

Connectivity to Google cloud service for ATLAS

LHCOPN-LHCONE meeting #50 - Prague, Apr 19th 2023

G. Merino, on behalf of the ATLAS Google Project

The ATLAS Google Project

ATLAS ongoing R&D project to demonstrate it can use Google Cloud resources to extend the current distributed computing infrastructure.

- Nov 2019 - Oct 2022, Phase 1: successful prototypes to demonstrate integration of ADC tools (PanDA, Rucio, etc) with GCP.
- Apr 2021 - Mar 2022, Phase 2: successful demonstration of usefulness of GCP as an ATLAS analysis facility

In July 2022 started a “production” phase of the project: running a Cloud site of $\approx 2.5\%$ size in production for an extended period of time of about 1 year. Evaluating a new GCP fixed-price “subscription agreement” model. Main goals are:

- Gain operational experience of running a Cloud site at scale , as close to a "standard" WLCG site as possible
- Evaluate TCO of a Cloud site at scale

ATLAS Google Site parameters

User subscription agreement






- Fixed-cost subscription with ample flexibility
- Duration: 15 months (July 2022 - Sept 2023)
- Compute: average 10k vCPU
- Storage: average 7 PB

Thanks to the experience gathered during previous PoC phases, the site was up & running within days after green light

Running in europe-west1 (Belgium) region

- Physical proximity to CERN and other EU sites
- Large size and many resource types available
- Low carbon emissions

Carbon data across GCP regions

Google Cloud Region	Location	Google CFE%	Grid carbon intensity (gCO ₂ eq/kWh)	Google Cloud net operational GHG emissions
europe-central2	Warsaw	20%	576	0
europe-north1	Finland	91%	127	0  Low CO ₂
europe-southwest1	Madrid	*	121	0  Low CO ₂
europe-west1	Belgium	82%	110	0  Low CO ₂
europe-west2	London	57%	172	0
europe-west3	Frankfurt	60%	269	0
europe-west4	Netherlands	53%	282	0
europe-west6	Zurich	85%	86	0  Low CO ₂
europe-west8	Milan	*	298	0
europe-west9	Paris	*	59	0  Low CO ₂

<https://cloud.google.com/sustainability/region-carbon>

Operational experience




Running any type of production

- Job types dictated by data availability

Storage Element at Google Site initially set up with no deletion for cached data.

- As new files were produced, data was transferred to other sites as job inputs
- Egress traffic was grew quickly to very high values

In December 2022 the configuration was changed to minimize the amount of data cached at the Google Site and hence the egress traffic.

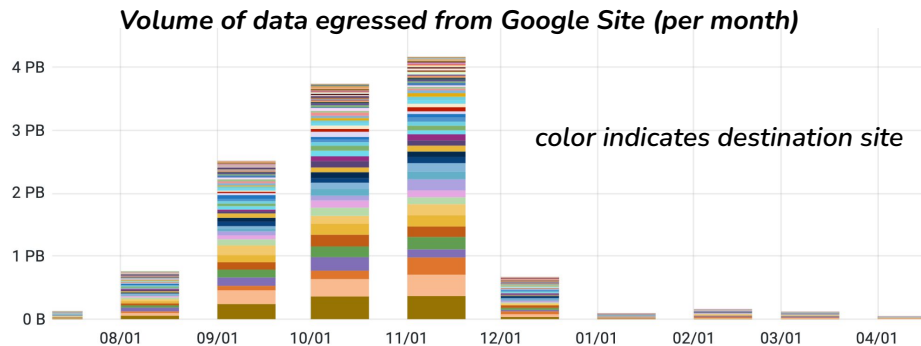


Egress traffic

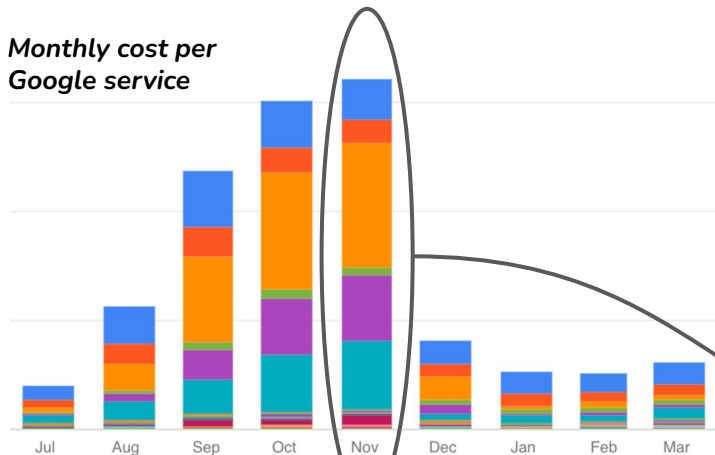
The egress traffic from the Google site reached **4PB/month** in November 2022.

Even if the subscription agreement works on a flat-rate, we can check the “list price” cost of the activity.

In Nov 2022, **~60% of the cost was egress**




Monthly cost per Google service



Google Site cost Nov 2022

- Egress EMEA-EMEA
- Egress EMEA-Americas
- Egress EMEA-APAC
- Compute
- Storage
- Other



Dedicated network links to the Cloud provider

For GCP, the direct connectivity product seems to be “**Cloud Interconnect**”, <https://cloud.google.com/network-connectivity/docs/interconnect>

*Cloud Interconnect extends your on-premises network to Google's network through a highly available, low latency connection. You can use **Dedicated Interconnect** to connect directly to Google or use **Partner Interconnect** to connect to Google through a supported service provider.*

<https://cloud.google.com/network-connectivity/docs/interconnect/pricing>

Dedicated Interconnect pricing	
Resource	Price
Interconnect connection	\$2.328 per hour per 10-Gbps circuit
Interconnect connection	\$18.05 per hour per 100-Gbps circuit
A 50-, 100-, 200-, 300-, 400-, or 500-Mbps VLAN attachment	\$0.10 per hour per VLAN attachment
A 1-, 2-, 5-, or 10-Gbps VLAN attachment	\$0.10 per hour per VLAN attachment
A 20-Gbps VLAN attachment	\$0.20 per hour per VLAN attachment
A 50-Gbps VLAN attachment	\$0.50 per hour per VLAN attachment

Partner Interconnect pricing	
Partner VLAN attachment capacity	Price
50 Mbps	\$0.05417 per hour per VLAN attachment
100 Mbps	\$0.0625 per hour per VLAN attachment
200 Mbps	\$0.08333 per hour per VLAN attachment
300 Mbps	\$0.1111 per hour per VLAN attachment
400 Mbps	\$0.1389 per hour per VLAN attachment
500 Mbps	\$0.1736 per hour per VLAN attachment
1 Gbps	\$0.2778 per hour per VLAN attachment
2 Gbps	\$0.5694 per hour per VLAN attachment
5 Gbps	\$1.25 per hour per VLAN attachment
10 Gbps	\$2.36 per hour per VLAN attachment
20 Gbps	\$3.61 per hour per VLAN attachment
50 Gbps	\$9.02 per hour per VLAN attachment

Egress pricing

The cost per GB for the egress traffic is substantially lower on a dedicated interconnect connection than on the internet.

For traffic of the order 3-4 PB/month the potential savings are large. For instance, using a 10Gbps dedicated interconnect to transfer 3PB in one month should be about 50% cheaper than using standard internet (~\$65k vs. ~\$145k per month)

<https://cloud.google.com/network-tiers/pricing>

Egress pricing is per GiB delivered. Egress pricing is based on the source region of the traffic. Ingress pricing is still free. [Always Free usage limits](#) do not apply to Standard Tier.

Belgium (europe-west1) ▾	
Monthly usage:	Price per GB delivered:
0-10 TB/month	\$0.085
10-150 TB/month	\$0.065
150-500 TB/month	\$0.045

<https://cloud.google.com/network-connectivity/docs/interconnect/pricing>

Egress traffic from a VPC network through an Interconnect connection	
Interconnect connection location	Price
Asia	\$0.042 per GB
Europe	\$0.02 per GB
North America	\$0.02 per GB
South America	\$0.08 per GB
Australia	\$0.042 per GB

Questions

Given the ATLAS computing model, sites of size >10k cpu cores export data to other Grid sites at rates 1-10 PB/month.

For running a site of this size on the Cloud, it seems interesting to explore options to deploy **dedicated network links**.

Possible advantages are:

- Potential cost reduction.
 - Data leaving a Cloud region over regular routes is charged at a much higher rate than data routed over dedicated links.
- Option to route the traffic through some path which is more optimal for destination sites.
 - We heard complaints about too high traffic arriving to ATLAS Grid sites through standard internet routes (non-LHCONE/LHCOPN)

Due to this, we would like to ask this forum about which technical options are available to test some of these setups

- Are there mechanisms in place to test “Cloud Interconnect” dedicated links with Google via some of the academic providers such as GEANT, Internet2, ESNET, or others in the short term? (so that we can fit it in our current ATLAS Google Project)
- Would it be feasible to test routing the traffic going out of a Google Cloud site into the LHCONE so that it would reach other Grid sites as if it was another trusted Grid site?

Thank you