DSR PROTOCOL · EU DIGITAL PRODUCT PASSPORT INFRASTRUCTURE

Why the bill stays flat

In one line: most platforms bill like a taxi meter — the price keeps ticking the more your products are scanned. DSR works like a monthly travel pass — you pay for the products you look after, and scanning is included. Here is why, in plain terms.

Two separate cost lines — and both stay flat

This page is about the everyday GS1 Digital Link scan — the ordinary consumer and supply-chain resolution that fires every time a code is scanned, served at the edge in typically under 30 ms and never metered by DSR. That is different from, and sits on top of, the one-off independent back-up (the ISO 27001 escrow vault that meets the Article 10(4) continuity duty, priced as its own thin layer). A usage-based platform meters the everyday scan; DSR does not — so the saving below is on the ordinary scan itself, before the back-up is even counted.

The usual way — the cost climbs as you grow

On a central-server platform, every single scan wakes up their cloud to do work and send data back. The two biggest meters are the computing it runs and the database read it performs on every scan; bandwidth (egress) is a third — real, but at a few-KB resolve usually the smaller euro line. All of it rises with the number of scans, so the bill grows exactly when a product does well: the hardest moment to budget for.

The DSR way — flat and predictable

DSR serves each scan from the edge — answered from a cache kept close to the shopper, bypassing origin app servers and the database read entirely — and charges per product you keep alive, not per scan. The meters below barely move as scans grow, so you can forecast next year from your product count, not from a guess about how popular a line becomes.

Where the money goes — compute first, egress second

1 · Computing — the read work (primary saving)

What it means: The origin app-server work to build and answer each scan — the largest cost at volume.

PER-SCAN PLATFORM

Every scan wakes central app servers (e.g. Node/Fastify); traffic spikes auto-scale them up, and the cost with them. At a few-KB payload this compute dwarfs the bandwidth.

DSR · EDGE-CACHED

Resolved in the edge worker with a high cache-hit ratio (90%+ by design), so most scans never touch an origin server — each extra scan adds next to no compute cost.

2 · Database reads — the read tax (primary saving)

What it means: Reading the product record from the database on each scan.

PER-SCAN PLATFORM

Each scan is a live database read (PostgreSQL / Redis), metered and scaled with read-replicas for spikes — a cost that climbs linearly with volume.

DSR · EDGE-CACHED

Served from the edge cache, so there is no central database read on each scan — this cost does not climb with traffic.

3 · Bandwidth / egress (secondary)

What it means: The charge when product data leaves the cloud to reach a phone.

PER-SCAN PLATFORM
Billed per gigabyte. For a few-KB resolve this is often the smaller euro line — but it grows with media, documents and busy seasons.

DSR · EDGE-CACHED
In-perimeter reads at the edge add no bandwidth charge — and, just as importantly, keep data in-region for residency.

4 · Shape of the bill

What it means: How predictable the total is, month to month.

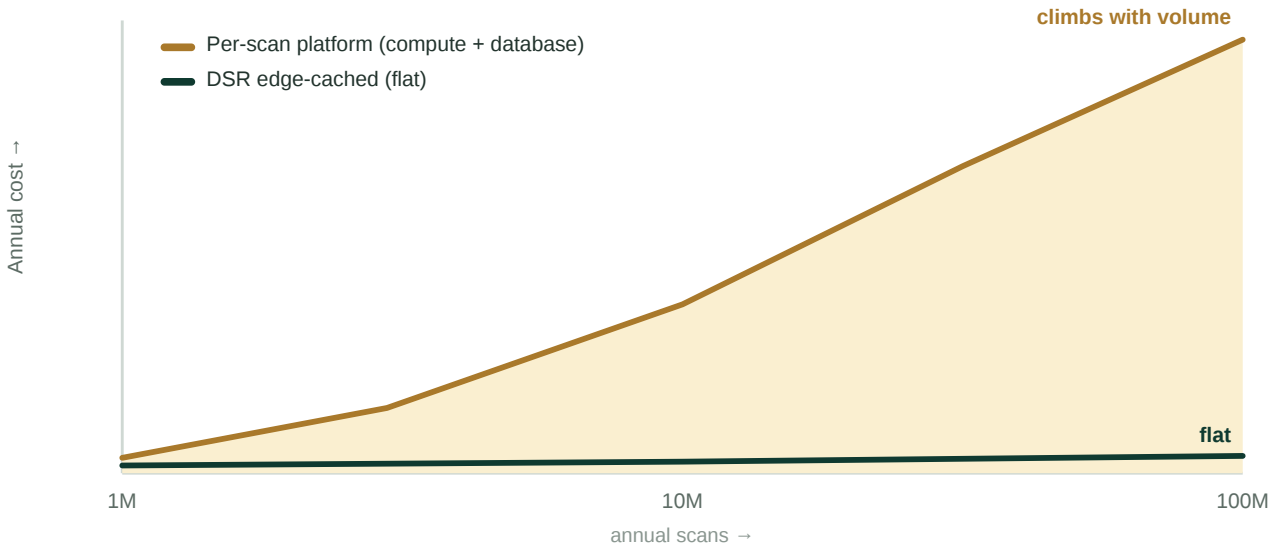
PER-SCAN PLATFORM
Rises with success — a surprise invoice exactly when you are winning.

DSR · EDGE-CACHED
Flat and predictable — set by how many products you store, not how often they are scanned.

What happens at retail scale — 1M to 100M+ scans

A usage-based cloud origin runs app-server compute and a database read on every scan (egress is a smaller, secondary line), so the bill climbs roughly in step with volume. An edge-cached engine resolves each scan at the edge from cache, so the cost stays effectively flat as volume grows. Figures are illustrative order-of-magnitude ranges at published cloud list rates (cloud infrastructure is billed in US dollars); an independent assessment — your actual rates depend on provider and traffic.

The read-tax curve — cost as scans grow



Illustrative shape, not to scale: compute and database reads drive the climb; DSR stays flat as scans grow.

ANNUAL SCAN VOLUME	USAGE-BASED CLOUD ORIGIN	DSR EDGE ENGINE	NET SAVING
1,000,000 scans / year	≈ $5,000 – $15,000	under ≈ $20	≈ 99% lower
10,000,000 scans / year	≈ $50,000 – $150,000	≈ $50 – $150	≈ $100k+ saved
100,000,000 scans / year	≈ $500,000 – $1,500,000+	≈ $500 – $1,000	≈ $1M+ saved

What this means for a national retailer (the scale of a Migros or CarrefourSA): thousands of private-label SKUs across fresh-food traceability, dynamic promotions and recall management generate scans at exactly this scale. A usage-based meter turns everyday customer engagement into an ongoing cloud tax; DSR's flat, in-perimeter edge resolution removes it — so unlimited engagement and compliance campaigns carry no variable cloud overhead. Illustrative; no affiliation.

Greener by architecture — a lower digital carbon footprint

The same edge-cached design that keeps the bill flat also cuts energy and carbon per scan: there are no idle servers waiting for a traffic spike, and each scan takes a shorter trip.

No idle infrastructure

A usage-based stack pre-provisions compute clusters, load balancers and database read-replicas that sit powered-up for peak traffic. DSR serves cached reads from the global delivery edge and keeps the origin and database off the scan path — no dedicated fleet burning power at idle.

Fewer network hops

A central-origin scan travels mobile network → load balancer → origin database in a central data centre and back. A repeat DSR read is returned from the nearest edge of the delivery network (a cache hit, with the origin and database kept off the scan path), so each scan makes a shorter trip and consumes less network and data-centre energy.

ESG & Scope 3 alignment

A lower digital carbon footprint per million queries supports an enterprise partner's net-zero and Scope 3 reporting — and strengthens the Teknopark R&D sustainability case.

Resource-efficient architecture — no data-egress fees

Structural, not statistical. Scale-to-zero serving means no idle servers draw power between scans; a repeat read is returned from an edge cache, keeping the origin and database off the scan path; and hosted artefacts sit on object storage (R2) with no data-egress fees. We do not publish a per-scan carbon figure — a defensible one would need an audited lifecycle assessment — so the point here is the architecture and the bandwidth it avoids, not an estimate.

What this means for you

- Predictable budgeting — the bill does not jump when a product takes off.
- Materially lower running cost, because you are not paying a meter on computing, database reads and bandwidth on every scan.
- EU-resident, edge-cached serving — data stays close to EU shoppers, with no third-country dependency in the custody path (good for KVKK / GDPR).
- One QR code covers both compliance and shopper rewards — one system to buy and run, not two.
- Savings compound with scale — at 100M scans the gap is hundreds of thousands against a flat edge cost in the hundreds, a 90%+ reduction.
- Lower digital carbon — no idle servers and sub-30 ms edge resolution cut energy per scan, supporting net-zero and Scope 3 goals.

Three lines for the board

1. Financial scale: at 100 million scans a usage-based platform can bill an enterprise retailer hundreds of thousands — into the millions — in cloud egress and compute; DSR handles the same volume for a flat edge cost in the hundreds, a materially lower, 90%+ reduction at scale.
2. Zero infrastructure overhead: we run no heavy origin servers or idle database clusters, so the architecture scales without adding server hardware or operational complexity.
3. Sustainability & ESG: resolving at the local edge in typically under 30 ms removes extra network hops and server compute cycles, lowering the digital carbon footprint for partners pursuing net-zero.

An honest note: the shape of the cost — a flat pass versus a climbing meter — is the point here. The scale figures are illustrative order-of-magnitude ranges at published cloud list rates, and the latency figures are architectural estimates; all are an independent assessment, not a quote, and your actual rates depend on provider and traffic. Named retailers are illustrative only, with no affiliation.
