GPCN™

GLOBAL PRIVATE CLOUD NETWORK

WHITE PAPER

Extending Private Cloud Infrastructure Globally

A controlled, API-driven infrastructure model for steady-state, regulated and globally distributed workloads

GLOBAL	99.99%	NO EGRESS	API + IaC
Distributed participating provider footprint	Compute availability service commitment	None charged by GPCN; external fees may apply	REST API and documented Terraform workflows

Public Edition | July 2026

Executive Summary

Global reach without surrendering infrastructure control

Enterprises increasingly operate across regions, countries and regulatory environments. Core applications, data platforms and emerging AI workloads must remain secure, performant and governed wherever users, systems and dependent services are located.

Public cloud remains valuable for elastic and provider-native services. However, steady-state, data-sensitive and latency-dependent workloads often require a different operating model: predictable economics, explicit placement, controlled connectivity and infrastructure that does not move without customer direction.

THE GPCN APPROACH GPCN™ extends private cloud infrastructure across a growing network of participating data center locations through a consistent portal, REST API and private-network model. Customers choose where workloads run, how they connect and when capacity changes.

The global infrastructure challenge

As organizations expand, applications may remain centralized while users, data sources and dependent services become globally distributed. Traditional approaches can introduce several recurring challenges:

- Long-haul public-internet paths that introduce latency variability and congestion.
- Data-transfer and egress charges that grow with traffic volume and data gravity.
- Availability designs that may require duplicated capacity across zones or regions.
- Limited visibility into infrastructure placement, movement and provider-driven changes.
- Regional inconsistency that increases operational complexity and fragments governance.

A complementary cloud strategy

GPCN™ is not intended to replace every public-cloud service. It provides an additional infrastructure layer for workloads that benefit from fixed, controlled compute while allowing organizations to retain public-cloud platforms where native services, elasticity or existing architectural dependencies make sense.

Architecture and Control Model

A distributed private-cloud fabric with one consistent operating layer

GPCN™ provides a consistent access and control layer across participating private-cloud locations without operating as a conventional hyperscale cloud. Infrastructure remains private and purpose-built, while global reach is delivered through controlled connectivity and standardized operations.

1 Local access	2 Private transit	3 Explicit placement
Users, applications and services connect through the nearest appropriate GPCN™ location using public or private connectivity.	Private networks provide controlled connectivity using customer-defined addressing and network services.	The portal and API provision resources in the selected location. Placement and capacity changes remain customer-directed.

Local entry → controlled private fabric → customer-directed infrastructure

Control by design

GPCN™ is operated as explicit, customer-directed infrastructure rather than an orchestrated public-cloud service. Location, network design and capacity are selected during deployment, and requested changes remain deliberate. This model improves auditability and helps align infrastructure operations with residency, security and change-control policies.

API and infrastructure as code

A documented REST API provides programmatic access to infrastructure across supported locations. Authentication guidance, endpoint references, error handling, rate limits and asynchronous-operation conventions support integration, automation and lifecycle management.

Documented Terraform workflows support declarative provisioning and version-controlled infrastructure practices. Teams can integrate GPCN™ resources into existing CI/CD processes while retaining placement control and governance; additional guides cover Kubernetes deployment and LLM-hosting patterns.

Operational building blocks

- Consistent fixed-VM profiles with included SSD storage and customer-selected placement.
- Private networks with custom CIDR ranges, DHCP, DNS, public IPs and managed network interfaces.
- Persistent block volumes that can be provisioned, attached, detached and resized as requirements change.
- Portal and API lifecycle management with two-factor authentication, API and SSH keys, browser console access and role-based permissions.

Service Model and Workload Fit

Predictable infrastructure for workloads that do not need hyperscaler abstraction

Availability without forced duplication

GPCN™ provides a 99.99% compute availability service commitment for a single fixed VM without requiring multi-zone duplication. Services are delivered from Tier-3 and Tier-3+ facilities and are designed for deterministic infrastructure behavior. Application-level redundancy and disaster recovery can still be added where business requirements call for it.

Predictable economics

The operating model combines a simple monthly VM price with included SSD storage. GPCN™ does not charge egress fees for data leaving or moving within GPCN. If a workload exchanges data with AWS, Azure, GCP or another external service, that provider may charge for data leaving its environment; those third-party charges are separate and are not imposed by GPCN.

BEST FIT Steady-state compute, compliance-sensitive applications, backup and recovery, centralized SaaS platforms, VDI, virtualization transitions and data-intensive workloads that value predictable behavior and explicit control.

When to evaluate GPCN™

- The workload runs continuously or has a predictable capacity profile.
- Data-transfer charges materially affect total cost of ownership.
- Residency, sovereignty or explicit placement must be demonstrable.
- The organization wants private infrastructure without a long hardware procurement cycle.
- The application is portable and not dependent on hyperscaler-native managed services.

Governance remains customer-defined

GPCN™ separates infrastructure delivery from the customer's security, governance and management model. Organizations can apply their own operational tooling, policies and controls while using the platform for consistent infrastructure access across regions.

Ideal Use Cases

Common infrastructure patterns where control and predictability matter

Backup, archive and long-term retention

Sustained data movement and long retention periods can make egress and access charges a dominant part of public-cloud storage economics. GPCN™ can provide private infrastructure for backup and archive platforms with predictable storage costs, controlled residency and no GPCN egress fees within the service model.

Disaster recovery and business continuity

Physically independent private-cloud locations can support geographic recovery designs without forcing every production workload into a duplicated multi-region architecture. Recovery capacity and operational procedures remain explicit and customer-directed.

Long-lived virtualized and application platforms

Long-lived application and virtualized environments benefit from deterministic performance and explicit change control. GPCN™ can support migrations, platform transitions and temporary capacity extensions without tying the workload to provider-specific managed services.

SaaS platforms and centralized applications

Centralized application environments can serve distributed users through geographically appropriate private-cloud entry points. This preserves one operational platform while improving connectivity consistency for users and dependent services.

Virtual desktop infrastructure

Local access points connected to centralized desktop environments through the private service fabric can improve session stability without requiring a complete VDI stack in every geography.

AI and data-intensive workloads

CPU, GPU and storage resources can be placed close to data, users or service dependencies while maintaining explicit infrastructure control. This supports private AI development, inference and data-processing patterns where residency and data-transfer economics matter.

A Deliberate Alternative

Global private infrastructure without silent movement or consumption-driven complexity

GPCN™ gives organizations another way to extend infrastructure globally. Instead of abstracting placement and charging separately for every movement of data, the platform combines private regional capacity, controlled connectivity and a common API-driven operating layer.

The result is a practical complement to public cloud: organizations can retain cloud-native platforms where they add value while placing steady-state, regulated and data-intensive workloads on infrastructure designed for predictability, control and clear economics.

Five questions for infrastructure leaders

- Is the workload steady-state or predictably sized?
- Do egress and data-transfer charges distort the real operating cost?
- Does the organization require explicit workload placement or residency?
- Would a 99.99% single-VM service commitment reduce unnecessary duplication?
- Can the workload run without deep dependency on provider-native managed services?

NEXT STEP Identify one steady-state or data-sensitive workload, select the required geography and compare its complete operating model: compute, storage, connectivity, availability, egress and commitment.

About GPCN™

The Global Private Cloud Network provides members-only access to private virtual machines, GPU compute, persistent block storage and private networks across a growing, globally distributed provider footprint. The platform combines portal and API operations with customer-directed placement, predictable service profiles and infrastructure designed for control and global reach.

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