End-to-End VMware Hardware Refresh and vSphere 7-to-8 Upgrade

Professional portfolio project summary

Amandio Pereira | UK Engineering Manufacturer | 2024

Prepared as a concise technical narrative of the planning, procurement, build, migration, risk reduction and decommissioning work completed as the lead engineer for a VMware platform refresh.

Project objective	Replace the legacy three-host VMware estate, move production workloads onto the new hardware, upgrade the platform from version 7 to version 8, and retire the old environment with minimal operational disruption.
My role	End-to-end technical lead: discovery, design translation, procurement input, switch and host configuration, VM migration, licence transfer, upgrade sequencing, resilience remediation, testing and handover.
Business value	Improved storage performance and headroom, stronger host and storage resilience, budget stability through a three-year VMware term, and avoidance of a five-figure external implementation charge.
Core technologies	VMware vSphere 7/8, Dell PowerEdge servers, all-flash SAN storage, Cisco Nexus switching and Aruba wireless.

1. Executive Summary

This project replaced a legacy three-host VMware estate built on older HPE servers with a new three-host Dell PowerEdge platform and an all-flash storage array. The work covered more than a simple hardware swap. It began with discovery and documentation of the existing host, network and storage configuration, moved through procurement and design decisions, and finished with live workload migration, licence transfer, a controlled vSphere 7 to 8 upgrade, and the retirement of the previous hosts.

A deliberate sequencing decision reduced risk: the new hosts were first built to match the running version of the old estate, separating the hardware migration from the platform upgrade. Production VMs were then moved to the new hosts using vMotion, which transfers running virtual machines between hosts. Once the workloads were on the new hardware, the platform was upgraded to version 8 and the VM compatibility level raised.

The project also uncovered and resolved an operational weakness. The site wireless controller VM had effectively been tied to a single network path on one legacy host. I mirrored the required WLAN configuration on the new estate and added the necessary switch connectivity, removing that single point of failure. Testing confirmed that the controller could move between hosts using vMotion without observed service interruption. The added connectivity also removed its original dependence on a single host and network path.

Project metrics and commercial position

Metric	Detail
Approved platform	Three-host Dell platform with all-flash SAN storage and direct-attach cabling.
Design refinement	The design was refined to add NIC resilience and use direct cabling rather than a more expensive transceiver-and-fibre approach.
Competitive position	The selected all-flash solution remained below the alternative vendor proposal.
Licensing decision	A three-year VMware term was recommended instead of the usual one-year term, reducing exposure to licensing and pricing uncertainty during Broadcom’s transition to its new commercial model.
Implementation saving	The work was delivered in-house and avoided a five-figure external implementation charge.

2. Starting Point and Discovery Work

The preliminary discovery work was essential. Before designing the new estate, I mapped the legacy hosts in detail so that I could understand exactly how traffic, storage and management were being handled. The old environment was a three-host HPE ProLiant Gen9 cluster with an established but ageing compute and memory configuration.

The network mapping showed that the live configuration separated management/server traffic from vMotion and vSAN traffic across dedicated virtual switches. It also showed that older wireless-related virtual switches could be removed because they were deprecated. That documentation provided the confidence to distinguish what genuinely had to be reproduced and what could safely be retired.

My storage assessment identified a mixed estate combining vSAN with additional SAN capacity, and utilisation leaving limited comfortable headroom for future growth once resilience and lifecycle considerations were taken into account.

Legacy environment snapshot

Area	Observed legacy state	Why it mattered
Hosts	3 x HPE ProLiant Gen9 hosts	Older compute platform with limited lifecycle value and less room for resilience improvements.
CPU and memory	Established Gen9 compute and memory platform	Sufficient for the current footprint but no longer attractive as a forward-looking base.
Storage	vSAN plus additional SAN capacity	Showed both dependency and growth pressure; the target design needed cleaner, faster consolidated storage.
Networking	Dedicated virtual switches for management/server and vMotion/storage traffic	Defined the service lanes that had to be preserved for a safe migration.
Deprecated paths	Older wireless-related vSwitches still present in documentation	Highlighted the need to separate live requirements from historic clutter.

3. Target Architecture and Procurement Logic

The approved target design used three modern Dell PowerEdge hosts backed by an all-flash storage array. The hosts moved to a newer platform generation with improved NIC flexibility, while the array provided substantially increased capacity and dual-controller resilience.

The planning material also captured the network design in practical terms rather than only at a conceptual level. The new layout provided redundant high-bandwidth host-to-SAN connectivity, resilient production VM networking and dedicated vMotion and storage-service paths. Direct-attach cabling was deliberately chosen because it delivered the required outcome without the much higher cost of a transceiver-and-fibre approach.

Procurement was not treated as an admin afterthought. The commercial decision making formed part of the engineering outcome. The selected all-flash solution remained below the alternative vendor proposal, while additional NIC resilience improved supportability and a three-year VMware term reduced short-term budget uncertainty.

Old estate versus new estate

Component	Old estate	New estate	Practical impact
Compute hosts	3 x HPE ProLiant Gen9	3 x Dell PowerEdge 16G	Modernised the platform while retaining a familiar three-host cluster model.
CPU per host	Legacy Xeon generation	Newer Xeon generation	More current host capability and a cleaner support position.
Memory per host	DDR4	DDR5	Same capacity, faster platform generation.
Primary storage	vSAN plus additional SAN capacity	All-flash SAN with increased usable capacity	More capacity, better performance profile, cleaner storage model.
Connectivity	Mixed legacy link speeds and dependencies	Redundant high-speed switch and SAN paths	Improved resilience and cleaner separation of traffic types.
Support	Legacy support footprint	Current vendor support coverage	Reduced operational risk in the event of hardware failure.

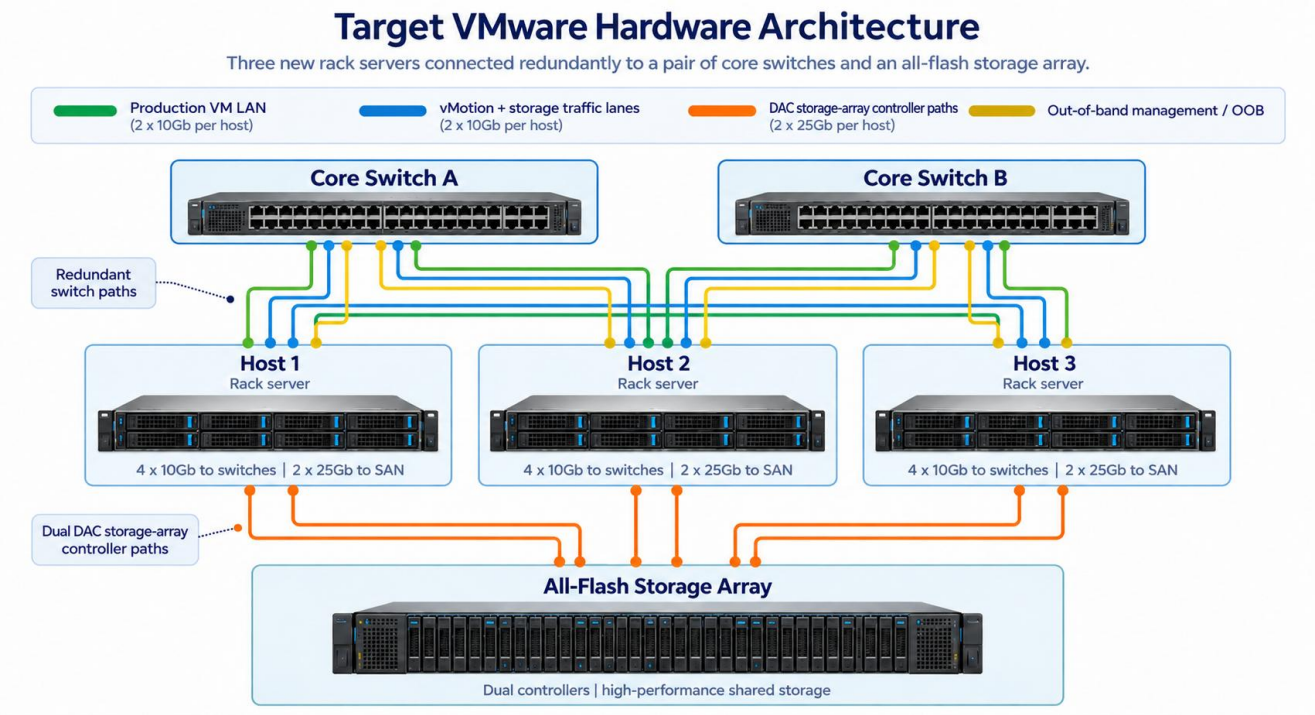

Figure 1. Simplified representation of the target host, switch and storage layout

4. Delivery Strategy and Sequencing

The high-level upgrade plan recorded in the original planning note was sound: build the new environment, connect it to the cluster, migrate the VMs, move the licensing, upgrade the platform, then remove the old hosts. The critical refinement in practice was sequencing the version alignment correctly so that the migration itself remained as safe and reversible as possible.

The new hosts were therefore introduced on a version compatible with the live old estate. This kept the hardware migration and platform upgrade as separate stages, making each change easier to validate. Once the workloads were running on the new hardware, the platform was upgraded from version 7 to version 8.

Migration sequence used on this project

The sequence reduced risk by preserving compatibility first and upgrading only after workload cutover.

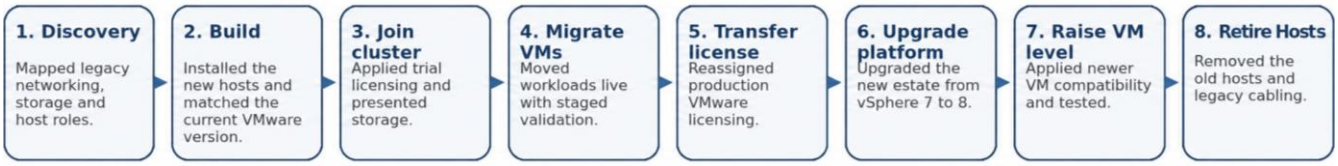

Figure 2. The staged sequence used for the refresh and upgrade.

Implementation phases

Phase 1 - Discovery and documentation: Confirmed the live host design, mapped VMkernel settings, identified deprecated virtual switches, captured current storage usage, and reserved switch ports for the new estate.

Phase 2 - Hardware and network preparation: Installed the new hosts and storage, allocated the required core-switch connectivity, and prepared the switch-side configuration needed for the host traffic lanes.

Phase 3 - Cluster onboarding: Configured trial VMware licensing on the new hosts, added the new hosts to the cluster, mounted the datastore, and validated storage and network presentation.

Phase 4 - Live migration: Migrated the production VMs to the new hardware while preserving operational continuity.

Phase 5 - Licence transfer and platform uplift: Moved the existing VMware licensing to the new hosts, upgraded the estate from version 7 to version 8, and then raised the VM compatibility level.

Phase 6 - Retirement: Removed the previous hosts from the cluster and decommissioned the old hardware and cabling from the core switches.

5. Network Design, Switching and Resilience Work

The switching work was a fundamental part of the project rather than a peripheral task. The new estate required clear port reservation, VLAN mapping and traffic separation so that management, production, storage and migration traffic would behave predictably.

In practice, this meant reserving dedicated host-facing connectivity on each core switch for the new estate, documenting the mapping, and connecting each host in a way that avoided dependence on a single switch or a single NIC. The resulting design also made troubleshooting and handover easier because the traffic roles were documented clearly.

Core-switch port allocation summary

Traffic lane	Core Switch A	Core Switch B	Purpose
Production VM LAN	3 dedicated ports	3 dedicated ports	Primary production / VM LAN links for all three new hosts.
vMotion / storage services	3 dedicated ports	3 dedicated ports	vMotion and storage-service lanes.
Storage-array management	Dedicated management port	Dedicated management port	Management connectivity for the storage-array controllers.
Out-of-band management	Dedicated management ports	Dedicated management ports	Host iDRAC and management connectivity.

6. Identifying and Removing the Wireless Controller Single Point of Failure

One of the most useful outcomes of the project was the identification of a hidden resilience gap before the old environment was decommissioned. The wireless controller VM that served building Wi-Fi was effectively tied to a single network path on one legacy host. A simple path failure could therefore have caused a wider Wi-Fi outage during the refresh.

To avoid that outcome, I recreated the necessary wireless-controller VLAN pathing on the new hosts as a temporary compatibility measure while the controller awaited wider standardisation. I first mirrored the older behaviour on one host, then extended it across the new cluster and added the required switch connectivity. The controller was no longer stranded on one host and one path.

The operational outcome was then tested properly, not assumed. Out-of-hours verification confirmed that the wireless controller VM could be vMotioned between hosts with no interruption to service while active wireless traffic was being exercised around the building.

Why this mattered

Risk	Action taken	Outcome
Wireless controller tied to a single network path on one legacy host	Mirrored the required wireless-controller VLAN configuration on the new estate and added the switch connectivity needed for the temporary workaround.	Controller no longer depended on one host and one network path during the migration window.
Potential Wi-Fi outage during host migration or failure	Enabled the controller VM to move across all three new hosts while the temporary configuration remained in place.	Improved continuity until the controller could be standardised onto the resilient VM LAN.
Risk of assuming resilience without testing	Performed multiple vMotion moves while streaming active traffic during an out-of-hours walk-through test.	No interruption of service was observed and the VM was left on the latest hardware compatibility level.

7. Outcome and Value Delivered

From an infrastructure point of view, the project delivered a more supportable host platform, materially stronger storage capability, clearer switch-level design and a cleaner upgrade path. From an operational point of view, it reduced dependency on ageing hardware and removed a previously hidden service-continuity risk around Wi-Fi delivery.

From a commercial point of view, the work combined technical and budget value. The selected all-flash solution remained below the alternative proposal, while the direct-attach cabling approach avoided a materially higher fibre/transceiver cost. With Broadcom's acquisition of VMware completed shortly beforehand and its licensing model beginning to change, I recommended securing a three-year VMware term rather than taking the usual one-year term, reducing exposure to further licensing and pricing changes. The refresh was also executed in-house rather than through an external implementation engagement, avoiding the associated service cost.

Engineering approach

- Reduced migration risk by moving the workloads to the new hardware first and only then upgrading VMware from version 7 to version 8.
- Documented the existing environment in detail before making changes, separating settings that were still required from older configuration that could be removed.
- Considered the servers, switches, storage and network connections as one system, designing resilience across them and testing that services could continue if workloads moved between hosts.
- Balanced resilience, maintainability and cost, choosing simpler options where they met the technical requirement without weakening the overall design.

End of case study