End-to-End Temporary Hyper-V Cluster Build and DEV/UAT Migration

Professional portfolio project summary

Amandio Pereira | UK Technology Services Organisation | 2026

Prepared as a concise technical narrative of the hardware recovery, build, configuration, troubleshooting, validation and migration work completed to create an interim Hyper-V platform during a datacentre transition.

Project objective	Build a temporary three-node Hyper-V cluster from repurposed legacy hardware, validate it for development and user acceptance testing (UAT) workloads, and migrate services safely while modern hosting equipment was being transferred to a datacentre and permanent replacement hardware was still months away.
My role	Hands-on infrastructure engineer: recovered and rebuilt legacy hosts, configured Windows Server 2022 Hyper-V and Failover Clustering, physical switch ports/VLANs, host networking and shared SAN storage; resolved compatibility issues, tested resilience, and prepared and executed Hyper-V Replica migration.
Business value	Maintained a usable local development/UAT platform through an infrastructure transition, using available assets rather than waiting for new equipment while preserving shared storage, failover and rollback capability.
Core technologies	HPE ProLiant DL380 G7, Windows Server 2022 Datacenter, Hyper-V, Failover Clustering, Ethernet switching/VLANs, Switch Embedded Teaming, MPIO, Cluster Shared Volumes, Hyper-V Replica and PowerShell.

1. Executive Summary

The organisation needed an interim platform for development and user acceptance testing (UAT) workloads while the existing modern host hardware was being moved to a datacentre and suitable replacement equipment would not arrive for several months. The practical option was to recover three older HPE ProLiant DL380 G7 servers, reuse serviceable components from retired hosts where necessary, and turn them into a stable temporary cluster.

The work therefore involved more than installing Hyper-V. The hosts had to run Windows Server 2022 on much older hardware, integrate with shared SAN storage, and be matched to switch-side VLAN and port configuration separating management, live-migration and VM traffic. The target build used three nodes, a disk witness, two Cluster Shared Volumes and a two-NIC Switch Embedded Teaming virtual switch for VM traffic.

A significant compatibility issue emerged during cluster validation: ordinary connectivity tests succeeded, but host-to-host WinRM, CIM, WMI and Hyper-V management calls did not. After excluding network, firewall, policy and authentication causes, the fault was narrowed to legacy management-NIC behaviour. Adjusting offload and virtualisation-related settings restored remote management and allowed validation to proceed.

Once stable, the cluster was exercised with Windows 11 and Ubuntu Generation 2 test VMs through live migration, node drain and controlled host reboot/recovery. The migration stage then used Hyper-V Replica with documented inventory, network mapping, checkpoint cleanup, readiness checks and rollback controls to move DEV/UAT workloads onto the temporary platform.

2. Starting Point and Engineering Constraints

The project was driven by a temporary capacity gap rather than a conventional refresh. The modern hosts in service had to be transferred to a datacentre before the planned replacement equipment was available, so the requirement was to create a dependable bridge platform from hardware already on site.

Constraint	Engineering response	Why it mattered
Temporary capacity gap	Built an interim three-node platform from available retired hardware.	Kept DEV/UAT workloads available during the datacentre transition rather than waiting months for replacement hosts.
Legacy host condition	Recovered three serviceable DL380 G7 nodes by reusing working components from previous hosts where required.	Made enough compatible hardware available to retain a three-node clustered design instead of falling back to a single standalone host.
Large generation gap	Installed Windows Server 2022 Datacenter and validated roles, drivers, firmware behaviour and management functions methodically.	The operating system and clustering stack were substantially newer than the underlying platform, so compatibility could not be assumed.
Network and storage continuity	Separated management, live migration and VM traffic; connected shared SAN storage through MPIO and used a dedicated witness plus CSVs.	Preserved the operational characteristics needed for clustered virtualisation rather than treating the interim platform as disposable.
Migration risk	Inventoried the old cluster, mapped VLANs, removed stale state and checkpoints, and retained the source VMs as a rollback position during cutover.	Reduced the chance of carrying hidden faults or ambiguous configuration into the replacement environment.

3. Temporary Target Architecture

The design used simple, separated traffic roles across Hyper-V and the physical switches. I mapped host NICs to switch ports by MAC address and configured the required VLANs, aiming for the most supportable temporary cluster the available equipment could reliably sustain.

Layer	Implemented design	Operational purpose
Compute	Three repurposed HPE ProLiant DL380 G7 hosts.	Three-node failover capacity using serviceable hardware already available on site.
Platform	Windows Server 2022 Datacenter with Hyper-V, Failover Clustering and management tools.	A consistent modern Windows virtualisation and cluster-management layer across all hosts.
Shared storage	SAN-presented disks through MPIO, a dedicated disk witness and two Cluster Shared Volumes.	Shared VM storage, path resilience and cluster quorum.
Management	1Gb management NIC on the management VLAN; DNS registration enabled.	Stable host administration and cluster/client communication.
Live migration	1Gb live-migration NIC on a dedicated host-to-host network.	Kept migration traffic away from management and VM data paths.
VM data	2 x 1Gb NICs in SET (Switch Independent / Dynamic); switch ports carried VM VLANs; no LACP.	VM traffic capacity and link resilience aligned with the physical switches.
SAN path	10Gb DAC connectivity to shared storage with multipath configuration.	Higher-bandwidth storage access than the 1Gb host-management and VM network links.

Switching and Host Network Configuration

I configured both the physical switching and the Hyper-V host networking supporting the temporary cluster. This included dedicated host-facing ports for management, live migration and VM data, with the required VLAN membership and spanning-tree edge-port configuration. A separate switch provided host iLO and SAN-management connectivity. Host NICs were correlated to physical switch ports by MAC address before being renamed and assigned their operational roles.

Within Hyper-V, the two VM-data adapters on each host were combined using Switch Embedded Teaming in switch-independent mode with Dynamic load balancing. Management and live-migration traffic remained on dedicated adapters, while VM VLAN tagging was coordinated with the corresponding physical switch configuration.

Figure 1. Rear cabling reference for the temporary hosts, showing dedicated data, management and live-migration connections alongside 10Gb DAC SAN connectivity.

4. Resolving Legacy Hardware Compatibility Issues

Cluster validation initially failed on Hyper-V configuration checks even though the cluster, shared storage, witness, CSV layout and VLAN configuration appeared sound. Basic host-to-host connectivity also worked: ping, DNS resolution, TCP checks and the WinRM listener were reachable. The failure only became visible when higher-level remote management was exercised between the rebuilt hosts.

WinRM/WinRS, PowerShell remoting, CIM, classic WMI and Hyper-V WMI calls all failed or behaved inconsistently between the legacy nodes. Firewall rules, Group Policy, secure channel, time synchronisation, Kerberos tickets and SPNs were checked without finding a relevant fault. The asymmetric behaviour between sources and targets then pointed away from authentication and towards the host NIC/driver stack.

Area	Finding	Action / outcome
Basic connectivity	Ping, DNS and TCP 5985/135/445 succeeded; Test-WSMan showed a listener.	Confirmed that the problem was not a simple reachability or listener failure.
Management plane	WinRM, WinRS, PowerShell remoting, CIM, WMI and Hyper-V WMI failed host-to-host.	Narrowed the investigation to behaviour affecting management traffic rather than general networking.
Policy and authentication	Firewall, GPO, secure channel, time, Kerberos and SPN checks did not explain the failure.	Avoided making unnecessary policy changes once those layers were eliminated.
Legacy NIC stack	HP/Broadcom NC382i management adapters had offload and virtualisation features enabled.	Disabled problematic Priority/VLAN, VMQ, LSO and checksum-offload settings on the management NICs and rebooted the affected hosts.
Result	The same remote-management tests succeeded after remediation.	Host-to-host management was restored and cluster validation could proceed normally.

5. Validation and Resilience Testing

The temporary nature of the platform did not remove the need to prove that it behaved predictably. Three Generation 2 test VMs were used: Windows 11 with Secure Boot and vTPM, plus Ubuntu Desktop and Ubuntu Server with Secure Boot. For the Windows 11 VM, the required Shielded VM Local Certificates were made available across the cluster before live-migration testing.

Test	Method	Result
VM compatibility	Built Windows 11 and Ubuntu Generation 2 guests using the target VM network and shared storage.	All test VMs operated normally on the cluster.
Live migration	Moved all three running VMs between each of the three nodes.	Migrations completed successfully; VMs remained online and responsive.
Planned maintenance	Drained one cluster node while the test VMs were active.	VMs evacuated to other nodes and remained available.
Controlled host reboot and recovery	Placed the test VMs on one node and performed a controlled host reboot.	Cluster roles moved to another available node and recovered cleanly.
Final health	Rechecked node state, Cluster Shared Volumes and cluster networks after testing.	All three nodes, both CSVs and the management/live-migration networks were healthy.

Validation Warnings and Engineering Judgement

The Microsoft cluster validation report contained warnings, but they were reviewed against the purpose of the environment rather than treated as automatic defects. This mattered because unnecessary remediation could have introduced more complexity into a deliberately temporary design.

Validation item	Decision	Rationale
Cluster-name permissions	Broad computer-object creation rights were not granted to the temporary cluster.	If a clustered client access point was required, the object could be pre-staged with targeted permissions rather than giving a DEV cluster broad rights over the server OU.
DCB / QoS	No remediation required.	The temporary hosts used 1Gb networking and were not using RDMA, SMB Direct or a DCB-dependent storage design.
Storage validation	Not Applicable was accepted for the candidate-disk test.	The SAN disks were already clustered and online, leaving no free candidate disks for that specific validation step.

6. Migration Preparation and Hyper-V Replica

Before migration, the old DEV Hyper-V environment was treated as a discovery and cleanup exercise rather than assuming that its recorded state was accurate. VM inventory, VHD paths, network adapters, VLAN configuration, checkpoints and cluster roles were captured from the reachable source hosts. One old node was no longer available and was taken out of scope.

The preparation work exposed several items that could have complicated migration. Old checkpoints were removed through Hyper-V so their differencing disks could merge correctly, and a stale duplicate VM registration and matching offline cluster role were removed by ID while the valid running VM was left untouched. Network mapping also showed that most workloads used the standard DEV/UAT VM network while one management-related VM followed a separate network path, so I updated the target switch configuration and VM mapping to preserve both cases.

Hyper-V Replica was then used as the migration mechanism. Source and target Replica Broker objects were prepared with targeted AD permissions, replication was configured between the two clusters, and workloads were moved through planned failover only after replication health and network mapping had been checked. The old source VM and disks were deliberately retained, powered off, until boot, network, application access and live migration had been confirmed on the target. The retained source provided a pre-cutover recovery point; any subsequent changes on the target would need to be reconciled before returning to it.

Controlled Migration Sequence

No.	Stage	Key control
1	Discover source state	Inventory source hosts, VMs, VHDs, VLANs, checkpoints and cluster roles.
2	Clean migration blockers	Merge old checkpoints correctly and remove stale duplicate records without touching the valid workload.
3	Map target networking	Preserve the standard DEV/UAT VM network and the identified management-network exception on the new VM-Data path.
4	Configure Replica Brokers	Pre-stage the required AD objects and use targeted permissions instead of broad OU rights.
5	Replicate and cut over	Start Hyper-V Replica, verify health, then use planned failover during an agreed migration window.
6	Validate and retain rollback	Confirm boot, network, service access and live migration before accepting the target and retiring the old copy.

Migration Readiness Checks

- No Hyper-V checkpoints or active AVHDX differencing disks remained on the reachable source hosts.
- The stale duplicate VM and cluster role were removed while the valid running instance remained intact.
- Normal DEV/UAT workloads were mapped to the target VM data network; the management-network exception was handled deliberately rather than discovered during cutover.
- A low-risk workload could be moved first and validated before continuing with the remaining migration wave.

7. Outcome and Value Delivered

The project delivered a functioning three-node Hyper-V Failover Cluster from hardware that had been retained beyond its normal lifecycle for temporary use. By recovering serviceable components, installing Windows Server 2022 and validating live migration, node drain and controlled host reboot/recovery, I turned the available equipment into a practical interim platform.

Operationally, the cluster provided shared SAN storage, configured switch-side VLANs, separate management and live-migration networks, resilient VM data connectivity through SET, and tested failover behaviour. It also demonstrated that the limitations of old hardware were understood rather than ignored: the most significant management-plane fault was diagnosed below the obvious networking and authentication layers and remediated specifically for the legacy NIC stack.

The migration work reduced risk further by cleaning the old environment before cutover, documenting network exceptions and using Hyper-V Replica with a controlled rollback position. This allowed DEV/UAT workloads to move onto the interim platform while the wider datacentre transition continued and permanent hardware remained pending.

Engineering Decisions That Mattered

Decision	Reason	Practical effect
Recover three legacy hosts rather than build a single temporary server	Enough serviceable components were available to reconstitute a three-node platform.	Retained clustering and failover capability during the interim period.
Use Switch Embedded Teaming in Switch Independent / Dynamic mode	The VM data path needed two physical links without dependence on LACP configuration.	Provided aggregate VM-data capacity and link resilience with a simple switch-side design.
Configure host and switch traffic roles together	Management, live migration and VM traffic needed clear roles across both Windows and the physical switch ports, with SET used instead of LACP for VM data.	Kept VLAN and port behaviour predictable end to end, reducing contention and simplifying fault isolation.
Treat validation warnings according to design intent	Not every warning represented a defect for a temporary non-RDMA development cluster.	Avoided unnecessary permissions and networking complexity.
Retain the old source after planned failover	A temporary target should not remove the rollback option prematurely.	Kept a clear recovery position until the migrated workload was accepted.

Engineering approach

- Treated the temporary cluster with the same care as a permanent system, including separate network functions, resilient shared storage, failover testing and a clear rollback plan.
- Troubleshot the remote-management fault methodically, ruling out basic networking, firewall, policy and authentication issues before tracing the problem to settings on the legacy network adapters.
- Reused older hardware pragmatically, after checking its operation with Windows Server 2022 and testing live migration, node drain and controlled host reboot/recovery.
- Prepared for migration by checking and cleaning the existing environment first, resolving old checkpoints, duplicate records and network exceptions before moving workloads and keeping the original VMs available until the new platform had been verified.

End of case study