refreshed Veeam GUI screenshots were captured during the final submission pass and were incorporated into Section 3.7 and Figure 41. The evidence set now reflected the repository inventory, including Site2Veeam on Z:\Site2AgentBackups, the offsite SMB target Site1OffsiteSmbShare at \\192.168.64.20\Site2OffsiteFromSite2, all four job families Ubuntu_Servers, Windows_Servers, C1_FileShare, and C2_FileShare, and the active Veeam service state. The earlier presentation gap therefore no longer existed.

C1FS TCP 5985 not reachable from MSPUbuntuJump. Resolved. TCP 5985 on C1FS (172.30.65.4) was confirmed reachable from the MSP management path during the final pass. WinRM-based inspection from MSPUbuntuJump was therefore available in addition to the Jump64 path documented in earlier validation sessions. The earlier evidence that had been collected through WMI and remote process execution from Jump64 remained technically valid and was retained in Section 3.4.3 as part of the system's validation history.

C2LinuxClient c1.local resolution failure. Resolved. During the March 25, 2026 validation pass, c1-webserver.c1.local returned REFUSED from C2LinuxClient because c1.local was absent from the resolver search scope. The resolver configuration was corrected in a later pass by adding both c1.local and c2.local to the search domain list. The final validated state confirmed that both hostnames resolved correctly and returned HTTP 200 from the client. The original failed observation was retained in Section 3.5.3 as troubleshooting history rather than being removed.

Appendix F. Supplemental Validation and Escalation Notes

This appendix provides supplemental configuration-level validation notes that may be used during planned change review, incident triage, or formal handover walkthroughs when a service area needs deeper inspection than the baseline operational summary.

Recommended validation pattern: open the relevant management view first, then run the matching command or user workflow, and finally confirm the expected operator-facing or service-facing result. This keeps configuration evidence and live behavior tied together during troubleshooting.

Manual emphasis for Service Block 4: for user-permission validation, the most authoritative evidence is to show one user reaching Public plus their own Private path and then failing to read another user's Private path.

Manual emphasis for DHCP failover: show both dhcpd.conf role definitions together with the active service state so the pair is clearly presented as one coordinated primary/secondary design.

Appendix G. Requirements Traceability Matrix

This appendix maps the project-brief requirements tracked in this integrated report to the sections, tables, figures, and operational notes that show how they were implemented, constrained by lab hardware, or documented for next-phase expansion.

Table 55. Project requirement to integrated-report evidence map

Project requirement	Primary evidence in this report	Implementation status and note
Separate company networks and a separate management network	Sections 3.1, 3.2, 3.3, 3.6; Appendix A	Tenant LANs, DMZs, MSP access paths, and storage network paths are documented as separate operating zones.
Storage server platform with dedicated workload disks and offsite backup storage	Sections 3.4, 3.5; Appendix B	Using the fixed eight-disk lab unit (2 OS + 6 RAID 10 storage), the report preserves the functional storage layout, repository volumes, iSCSI mappings, and offsite copy paths.
Jump servers and approved administrative access	Sections 3.1 and 3.6; Appendix A; Appendix C	Administrative access remains anchored to approved jump systems and MSP-facing management channels.
VM, file-share, and client backup coverage	Section 3.5; Appendix B	Agentless VM backup, client backup, file-share protection, and offsite copy handling are all documented in one backup section.
Template-driven provisioning and automated deployment readiness	Sections 3.1 and 3.7	Template virtual machines and reproducible deployment workflows are documented for the private-cloud and public-cloud portions of the project.
Users must not receive direct access to administrative tools	Sections 3.1, 3.6, and 3.7	Management tools are presented as MSP or approved administrative surfaces rather than general user-facing services.

Project requirement	Primary evidence in this report	Implementation status and note
Value-added services implemented beyond the core baseline	Section 3.7	The merged report keeps the monitoring, management, and cloud-delivery extensions as a distinct main-body section with clear operational relevance.
High-availability and migration preparedness	Sections 3.1 and 3.6	Current cold-migration and recovery readiness are documented through backups, templates, and rebuild workflows, while second-node Proxmox clustering is identified as the next-phase path to true live migration.
Tenant user accounts and separation between user and administrative workflows	Sections 3.2, 3.3, 3.4; Appendix C	Company user accounts, per-user private access, and local-admin-versus-domain-user workflows are documented for both tenants through validated client sessions and file-share controls.
Remote access enabled through approved protocols and bastion paths	Sections 3.1, 3.2, 3.3, 3.6; Appendix A; Appendix C	RDP, SSH, and browser-based administrative access are documented through approved management paths rather than direct unmanaged exposure.
Business alignment for Company 1 compliance and Company 2 open-source preference	Sections 2.1, 3.2, 3.3, 3.4, and 3.5	Platform and storage choices are tied to Company 1's controlled Canadian data-handling expectations and Company 2's open-source, low-vendor-lock-in operating model.

6. References

- [1] Proxmox Server Solutions GmbH, "Proxmox Virtual Environment Documentation," [Online]. Available: <https://pve.proxmox.com/pve-docs/>. [Accessed: Mar. 8, 2026].
- [2] OPNsense Project, "OPNsense Documentation," [Online]. Available: <https://docs.opnsense.org/>. [Accessed: Mar. 8, 2026].
- [3] Microsoft, "Windows Server Documentation," [Online]. Available: <https://learn.microsoft.com/windows-server/>. [Accessed: Mar. 8, 2026].
- [4] Samba Team, "Setting up Samba as an Active Directory Domain Controller," [Online]. Available: https://wiki.samba.org/index.php/Setting_up_Samba_as_an_Active_Directory_Domain_Controller. [Accessed: Mar. 8, 2026].
- [5] Veeam Software, "Veeam Backup and Replication Documentation," [Online]. Available: <https://helpcenter.veeam.com/>. [Accessed: Mar. 8, 2026].
- [6] Microsoft, "Windows Admin Center overview," [Online]. Available: <https://learn.microsoft.com/en-us/windows-server/manage/windows-admin-center/overview>. [Accessed: Mar. 26, 2026].
- [7] Microsoft, "Distributed File System Namespaces overview," [Online]. Available: <https://learn.microsoft.com/en-us/windows-server/storage/dfs-namespaces/dfs-overview>. [Accessed: Mar. 26, 2026].
- [8] Microsoft, "iSCSI Target Server overview," [Online]. Available: <https://learn.microsoft.com/en-us/windows-server/storage/iscsi/iscsi-target-server>. [Accessed: Mar. 26, 2026].
- [9] Grafana Labs, "Get started with Grafana Open Source," [Online]. Available: <https://grafana.com/docs/grafana/latest/fundamentals/getting-started/>. [Accessed: Mar. 26, 2026].
- [10] Cockpit Project, "Cockpit documentation," [Online]. Available: <https://cockpit-project.org/documentation.html>. [Accessed: Mar. 26, 2026].
- [11] OPNsense Documentation, "Firewall Rules," [Online]. Available: <https://docs.opnsense.org/manual/firewall.html>. [Accessed: Mar. 24, 2026].
- [12] OPNsense Documentation, "Network Address Translation," [Online]. Available: <https://docs.opnsense.org/manual/nat.html>. [Accessed: Mar. 24, 2026].
- [13] OPNsense Documentation, "Setup SSL VPN Road Warrior," [Online]. Available: https://docs.opnsense.org/manual/how-tos/sslvpn_client.html. [Accessed: Mar. 24, 2026].
- [14] Ubuntu Server Documentation, "Set up Samba as a file server," [Online]. Available: <https://documentation.ubuntu.com/server/how-to/samba/file-server/>. [Accessed: Mar. 24, 2026].

- [15] Ubuntu Server Documentation, "iSCSI initiator (or client)," [Online]. Available: <https://documentation.ubuntu.com/server/how-to/storage/iscsi-initiator-or-client/>. [Accessed: Mar. 24, 2026].
- [16] Microsoft Learn, "IIS binding Element for bindings for site," [Online]. Available: [https://learn.microsoft.com/en-us/previous-versions/iis/settings-schema/ms691267\(v=vs.90\)](https://learn.microsoft.com/en-us/previous-versions/iis/settings-schema/ms691267(v=vs.90)). [Accessed: Mar. 24, 2026].
- [17] Veeam Help Center, "Configuring Backup Repositories," [Online]. Available: https://helpcenter.veeam.com/docs/vbr/userguide/sch_configure_repository.html. [Accessed: Mar. 24, 2026].
- [18] nginx Documentation, "Server names," [Online]. Available: https://nginx.org/en/docs/http/server_names.html. [Accessed: Mar. 27, 2026].
- [19] Microsoft Learn, "About WMI," [Online]. Available: <https://learn.microsoft.com/en-us/windows/win32/wmisdk/about-wmi>. [Accessed: Mar. 27, 2026].
- [20] Microsoft Learn, "SMB features in Windows and Windows Server," [Online]. Available: <https://learn.microsoft.com/en-us/windows-server/storage/file-server/smb-feature-descriptions>. [Accessed: Mar. 27, 2026].
- [21] HashiCorp, "Terraform AWS Provider Documentation," [Online]. Available: <https://registry.terraform.io/providers/hashicorp/aws/latest/docs>. [Accessed: Mar. 29, 2026].
- [22] Kubernetes Documentation, "kubeadm init," [Online]. Available: <https://kubernetes.io/docs/reference/setup-tools/kubeadm/kubeadm-init/>. [Accessed: Mar. 29, 2026].
- [23] Caddy Documentation, "Reverse Proxy," [Online]. Available: https://caddyserver.com/docs/caddyfile/directives/reverse_proxy. [Accessed: Mar. 29, 2026].
- [24] Docker Documentation, "Dockerfile reference," [Online]. Available: <https://docs.docker.com/reference/dockerfile/>. [Accessed: Mar. 29, 2026].
- [25] AWS Documentation, "Amazon Route 53 Developer Guide," [Online]. Available: <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/>. [Accessed: Mar. 29, 2026].