SYMBEEO

PUBLIC

DESIGN STUDY

Private AWS-to-Google Cloud link for a game studio's live telemetry

Complete design on current platforms

CAPABILITY	INDUSTRY	SETTING
Cloud	Media & Gaming	European game studio
DOCUMENT	**VERSION**	**DATE**
CLD-003	1.1	Oct 10, 2026

WRITTEN BY

Steve Phillips, Principal Data Center Architect

© 2026 Symbeeo. All rights reserved.

Document control

Document	CLD-003	Version	1.1
Issued	Oct 10, 2026	Status	Issued by Symbeeo
Classification	Public	Supersedes	1.0, Sep 27, 2026
Capability	Cloud	Industry	Media & Gaming
Author	Steve Phillips	Role	Principal Data Center Architect
Credentials	CCSP, CCIE, AWS Certified Security – Specialty		
Approval	Approved by Edwin Ozer, Founder, Oct 10, 2026. Issued by Symbeeo, Oct 10, 2026. Design study: every figure is vendor-published, a design target or a marked assumption.		
Sign-off	"An engagement under this design is signed by the practice lead before any build proceeds."		

Contents

1 · Scope 3
1c · Public implementations this design follows 4
2 · Requirements register 4
3 · Architecture overview 5
4 · Design by subsystem 6
5 · Failure behavior 7
6 · Limits 8
7 · Lessons learned 8
8 · The author 8

© 2026 Symbeeo. All rights reserved. · Public

CONVENTIONS

Source tags. [S] figure read from the vendor's own current page on Sep 24, 2026, or on Sep 27, 2026 where marked · [S*] figure from the vendor's page summary where the page itself could not be opened · [A] design assumption, to be confirmed by the signing engineer · [T] design target: the record's outcome figure and anything derived from it. No figure comes from any engagement, file or asset. A row with more than one tag is counted once: as a published figure if any tag is [S] or [S*], otherwise as an assumption if any tag is [A], otherwise as a design target.

Table 1. Revision history

Date	Version	Change
Sep 27, 2026	1.0	First issue.
Oct 10, 2026	1.1	Approval line now reads: approved by the Founder, issued by Symbeeo. Steve Phillips remains the author.

1 · Scope

Who it is for. The head of platform engineering at a game studio whose live-service backends run on AWS and whose player analytics and model training run on Google Cloud, with telemetry crossing between the two over IPsec tunnels on the internet that saturate on launch weekends.

Typical buyer [A]. A European studio with three live titles. Game backends on Amazon EKS in AWS Europe (Frankfurt); analytics in BigQuery and Cloud Storage in Google Cloud Frankfurt. Telemetry from backends to analytics: 2.5 Gbps average, 6 Gbps at launch-weekend peak. A nightly batch of 20 TB from Amazon S3 to Cloud Storage. The studio's offices and build farm reach both clouds from one Frankfurt colocation facility through an existing edge router pair with 2 × 10 Gbps Direct Connect and 2 × 10 Gbps Dedicated Interconnect.

What is in. The cloud-to-cloud link, its attachment on each side, the address summaries and private name resolution across it, the fallback through the colocation edge, and monitoring. This is SRA-CLD-001 decision D5 and requirements R7–R11 applied to one pair of clouds. **What is out.** Landing zones (assumed in place on both sides); the colocation edge itself; game-server capacity; the analytics pipeline.

Assumptions. Both landing zones exist with a network hub account on AWS and a host project on Google Cloud [A]. The two clouds' address plans do not overlap today [A]. Telemetry producers can queue for up to 30 minutes [A]. The studio stays on these two clouds; Azure is not in scope [A].

Related documents, read Sep 24, 2026.

- AWS, AWS Interconnect is now generally available — https://aws.amazon.com/blogs/aws/aws-interconnect-is-now-generally-available-with-a-new-option-to-simplify-last-mile-connectivity/
- AWS, AWS and Google Cloud collaborate to simplify multicloud networking — https://aws.amazon.com/blogs/networking-and-content-delivery/aws-and-google-cloud-collaborate-to-simplify-multicloud-networking/
- Google Cloud, Partner Cross-Cloud Interconnect for AWS overview — https://docs.cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-cci-for-aws-overview
- Google Cloud, Cloud Interconnect release notes — https://docs.cloud.google.com/network-connectivity/docs/interconnect/release-notes
- AWS Direct Connect quotas — https://docs.aws.amazon.com/directconnect/latest/UserGuide/limits.html
- AWS Direct Connect, BGP settings (BFD; default hold timer 90 s, read Sep 27, 2026) — https://docs.aws.amazon.com/directconnect/latest/UserGuide/bgp-specific-settings.html

© 2026 Symbeeo. All rights reserved. · Public

- AWS Interconnect Service Level Agreement (last updated Apr 15, 2026), read Sep 27, 2026 — https://aws.amazon.com/interconnect/sla/
- AWS Interconnect FAQs (the 99.99 % SLA "up to the Direct Connect port" is stated for Interconnect – last mile), read Sep 27, 2026 — https://aws.amazon.com/interconnect/faqs/
- Google Cloud, Cloud Interconnect SLA (names Dedicated, Partner and Cross-Site Interconnect), read Sep 27, 2026 — https://cloud.google.com/network-connectivity/docs/interconnect/sla

1c · Public implementations this design follows

Table 2. Public implementations this design follows

Implementation	Who published it · date	What it settles for this design	URL
A SaaS company's published statement on AWS Interconnect – multicloud: bridges to Google Cloud "with the same ease as deploying internal AWS resources"; quad redundancy across facilities; MACsec between the providers' edge routers; API specification open on GitHub	AWS Networking blog · Dec 8, 2025 · public deployment	§4.1: a published cross-cloud user on the same product and the same pair	https://aws.amazon.com/blogs/networking-and-content-delivery/aws-and-google-cloud-collaborate-to-simplify-multicloud-networking/
General availability with Google Cloud in five region pairs including Europe (Frankfurt); MACsec on every physical link; links over at least two facilities (the post's 99.99 % SLA "up to the Direct Connect port" is stated for Interconnect – last mile, not multicloud)	AWS News blog · Apr 14, 2026 · vendor documentation, used to build our own design	R3–R5, R7	https://aws.amazon.com/blogs/aws/aws-interconnect-is-now-generally-available-with-a-new-option-to-simplify-last-mile-connectivity/
Partner Cross-Cloud Interconnect for AWS with VPC Network Peering generally available; Network Connectivity Center support in Preview; locations added July 14, 2026	Google Cloud release notes · Apr 14, 2026 and Jul 14, 2026 · vendor documentation, used to build our own design	R8 and §4.3: peering now, NCC later	https://docs.cloud.google.com/network-connectivity/docs/interconnect/release-notes
AWS Interconnect SLA: 99.99 % Monthly Uptime Percentage per covered interconnect, service credits the sole remedy; excludes problems beyond the demarcation point of AWS Interconnect; "Unavailable" means no data sent or received for at least 120 consecutive seconds	AWS · last updated Apr 15, 2026 · vendor documentation, used to build our own design	R6, R11 and §6	https://aws.amazon.com/interconnect/sla/

2 · Requirements register

Table 3. Requirements register

#	Requirement	Value	Source
R1	Cross-cloud throughput	≥ 6 Gbps sustained at launch-weekend peak	[A]
R2	Capacity steps	pre-approved speeds from 1 Gbps to 100 Gbps; 10 Gbps ordered	[S] Google partner CCI; [T] size

© 2026 Symbeeo. All rights reserved. · Public

#	Requirement	Value	Source
R3	Region pair	AWS Europe (Frankfurt) with Google Cloud	[S] a GA pair since April 14, 2026
R4	Encryption	IEEE 802.1AE MACsec on every physical link between the providers' routers	[S]
R5	Resilience	logical links across at least two physical facilities	[S]
R6	Availability	99.99 % monthly uptime per covered interconnect on the AWS side, service credits only; Google covers its own portion under its own SLA	[S] AWS Interconnect SLA, Apr 15, 2026; Google partner CCI overview
R7	AWS attachment	Direct Connect gateway to the existing Transit Gateway	[S] attaches to a DX gateway, then VGW, TGW or Cloud WAN
R8	Google attachment	VPC Network Peering to the analytics VPC; NCC only after GA	[S] peering GA, NCC Preview
R9	Address summaries	0 overlapping prefixes; ≤ 8 summaries advertised by each side	[T]
R10	Names	private zones on each side resolve from the other, both directions	[T]
R11	Fallback	telemetry continues through the colocation edge at ≥ 2 Gbps while the link is down; traffic on the fallback ≤ 90 s after the whole link is lost	[T]; [A] planned on the 90 s Direct Connect default BGP hold timer [S]; the managed link's own timers are not published
R12	Fallback detection	BFD 300 ms × 3 on the Direct Connect sessions	[S] Direct Connect defaults
R13	Provisioning	link carrying traffic ≤ 1 working day after the order	[T]; [S] Google lists provisioning in minutes
R14	Monitoring	synthetic probes over the link, alarm on loss or latency	[S] CloudWatch Network Synthetic Monitor bundled

Fourteen rows; eleven carry a vendor figure, one is a design assumption, two are design targets.

3 · Architecture overview

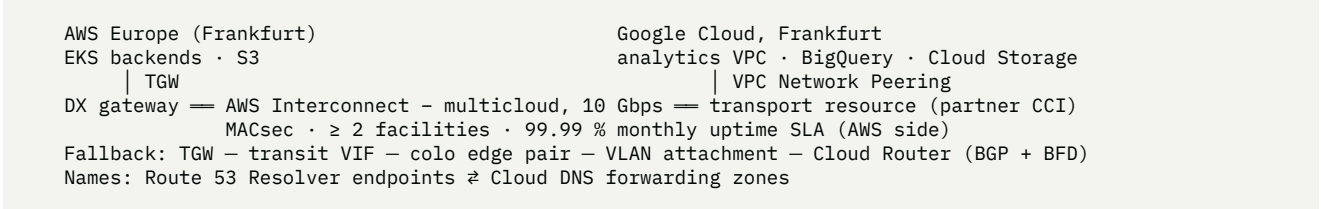


Figure 1. Private AWS-to-Google Cloud link

One managed 10 Gbps link between the two providers' routers replaces the internet tunnels. The studio orders capacity instead of ports and cross-connects; the colocation edge it already owns becomes the fallback path, not the primary

© 2026 Symbeeo. All rights reserved. · Public

one.

Table 4. Architecture overview

Role	Product (exact model)	Qty	Capacity used	Why this unit
Cross-cloud link	AWS Interconnect – multicloud	1 connection	6 Gbps peak of 10 Gbps [A] [T]	MACsec, two facilities, 99.99 % monthly uptime SLA on the AWS side [S]
Google side	Partner Cross-Cloud Interconnect for AWS (transport resource)	1	one transport per project per region [S]	GA with VPC Network Peering [S]
AWS hub	AWS Transit Gateway with Direct Connect gateway	1 + 1	game VPCs and the link [A]	up to 6 Transit Gateways per DX gateway [S]
Fallback edge	existing colocation edge pair	2 routers	2 × 10 Gbps DX, 2 × 10 Gbps Dedicated Interconnect [A]	already carries office and build-farm traffic [A]
Link monitoring	Amazon CloudWatch Network Synthetic Monitor	1	the link path [T]	bundled with the interconnect [S]

4 · Design by subsystem

4.1 The link.

AWS Interconnect – multicloud became generally available on April 14, 2026 with Google Cloud in five region pairs, Europe (Frankfurt) among them [S]. "Every connection uses IEEE 802.1AE MACsec encryption" on the physical links between the two providers' routers, and each connection spans logical links in at least two facilities [S]. Pricing is a flat hourly rate for the requested capacity [S]. 10 Gbps carries the 6 Gbps peak at 60 % **[A][T]**; the 20 TB nightly batch, rate-limited to 7 Gbps so night telemetry keeps 3 Gbps, takes about 6.4 hours, inside the night [A]. Rejected: HA VPN over the internet — no committed throughput, fails R1 at peak. Rejected: Google Cross-Cloud Interconnect on dedicated ports — the studio orders and runs the ports on both sides and waits weeks, where the partner product provisions in minutes [S].

4.2 AWS side.

The connection attaches to a Direct Connect gateway associated with the existing Transit Gateway [S]. A dedicated Transit Gateway route table for the game VPCs propagates the Google summaries; the build-farm and office routes stay on their own table [A]. The colocation edge prepends its AS path to the Google summaries so the direct link is always preferred (R11) [A]. Rejected: Cloud WAN in place of the Transit Gateway — one region, one hub; a global core network adds cost here without a second region to join [A].

4.3 Google side.

The link appears as a transport resource; each project is limited to one transport per region [S], so it lives in the connectivity host project. Partner Cross-Cloud Interconnect for AWS is generally available with VPC Network Peering, while Network Connectivity Center support is in Preview [S]; the design peers the transport to the analytics VPC only. Peering is not transitive [A], so a second Google VPC needs its own peering until NCC reaches GA. Rejected: an NCC spoke now — Preview, not for a production path.

© 2026 Symbeeo. All rights reserved. · Public

4.4 Addresses and names.

AWS side 10.64.0.0/12, Google side 10.80.0.0/12, as in SDD-CLD-001 §4.2 [A]; each side advertises at most eight summaries (R9). Route 53 Resolver inbound and outbound endpoints in the hub VPC, Cloud DNS forwarding zones in the host project; each forwards the other's private zone (R10) [A]. Rejected: public DNS names for internal endpoints — leaks the internal map and fails R10's private intent.

4.5 Fallback.

With the link down, the Transit Gateway learns the Google summaries only from the colocation edge, and traffic hairpins through the facility (SRA-CLD-001 F7). Direct Connect sets BFD at 300 ms minimum interval and multiplier 3 by default [S]; the edge routers match it, so a circuit loss inside the fallback path is detected in under a second. Telemetry is rate-limited to 2 Gbps on the fallback and the nightly batch waits (R11) [T]. Rejected: a second interconnect in the London pair [S] — the game backends are not in London, so it would still cross regions.

4.6 Monitoring and capacity.

Network Synthetic Monitor is bundled with the link [S] and probes from the game VPC to a target in the analytics VPC [A]. The capacity rule reads live telemetry only: the nightly batch is excluded, because it is rate-limited to 7 Gbps by design and would otherwise sit inside the 95th percentile every day (6.4 h of 24 is 27 % of the samples). Telemetry is measured from VPC flow logs with the batch's transfer endpoints filtered out [A]. Capacity rises to the next pre-approved step when the 95th percentile of live telemetry passes 60 % of 10 Gbps for three consecutive days [T], or when the batch no longer finishes inside the night at 7 Gbps [T]. The hourly charge is set against the current internet egress and VPN cost at each renewal.

5 · Failure behavior

Table 5. Failure behavior

Failure	Behavior	Source
One facility carrying the link lost	the logical links in the other facility carry the traffic	[S] at least two facilities
Whole cross-cloud link down (both facilities)	BGP withdraws; telemetry via the colocation edge, rate-limited to 2 Gbps, target ≤ 90 s (R11); producers queue meanwhile; batch waits; an outage under 120 s does not count against the SLA	[T] R11; [A] 90 s planned on the DX default hold timer; [S] 120 s SLA threshold; [A] queue
One Direct Connect circuit down while on fallback	BFD detects it in under 1 s; the second circuit carries traffic	[S] 300 ms × 3
Peak above 10 Gbps	producers queue up to 30 min; capacity step ordered	[A] queue; [T] rule
One DNS resolver endpoint down	the second endpoint answers; IP paths unaffected	[A]
Google analytics VPC re-created	peering to the transport re-applied by pipeline	[A]

© 2026 Symbeeo. All rights reserved. · Public

6 · Limits

- NCC support for the Google side is in Preview; until it is GA, each extra Google VPC needs its own peering, and peering is not transitive.
- Region pairs are limited to those AWS and Google list (five at GA on April 14, 2026; Google added australia-southeast1 and europe-north2 on July 14, 2026) [S]; a title hosted in another AWS region uses the fallback path or waits for a new pair.
- Oracle Cloud Infrastructure is generally available with AWS Interconnect – multicloud since July 2026, and Microsoft Azure has been in Preview since August 31, 2026 [S]; neither is in this design.
- The AWS Interconnect SLA gives 99.99 % monthly uptime per covered interconnect, with service credits as the sole remedy; it excludes problems beyond AWS's demarcation point and counts unavailability only after 120 consecutive seconds [S]. Google's partner Cross-Cloud Interconnect page says each provider offers its own SLA for its portion [S], and Google's Cloud Interconnect SLA page names Dedicated, Partner and Cross-Site Interconnect only (read Sep 27, 2026), so no Google-side figure is stated here.
- The BGP timers of the provider-managed link are not published; the 90 s target plans on the Direct Connect default hold timer [S] as an assumption [A], not a published bound, and the first failover test measures the real figure.
- The 6 Gbps peak is an assumption; 30 days of VPC flow logs around one launch set the real figure.

7 · Lessons learned

None yet — this is a design study; populated after the first build against it.

8 · The author

WRITTEN BY

Steve Phillips

Principal Data Center Architect

CCSP, CCIE, AWS Certified Security – Specialty

Signed off by Edwin Ozer, Founder.

© 2026 Symbeeo. All rights reserved. · Public

SYMBEEO

Each vendor's tool shows you their option. We show you all of them against one requirement.

Compared. Then Delivered.

hello@symbeeo.com · symbeeo.com

© 2026 Symbeeo. All rights reserved.

Product and company names and logos are their owners' trademarks, shown only to identify the lines Symbeeo sources. Naming a vendor or showing its logo does not mean the vendor sponsors or approves Symbeeo's comparisons or any recommendation Symbeeo makes.

Figures are as published by their sources on the dates given in this document; figures marked [A] are assumptions.

CLD-003 · Version 1.1 · Oct 10, 2026