

RIGHT CLOUD RETAIL

The right cloud for every workload, at the store edge, in private cloud or in public cloud, with elastic scale for every peak at a cost you can predict.



A GTT SOLUTION BRIEF
June 2026

RIGHT CLOUD RETAIL

The right cloud for every workload, at the store edge, in private cloud or in public cloud, with elastic scale for every peak at a cost you can predict.

CONTENTS:

THE LANDSCAPE 3

WHERE GTT COMES IN 4

THE APPROACH 5

 Where each workload belongs 5

 From cloud sprawl to one operating model..... 6

THE SOLUTION..... 6

 Core Capabilities of Managed Hybrid Cloud Services 7

THE GTT PLATFORM ADVANTAGE 7

THE PROOF 8

WHY GTT..... 9

RETAIL SUB-VERTICALS SERVED 10

ABOUT GTT..... 11




THE LANDSCAPE

Retail has spent the past few years putting AI to work in the store aisles: computer vision for loss prevention, demand forecasting that adjusts by the hour, personalization that follows a shopper across channels. The retailers who have made it work are pulling away from the ones still running pilots. In fact, AI deployers achieve 2.3 times the sales growth and 2.5 times the profit growth of their peers, while fewer than a quarter of retailers have scaled AI (*IHL Group, 2025*). McKinsey sizes the generative-AI opportunity in retail at \$240 to \$390 billion a year (*McKinsey, 2024*). And the gap is widening, not closing.

What separates the two groups is rarely the AI itself. It is whether the systems underneath can get that AI the current information it needs, right at the store, the moment it is needed. An AI tool is only as good as the data reaching it and most store networks were never built to move that much, that fast. The cost of getting this wrong is easy to see; retailers are sitting on \$1.73 trillion in stock in the wrong place at the wrong time (*IHL Group, 2025*) and the fix is less about smarter forecasting than about seeing what is selling and where, as it happens. Getting AI to work in the store is a network problem before it is a math problem.

AI is the sharpest problem every retailer now has. Hybrid cloud has quietly become the default operating model, which is to say running workloads across on-premises, private and public cloud as one estate. Gartner expects 90% of organizations to adopt a hybrid approach by 2027 (*Gartner, 2024*), and 87.5% of IT decision-makers already call it the ideal model (*Nutanix, 2025*). Spending on private and hybrid cloud is growing nearly twice as fast as public cloud, pushed along by cost control, data sovereignty and the slow drift of workloads back from public cloud where the economics or the compliance never quite met expectations. Right-cloud is not a single destination. It is the discipline of running each workload where it performs best.

For retailers, the workload list keeps growing. Cloud-hosted POS such as Oracle MICROS and Shopify, cloud ERP led by SAP S/4HANA and its 2027 migration deadline, cloud warehouse and order management systems, AI and machine-learning models, e-commerce platforms, loyalty programs and customer data platforms. Each one needs managed infrastructure that reaches across the store edge, private cloud and public cloud. IDC expects 90% of the top 2,000 retailers to run edge computing, processing data at or near the store rather than backhauling it to a central cloud, and projects global edge spending to reach \$450 billion by 2029 (*IDC, 2025*).



CONNECT



SECURE



SIMPLIFY

THE CHALLENGES FOR RETAIL	
Latency kills real-time AI Backhauling every camera frame and sensor reading to central cloud adds delay that breaks computer vision, shelf monitoring and demand forecasting at the moment they need to act.	Two migrations at once SAP S/4HANA's 2027 deadline forces ERP and network modernization into the same window, with no slack for either to slip.
Peaks you pay for all year Black Friday and Singles Day need elastic capacity that on-premises hardware can only meet by over-provisioning twelve months for two.	Cloud runs ahead of the team Running across AWS, Azure, private and on-premises needs cloud-architect, Kubernetes and FinOps skills most retail IT teams lack, while a third of organizations now run more than \$12M a year on public cloud, much of it un-optimized <i>Flexera, 2026</i> .

IN RETAIL, CLOUD HAS STOPPED BEING A PLACE YOU SEND WORKLOADS. IT IS THE OPERATING MODEL THE BUSINESS RUNS ON, AND THE RETAILERS WHO GET PLACEMENT RIGHT CAN SAY YES TO AI IN EVERY STORE AND TRIPLE CAPACITY FOR A FLASH SALE, IN WEEKS RATHER THAN QUARTERS.

WHERE GTT COMES IN

GTT operates one managed hybrid cloud across the store edge, private cloud and public cloud, built on the GTT Envision platform. The operating principle to put to a customer is direct: you decide what runs where, and GTT runs it, secures it and keeps it performing, under one operating model and one accountable team. Network and cloud are not separate layers stitched together after the fact. Where most providers assemble hybrid cloud on top of a network, GTT operates hybrid cloud through it, on a global Tier 1 backbone GTT owns and runs.

That matters because hybrid cloud performance starts with connectivity, and the path from a store to a cloud application is only as good as the network beneath it. GTT already runs the network for most of the retailers it serves. More than 90% of GTT's distributed customers run their networking as a managed service through GTT. Right-cloud extends that same managed model up the stack, to the compute and the applications, so a retailer adds cloud capability as a partner relationship rather than another set of suppliers to manage.

THE RETAILERS PULLING AHEAD ARE NOT THE ONES WITH THE MOST CLOUD. THEY ARE THE ONES WHO RUN ANY WORKLOAD WHEREVER IT PERFORMS BEST, UNDER ONE OPERATING MODEL.

THE APPROACH

Most retail cloud estates were not designed. They accumulated. A POS supplier moved to its own cloud, a marketing team spun up another and ERP modernization added a third. Each arrived with its own contract, console and bill, and no single team owns the result end to end. GTT starts from the other end, with one operating model that decides placement on purpose.

FRAGMENTED MULTI-CLOUD	GTT: ONE OPERATING MODEL
Network and cloud bought and run as separate layers	Network and cloud operated as one platform, on GTT's global Tier 1 backbone
Placement decided ad hoc, supplier by supplier	Placement decided on purpose: edge, private or public, by performance and compliance
Variable consumption billing no one can forecast	Simple, predictable pricing finance can budget against
Specialist cloud skills the retail IT team has to hire	Cloud Advisory and a 24/7 managed service stand in for the skills gap
Several contracts, consoles and bills	One accountable team, one platform view in GTT EnvisionDX
Compliance posture varies by environment	Consistent governance and data residency across every environment

Where each workload belongs

The placement call is not guesswork. GTT's Cloud Advisory sorts workloads by what they actually need, latency, performance, compliance and cost, and puts each one in the tier that fits.

At the store edge	In private cloud	In managed public cloud
Computer vision, POS failover, real-time inventory and analytics	PCI-scoped systems, ERP steady-state, in-region customer data	Black Friday and Singles Day burst, dev and test, AI and ML model training
Latency-sensitive, and the data needs to stay local	Predictable performance, compliance and data residency	Elastic scale on demand, with no year-round over-provisioning

From cloud sprawl to one operating model

The shift is less about moving more workloads to the cloud and more about running the ones a retailer already has under a single model. GTT takes operational responsibility for provisioning, monitoring, patching, backup, disaster recovery, security governance and cost optimization, while the retailer keeps ownership and control through GTT EnvisionDX. Public cloud is the clearest example: GTT runs the day-to-day, and the retailer keeps its direct relationship with AWS or Azure rather than reselling capacity through anyone. One model, applied consistently, is what turns an accumulated estate into a designed one.

THE SOLUTION

GTT Managed Hybrid Cloud is our managed service that runs workloads across edge, private and public cloud as a single system. It is built from a small set of building blocks, all operated under one model and one platform, so a retailer can place each workload where it belongs without taking on a new way of working for each one.

At the store, edge cloud puts compute close to where data is created, for low-latency AI, POS failover and local data control. Retailers run their own workloads, computer vision, demand forecasting, POS failover, directly on the same GTT EnvisionEDGE device that handles networking and security, with no extra hardware in the aisle.

In **managed private cloud**, you get shared or dedicated capacity sized and secured to each workload, on GTT EnvisionCloud CORE for shared multi-tenant needs or GTT EnvisionPOD for the tier-1 retailer that requires fully isolated, dedicated infrastructure. The shared private cloud holds a PCI DSS v4.0.1 Attestation of Compliance assessed in January 2026, alongside ISO 27001, ISO 20000, ISO 22301, SOC 2 and CREST.

For elastic scale, **managed public cloud** lets GTT run the day-to-day across AWS, provisioning, monitoring, patching, backup, disaster recovery, security governance and cost optimization. Google Cloud and Azure are reachable today over Cloud Connect peering, with managed operations on the roadmap, so as a Google- or Azure-primary estate are connected now, they can be managed as the service expands. GTT does not resell hyperscaler capacity; the retailer keeps its direct provider relationship while GTT owns the operational outcome. Underneath it all, **Cloud Connect** keeps the path between environments private and low latency over GTT's global Tier 1 I backbone, and database services run the e-commerce, loyalty, customer and AI training data that sit behind the storefront.

Around those blocks sit the managed capabilities retail asks for first: encrypted Backup-as-a-Service for POS, inventory and customer data across hundreds of stores; Disaster-Recovery-as-a-Service that replicates critical workloads with cross-region snapshots for GDPR-, HIPAA- and ISO 27001-bound systems; VMware- and Kubernetes-as-a-Service for application consistency across edge, private and public; and GPU compute for training demand-forecasting and computer-vision models. The commercial design is the part buyers tend to remember. Where most providers bill on variable consumption, Managed Hybrid Cloud services give retail finance one clear view of spend and a budget that holds. Predictable, which is rarer in cloud than it should be.

That predictability is not a billing trick; it is an operating model. GTT manages capacity planning, cost governance and continuous optimization across every environment, watching for idle capacity, oversized instances and the slow creep of un-optimized spend that pushes retail cloud budgets an average of 17% over target (Flexera, 2026). Spend, anomalies and optimization actions surface in GTT EnvisionDX next to the network view, so finance and IT are reading from the same numbers rather than reconciling a stack of hyperscaler invoices after the fact.

Core Capabilities of Managed Hybrid Cloud Services

- **Edge (GTT EnvisionCloud EDGE):** compute at stores and distributed sites for in-store AI, POS failover, real-time analytics and local data control, with no added per-store hardware.
- **Private Cloud (GTT EnvisionCloud CORE shared, EnvisionPOD dedicated):** sized to each workload; shared cloud PCI DSS v4.0.1 attested (Jan 2026) plus ISO 27001, ISO 20000, ISO 22301, SOC 2 and CREST.
- **Public Cloud:** GTT runs AWS and Azure day-to-day, with the retailer keeping the hyperscaler relationship; Google Cloud reachable via Cloud Connect and on the roadmap.
- **Cloud Connect:** private peering to AWS Direct Connect, Microsoft ExpressRoute, Google Cloud Interconnect and Oracle FastConnect from 15 VDC nodes, scalable to 400Gbps.
- **Database Services:** managed databases for e-commerce, loyalty, customer and AI training data.
- **Hybrid Cloud and VDC:** 15 global Virtual Data Center locations across three continents; PCI-DSS compliant infrastructure-as-a-service with European data residency and free inter-zone transfer.
- **Value-add services:** Backup-as-a-Service, Disaster-Recovery-as-a-Service, VMware-as-a-Service, Kubernetes-as-a-Service and GPU compute, all under one operating model.

THE GTT PLATFORM ADVANTAGE

Right-cloud works because the cloud runs on the same platform as the network, not beside it. Cloud operations live in GTT EnvisionDX next to networking and security, on infrastructure GTT owns end to end.

GTT ENVISIONEDGE

The compute layer of cloud, in the store itself. The 3500, 5500 and 7500 models are available today, with 8500 and 9500 models adding redundant SSD storage and up to 128GB RAM in the second half of 2026. Retailers deploy their own workloads, AI models, POS failover, inventory processing, directly on the device that already runs networking and security, with capabilities service-chained on remotely after deployment. No second box in the back room.

GTT ENVISIONCORE

Nodes across three regions, alongside the 15 VDC locations, doing double duty: orchestrating virtualized network functions and hosting private cloud. Cloud Connect is built into the backbone, with direct peering to the hyperscalers from GTT EnvisionCORE nodes, so the path from edge to cloud stays on GTT's own network end to end.

GTT ENVISIONDX

The single management view across every cloud environment, edge, private and managed public, sitting alongside networking and security. Cloud operations run on the same platform as everything else, not a separate console bolted on the side.

THE PROOF

These are GTT's own retail customers, multinational operators running live workloads on real infrastructure, not reference architectures.



Distrelec: cloud applications that hold under load

Distrelec runs its ERP on AWS over GTT's Tier 1 global backbone, with the cloud path kept private and low latency rather than routed across the open internet. The platform serves more than 60,000 daily visitors and 7,000-plus daily shipments at near-100% uptime. Proof that a cloud-hosted core application can carry an e-commerce business at scale.



Kiabi: compliant cloud infrastructure across the estate

Kiabi runs across 431 stores in 26 territories on GTT infrastructure that holds PCI Level 1 and better than 99.7% uptime. It is the operational proof that PCI-compliant, data-resident cloud capacity can sit under a distributed retail estate without store-by-store exceptions, the same shared private cloud foundation GTT extends.



Also proven

- **Greenyard** consolidated 30-plus providers to one across 70 sites in 25 countries, saving €1M a year, the network foundation a managed hybrid estate sits on.
- **HMY** doubled bandwidth and cut cost 22% across four continents in a six-month migration.
- **Cloetta** runs a reliable global WAN on GTT, the connectivity layer every cloud workload depends on.

WHY GTT

Three things separate GTT in a right-cloud conversation, and they compound. Most providers can sell a retailer cloud capacity or network connectivity. Far fewer own both layers and run them as one accountable service. That ownership is what turns hybrid cloud from a stack of separate contracts into a single operating model, and it is where GTT's advantage starts.

CONNECT WITH CONFIDENCE ON A GLOBAL TIER 1 BACKBONE AND FLEXIBLE ACCESS	SECURE ONE PLATFORM FOR CLOUD AND NETWORK	SIMPLIFY WITH END-TO-END DELIVERY OWNERSHIP
GTT owns and operates a top-ranked global Tier 1 IP backbone reaching 170-plus countries, with 80% of customer traffic staying on-net and 140,000 customer locations already connected. Cloud Connect peers privately to AWS, Azure, Google and Oracle from 15 Virtual Data Center locations across three continents, so a workload in any region reaches its hyperscaler over a private, low-latency path rather than the open internet. The path from a store to a cloud application runs on GTT's own infrastructure end to end, which is why hybrid cloud performance starts here and not at the hyperscalers front door.	Edge, private cloud and managed public cloud run on the GTT Envision platform, with operations, security governance and a single management view in GTT EnvisionDX. That means consistent policy, unified visibility and one compliance posture across every environment, rather than a different security model for each cloud a retailer happens to use. The shared private cloud carries a PCI DSS v4.0.1 attestation alongside ISO 27001, ISO 20000, ISO 22301, SOC 2 and CREST, and because network and cloud sit on one control plane, that posture reaches from the store edge all the way to the hyperscaler.	GTT runs the cloud the way it runs the network: co-managed, with GTT owning provisioning, monitoring, patching, backup, disaster recovery and cost optimization while the retailer keeps full control and visibility through GTT EnvisionDX. Cloud Advisory guides what runs where, Professional Services run the large migrations including SAP S/4HANA, and the commercial model replaces variable consumption billing with one predictable number finance can budget against. For a lean retail IT team, that is the difference between hiring cloud architects, Kubernetes and FinOps specialists and then pointing the people they have at the work that differentiates the business. One accountable team, not a contract for every layer. This is also what separates GTT from the obvious alternatives. A hyperscaler's own managed service runs only its own cloud and stops at the edge of its network; a reseller marks up someone else's capacity and owns none of the path. GTT owns the backbone, runs the cloud on top of it and stands behind both with a single SLA, which is why the retailer gets one accountable team for the whole route from store to application rather than a different number to call for each layer

RETAIL SUB-VERTICALS SERVED

Where provider sprawl is worst depends on your corner of retail. The case for one partner stays the same:

SUB-VERTICAL	WHERE RIGHT CLOUD SHOWS UP
Grocery & Food Retail	Edge compute for demand forecasting and perishable waste reduction across thousands of stores, with cloud ERP modernization underway in most large chains.
Fashion & Apparel	Returns, AI and clienteling models that need GPU compute and customer-data platforms close to the storefront, scaled elastically for seasonal peaks.
General Merchandise & Big Box	Computer vision for shelf monitoring and loss prevention at the edge, with public-cloud burst capacity for high-volume e-commerce peaks.
Luxury & Specialty	Clienteling data kept in-region, in private cloud for residency, with the personalization models luxury buyers expect run close to the boutique.
Health, Beauty & Pharmacy	HIPAA- and GDPR-bound customer and prescription data protected by DRaaS and in-region private cloud, with edge compute for in-store services.
Convenience & Fuel	Forecourt and POS failover at the edge, where lost connectivity stops payment, with cloud-managed estate visibility across thousands of sites.
Consumer Electronics	Cloud-hosted ERP and high-traffic e-commerce proven at Distrelec, with elastic public-cloud capacity for launch-day and promotional spikes.
Wholesale & Distribution	Edge compute at warehouses and distribution centers for real-time fulfillment, with EDI (electronic data interchange, the structured messaging that moves orders between trading partners) and order-management workloads run across private and public cloud under one model.

NEXT STEPS

Right-cloud is a placement decision before it is a procurement one. GTT starts with where the workloads should run, then runs them.



- Assess.** GTT maps the current estate, what runs at the edge, in private cloud, in public cloud and on-premises, and identifies the workloads where placement is costing performance, money or compliance headroom.
- Place.** Cloud Advisory recommends the right home for each workload by latency, cost, performance and data residency, and sequences any migration, including SAP S/4HANA, against the retailer's peak-trading calendar so nothing moves during the trading window that matters most.
- Talk to GTT.** One conversation covers the full four-pillar suite, resilient connectivity, security, right-cloud and simplified operations, with a complimentary cloud workload-placement and TCO assessment to put concrete numbers against the move.

CONNECT. SECURE. SIMPLIFY.

ONE OPERATING MODEL FOR EVERY WORKLOAD, AT THE STORE EDGE, IN PRIVATE CLOUD OR IN PUBLIC CLOUD, AT A COST YOU CAN PREDICT.

ABOUT GTT

GTT is a leading networking and security as a service provider for multinational organizations, simply and securely connecting people and agents to data and applications anywhere in the world. We serve thousands of organizations, bringing together the right people, partners and technology to reduce the burden on IT teams and solve the most pressing networking and security challenges. Built on our top-ranked global Tier 1 network, GTT Envision is a single global technology platform to connect, orchestrate, virtualize and automate enterprise networks, enabling customers with consumable solutions to achieve business missions and meet ongoing demand when, where and how needed. Our portfolio includes SASE, SD-WAN, security, internet, voice and other connectivity options, complemented by a suite of professional services and sales and support teams in local markets around the globe. We partner with our customers to deliver Greater Technology Together. For more information, visit www.gtt.net.

ENDNOTES

External market-stat sources, numbered. GTT facts and customer outcomes are GTT-sourced and covered by The Proof.

1. IHL Group, “Retail Inventory Crisis Persists Despite \$172B in Improvements,” September 2025.
2. McKinsey, “LLM to ROI: How to scale gen AI in retail.”
3. Gartner, worldwide public cloud forecast, November 2024.
4. Nutanix, Enterprise Cloud Index, 2026
5. Flexera, State of the Cloud Report, 2026.
6. IDC FutureScape: Worldwide Retail Predictions.
7. IDC. Global edge computing spending to reach \$450 billion by 2029, growing at ~15% CAGR.
8. SAP, S/4HANA 2027 end-of-life

