

AVS is not an architecture choice.

It's a migration strategy.

Treating it as a destination is where most teams get it wrong.

You're not choosing
where to run workloads.

You're choosing how hard
it will be to leave later.

AVS feels like staying on-prem — just relocated into Azure's billing model.

“

**AVS doesn't
remove lock-in.
It changes where
the lock-in lives.**

”

You're not escaping VMware. You're relocating it into a metered, provider-controlled environment.

The published price is for compute.
The real cost is in everything around it.

- ▶ 3-node minimum floor
- ▶ vSAN storage overhead
- ▶ NSX-T embedded licensing
- ▶ AVS ↔ Azure service egress (not always free)
- ▶ Dual control plane operational cost

AVS

Transition strategy

- ▶ Familiar VMware ops model
- ▶ Preserved constructs
- ▶ Higher hidden costs
- ▶ Exit = 2nd transformation
- ▶ Lock-in relocated

Native Azure

Destination architecture

- ▶ New ops model (front-loaded)
- ▶ Full Azure integration
- ▶ Standard egress model
- ▶ Exit = migration only
- ▶ Higher ceiling long-term

Native Azure exit is expensive. AVS exit is expensive and complex.

Are you using AVS to:

VALID

Buy time for a defined migration?

This is what AVS is designed for.

RISKY

Avoid retraining your team?

You're deferring a cost that compounds.

EXPENSIVE

Delay re-architecting legacy workloads?

Architecture debt doesn't dissolve in Azure.

Only one of these is a strategy. The others are deferrals.

AVS exit is not a migration. It's a second transformation.

On-prem exit

Physical + operational

Native Azure exit

Expensive

AVS exit

Expensive + complex + contractual

Model the exit before you commit to the entry.



If you don't have an exit plan, AVS becomes your destination by default.

Full breakdown — trade-off table, hidden costs, and exit risk architecture:

rack2cloud.com/azure-vmware-solution-vs-native-azure/