



WHITE PAPER

Data Migration

A controlled path from hyperscale cloud, VMware and legacy infrastructure to private, customer-directed compute



Public Edition | July 2026

Executive Summary

Migration is not a file-transfer project; it is a controlled change to a working service

Organizations migrate workloads to control cost, modernize infrastructure, improve placement, reduce platform dependency or align technology with data-residency and operating requirements. The transfer method matters, but the outcome is broader: the application must start, connect, perform and remain supportable in the destination environment.

GPCN™ provides a private destination for fixed, customer-directed compute, storage and networking. A migration can be executed as a complete virtual-machine move, a data-only transition into a reprovisioned environment, or a continuously synchronized cutover for high-volume and low-downtime scenarios.

MOVE THE WORKLOAD	MOVE THE DATA	MOVE AND SYNCHRONIZE
Capture and migrate complete virtual machines when operating-system, application and data consistency must be preserved.	Rebuild or modernize the target environment, then move files, databases and application data through native tools.	Use staged replication, change-data capture or parallel transfer when volume is high or cutover downtime must be minimized.

A defensible migration position

- Select the method according to workload dependencies, data volume, downtime tolerance and destination requirements.
- Preserve rollback until the destination has passed technical validation and user acceptance.
- Treat network, security, identity, licensing and operating procedures as part of the migration scope.
- Use familiar migration, backup, replication, file-transfer and database-native tools where they fit the workload.
- Validate the complete service at the destination before decommissioning the source.

THE GPCN APPROACH A structured migration begins with discovery and target design, proceeds through controlled movement and validation, and ends only when the destination is operating as intended.

Choose the Migration Path

The best method follows the workload rather than a single preferred tool

1 Discover	2 Design	3 Prepare	4 Move	5 Validate
Inventory workloads, dependencies, data, licensing and business owners.	Choose destination capacity, placement, network, security and migration method.	Build the landing zone, establish connectivity and test the transfer path.	Replicate, export, restore or transfer using the selected workload workflow.	Test access, performance, data integrity and rollback before final cutover.

Four questions shape the plan

What is moving? Complete VMs favor image-based backup, replication or conversion. Files and databases may move through native export, replication or transfer workflows.

How much data is involved? Moderate volumes may use existing connectivity. High-volume datasets may require dedicated bandwidth, parallel transfer, seeding or a longer synchronization window.

How much downtime is acceptable? A scheduled outage can simplify cutover. Low-downtime workloads usually require incremental replication, change-data capture and a shorter final synchronization.

Where must it run? Placement should reflect users, dependent services, residency, recovery design, private connectivity and operational ownership.

Tooling categories

- Backup and replication: complete workload capture, restore and staged cutover.
- File synchronization and bulk transfer: directory, object and unstructured-data movement.
- Database-native migration: export/import, replication and change-data capture.
- VM conversion and workload migration: image conversion, driver preparation and orchestration.

OPERATING RULE The migration tool does not replace the runbook. Define owners, checkpoints, acceptance criteria, rollback triggers and the point at which the source can be safely retired.

Migration from Hyperscale Cloud

Reuse what is portable, identify what is provider-specific and move in controlled stages

Infrastructure defined through Terraform or APIs can reduce the effort required to establish a new landing zone, but a cloud-to-private migration is not automatically a one-click redeployment. Virtual machines, open-source software and portable automation are often straightforward. Provider-native managed services, networking constructs, identity dependencies and Terraform state require explicit assessment.

Workload element	Typical treatment	Migration consideration
Linux and Windows VMs	Recreate or convert	Validate boot method, drivers, agents, storage layout and network configuration.
Terraform and API workflows	Reuse selectively	Portable modules can accelerate deployment; provider-specific resources and state require review.
Open-source databases	Rehost or replicate	Use native export, replication or CDC and validate consistency at cutover.
Public-cloud managed services	Retain or refactor	A managed-service dependency may remain in place or require an equivalent service and application change.
Networking and identity	Rebuild deliberately	Recreate addressing, routing, firewall, DNS, access and trust relationships in the target environment.

A practical transition sequence

Define. Inventory workloads, data flows, managed services, performance baselines and placement requirements.

Deploy. Build target compute, storage and private networking through supported portal, API or infrastructure-as-code workflows.

Validate. Confirm access, security, network paths, monitoring, backup and application behavior.

Connect. Replicate or transfer data and reconnect required public-cloud or external services.

Go live. Complete final synchronization, transition traffic, monitor the service and retire source resources only after acceptance.

HYBRID IS A VALID OUTCOME A workload may move to GPCN™ while continuing to consume selected AWS, Azure, GCP or SaaS services. The migration plan should account for latency, security, dependency and third-party data-transfer charges.

Migration from VMware

Start with evidence from the current environment, then build and validate the target before cutover

Organizations reassessing VMware cost, licensing or platform strategy need a migration path that is measurable and reversible. An RVTools-based assessment can establish the current inventory and inform target sizing, storage, network design and migration sequencing.

VMware source element		GPCN™ target consideration
vSphere / ESXi hypervisor	→	GPCN virtual machine infrastructure
vCenter management	→	GPCN™ management portal and supported automation
VMDK virtual disks	→	Destination-compatible VM images and validated guest boot
vSwitch / dvSwitch networking	→	Customer-defined private network configuration
VMware Tools and guest drivers	→	Destination-compatible guest drivers and agents
SAN, vSAN or VMFS storage	→	Right-sized SSD or NVMe-backed block storage

Five controlled phases

- 1 ASSESS** Collect RVTools output and identify workload owners, dependencies, exceptions and migration groups.
- 2 QUOTE & APPROVE** Confirm target profiles, storage, network, placement, responsibilities and commercial scope.
- 3 BUILD** Create the landing zone, connectivity, security controls and migration tooling.
- 4 MIGRATE** Replicate or convert workloads, prepare guest drivers and validate data and application function.
- 5 GO LIVE** Complete UAT, cut over DNS or traffic, monitor and retain rollback until acceptance.

MIGRATION BOUNDARY Compatibility must be confirmed workload by workload. Legacy operating systems, appliances, licensing controls, specialized drivers and tightly coupled VMware features may require remediation or an alternate path.

Data Movement and Cutover

Match the transfer method to volume, change rate, validation needs and acceptable downtime

Environment and VM-based migration

Use image-based backup, replication or conversion when the operating system, application and data must move together. This path is often appropriate for legacy systems, mixed Windows/Linux estates and workloads that are impractical to rebuild.

File and database migration

Move only the required data when the target compute environment can be reprovisioned. File-transfer, database export/import, replication and change-data-capture workflows can support modernization rather than preserving every source component.

High-volume and continuous synchronization

Use dedicated connectivity, parallel transfer and incremental synchronization when datasets are large or production must remain available until final cutover. The final outage can then be limited to the last synchronization and traffic transition.

Cutover acceptance checklist

- Data integrity and application transactions are validated against agreed criteria.
- Routing, DNS, firewall, identity, certificates and administrative access operate as designed.
- Performance and monitoring are compared with the approved baseline.
- Backup, recovery and support procedures are active in the destination.
- Business owners complete UAT and explicitly approve the transition.
- Rollback remains available until the agreed acceptance window has passed.

DATA-TRANSFER ECONOMICS GPCN™ does not charge egress fees for data leaving or moving within GPCN. AWS, Azure, GCP or another external provider may charge for data leaving its environment; those third-party charges are separate and are not imposed by GPCN.

Governance, Workload Fit and Next Step

A migration is complete when the destination is accepted, supportable and recoverable

Migration governance

- Name an owner for every workload, dependency, approval and cutover decision.
- Define security, residency, connectivity and access requirements before the landing zone is built.
- Document source and target baselines for compute, memory, storage, network and application performance.
- Use migration waves that group workloads by dependency and business impact.
- Retain execution evidence, exceptions, test results and formal acceptance.
- Verify backup and recovery in the destination before source infrastructure is retired.

Where the model fits

GPCN™ is well suited to steady-state and data-sensitive workloads that benefit from dedicated private infrastructure, customer-directed placement, predictable service profiles and explicit network design. Public cloud can remain appropriate for elasticity, provider-native services and existing architectural dependencies.

RECOMMENDED FIRST STEP Select one representative workload and complete a migration assessment covering inventory, dependencies, target sizing, network, transfer method, validation, rollback and operating ownership.

About GPCN™

The Global Private Cloud Network provides members-only access to private compute, storage and networks across a growing global provider footprint. Portal and API operations, customer-directed placement and predictable service profiles support controlled infrastructure delivery.

**BUILT WITH
TRUST**



 **99.99%**
availability

 **Tier 3**
Data centers

Available through EnTelegent Solutions | gpcn.entelegent.com | 800-975-7192