LuftBlick Report 20240828

Fiducial Reference Measurements for Air Quality

Data Processing Evolution Report

Version 2 - Aug 28, 2024

Prepared by

Name	Institution
Daniel Santana	LuftBlick
Michaela Nickel	LuftBlick

Contents

1 Document Change Record

Issue	Date	Section	Update
1	31 August 2023	All	First version
2	28 August 2024	Contents	Added and renamed sections
		3 Infrastructure Requirements Review	Updated “Processing Load” figure.
		4 Traditional Infrastructure	Updated to current situation
		7 Migration to centralized server	New section
		8 Summary and Future Plans	Updated to current situation. Added "PGN IT infrastructure" figure.

1.1 Acronyms and Abbreviations

ESA: European Space Agency
FRM4AQ-2: Fiducial Reference Measurements for Air Quality 2
GCP: Google Cloud Platform
OVH: OVHCloud.com
PGN: Pandonia Global Network
STRATO: Strato server hosting

2 Introduction

This report represents deliverable D11 of the ESA project FRM4AQ-2. It describes the evolution, current state and future plans for the server infrastructure which powers the PGN. In the proposal of the FRM4AQ-2 project we have laid out basic requirements for the future infrastructure, and in the following sections we will elaborate on these requirements, discuss challenges and potential solutions, and finally present our current architecture and path forward.

3 Infrastructure Requirements Review

Since the birth of the PGN the demands for computational power and data storage have risen dramatically. Naturally this was to be expected due to the inherent scaling nature of the network. Virtually every goal that the PGN has implies an increased demand in computation and storage which has to be taken into account. The most important scaling factors are:

- **Time:** Ideally an instrument will continue measuring data for a long period of time and thus accumulate more and more data which needs to be stored and (re-)processed.
- **Number of Instruments**: Each instrument added to the network adds a new source of raw data and is the most obvious dimension of scaling next to time itself.
- **Number of Calibrations:** New calibrations require a reprocessing of data and present a large portion of our processing workloads.
- **New Data Products and Algorithms:** Similarly to new calibrations the introduction of a new data product or processing algorithm version requires a reprocessing of data, but instead of a single time series this can affect large parts or even all of the processed data.
- **Number of Users:** When our data is visualized, queried and downloaded more often this also adds to the demand on our infrastructure.

Looking at these scaling factors and how the network works in general we can group the infrastructure requirements into 3 distinct pillars:

Storage: Our ever growing record of raw and processed data needs to be stored in an organized and scalable way which can serve the data for analysis, visualization and distribution. Storing data in structured databases is strongly desired over simple text file formats in order to facilitate large scale data analysis, improve data consistency and maintain traceability.

Baseline Computation: The rather steady daily influx of data and automated data analysis related to the operational upkeep of the network form a constant baseline of computation required during a typical day of operating the network. This workload is mostly predictable and is well suited to be run on a permanently available server.

On-Demand Computation: This is the complete opposite of the constant baseline demand. Every workload that is concerned with batch data reprocessing (e.g. due to new calibrations, data products, algorithms) falls in this category. Suddenly large amounts of data need to be processed or analyzed, and therefore resources need to be allocated. Often these workloads can be parallelized (e.g. processing many days of data simultaneously) which is hugely beneficial from the perspective of the calibration staff but further adds to the heterogeneity (meaning spikes in demand) of the workload.

The figure “Processing Load” illustrates the difference between the constant operational baseline demand (green, labeled “operational”) and the spiky on-demand workloads (yellow, labeled “cloud”) representing batch reprocessing of data. Notably the on-demand batch processing is already requiring about 3x as many resources as the daily operational processing, and this ratio will shift even more the more raw data we accumulate and have to reprocess eventually. Additionally the figure implies that we utilize cloud computing for our on-demand workloads which we will describe in section 5.

Figure “Processing Load”

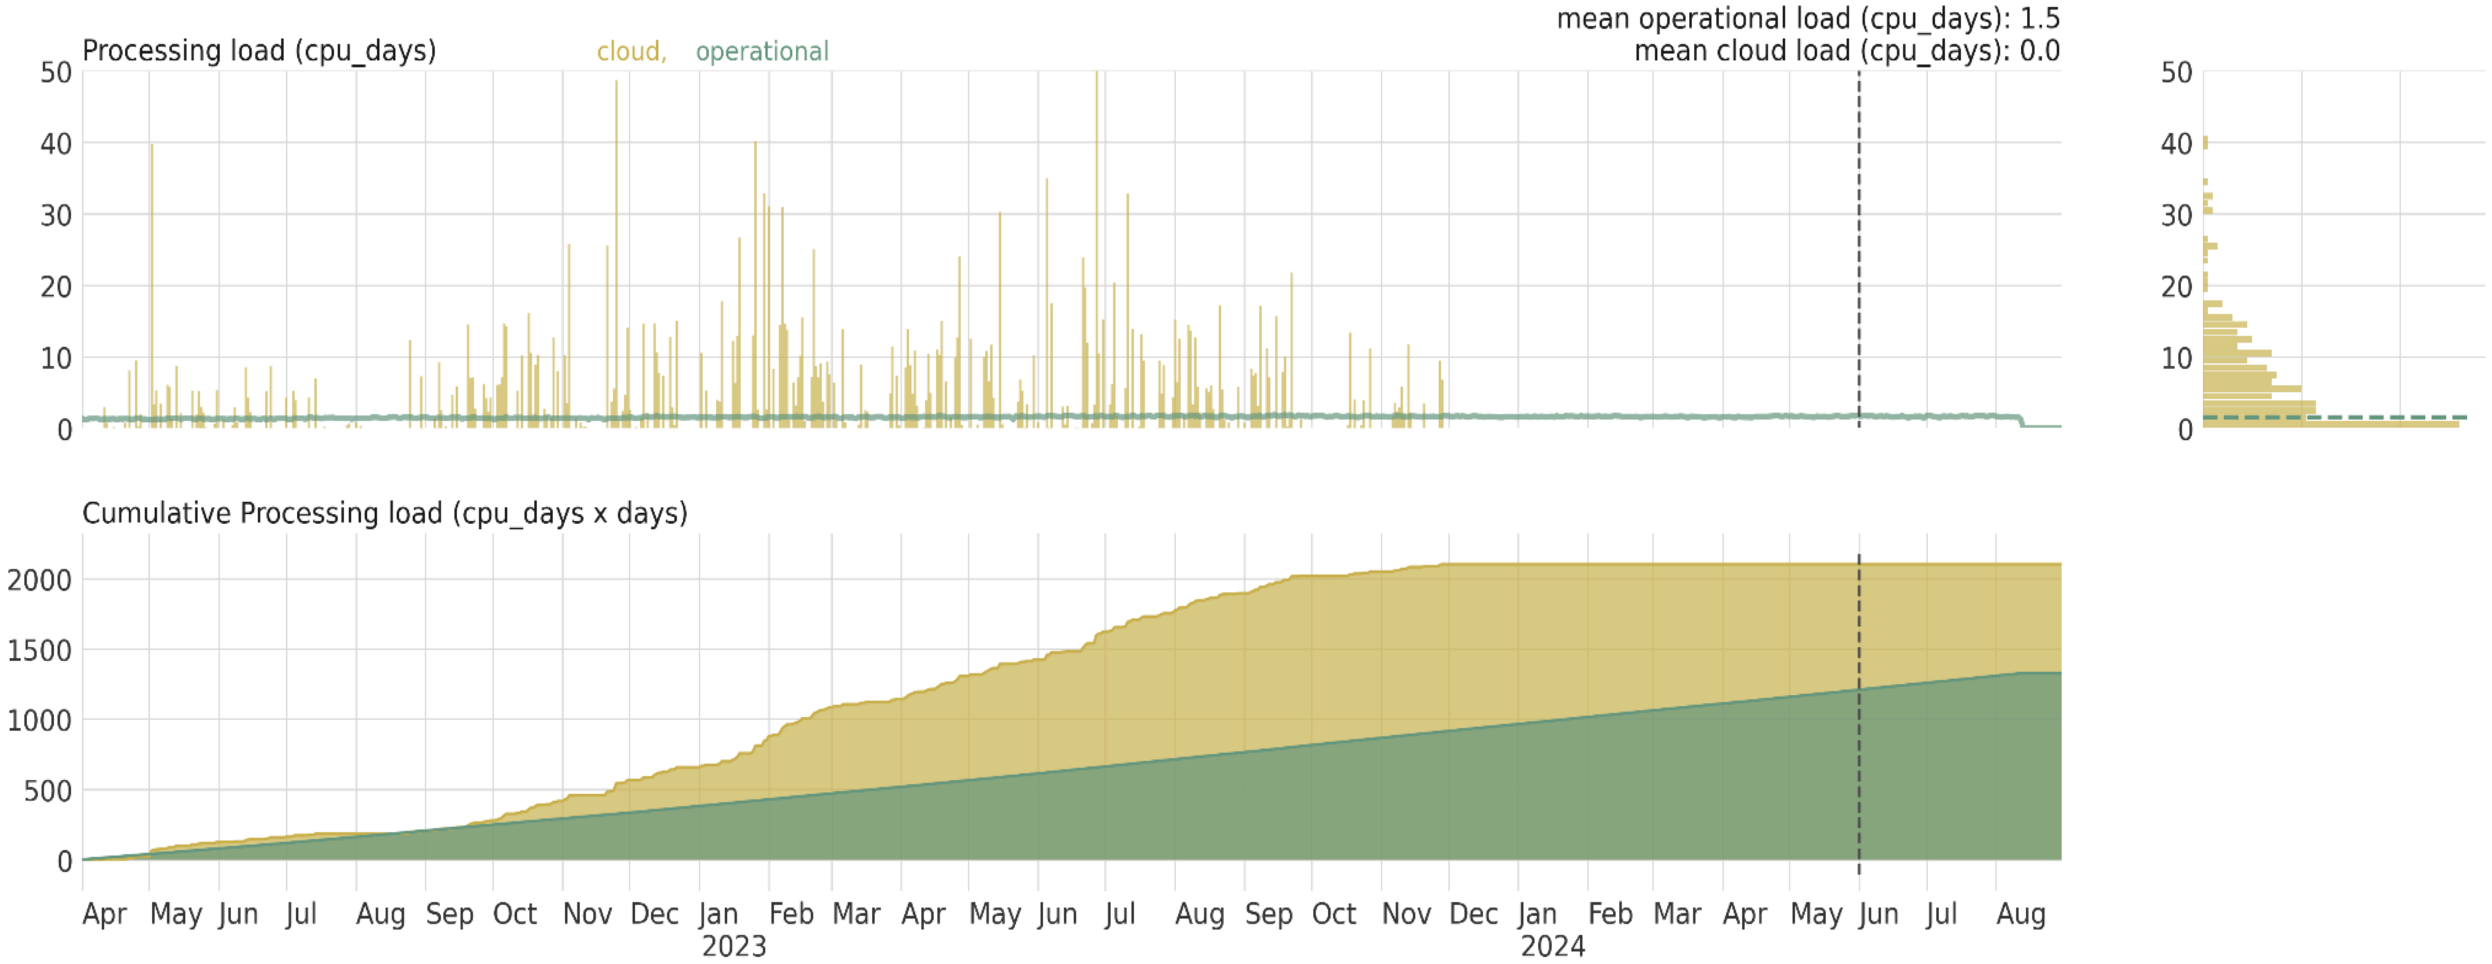

4 Traditional Infrastructure

Since the beginning of the PGN LuftBlick has been renting dedicated servers from the German provider STRATO. Until the end of 2023 we were working with 3 medium sized machines (lb3, lb4 and lb5) and an additional storage volume (hidrive). This setup provided us with 20 TB of total storage, and hosted most of our services like BlickV and BlickM, and dealt with the operational baseline processing, in the past.

For many years this infrastructure did its job, but in recent years we faced a number of challenges related to it. First of all, the distributed nature of our server landscape, which has evolved by subsequently adding server machines, increases the complexity and maintenance effort required to orchestrate the machines together. In general it is easier and more reliable to maintain one central storage server than many smaller ones connected together. Especially since STRATO was repeatedly throttling down our network throughput and in short simply not delivering the services as promised. And most importantly, as the figure "Storage on STRATO" shows that we were reaching the limit of our 20 TB capacity and we would need to add more servers or storage devices to the STRATO cluster soon.

In the beginning of 2024 we made the move to one large dedicated server that we are renting from the provider OVH which eliminated a lot of the shortcomings of our previous system.

Figure "Storage on STRATO"

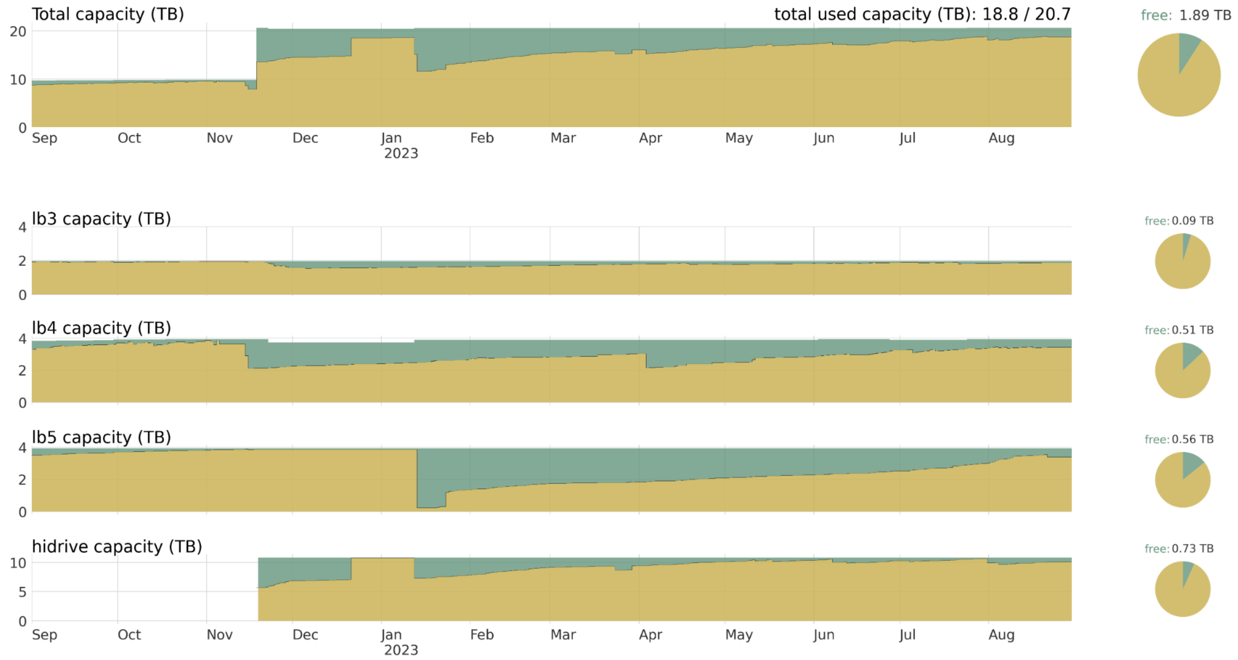

5 Cloud Integration

In 2021 we started integrating services from the Google Cloud Platform (GCP) to satisfy the infrastructure demands of the PGN. Renting resources from cloud providers is generally more expensive than renting dedicated servers from other providers, but the huge advantage is high levels of abstraction and on-demand scaling. Concretely this means we utilize serverless offerings like Google Cloud Functions and Google Cloud Run with which we orchestrate the parallel execution of containerized jobs. These jobs grab raw data from our storage infrastructure, process it and then the resulting data is collected and integrated into the storage layer. The two big advantages of this approach are that processing jobs can be finished fast due to the high parallelism, and we don't have to pay and maintain large amounts of infrastructure which would be idle most of the time.

The only disadvantage of this approach is that data has to be moved between our dedicated servers and the cloud. That is not a big problem, but all big cloud providers charge for transferring data out of the cloud, which is called network egress. This is due to the fact that the major cloud providers are incentivizing their customers to keep their data within a cloud provider's internal data network. Hence the question, why not move the whole infrastructure into one cloud provider?

There are some important reasons against it: First, for the baseline resources which are needed every day such as storage, databases and constant computation loads a dedicated server is significantly cheaper. Second, when serving the data for download, visualization and analysis, egress fees will be applied, which is hard to predict and especially contrarian to our company goals as we want to avoid having to restrict the distribution of data for cost reasons. And third, committing fully to one vendor increases the so-called "lock-in" and increases the dependance on a single company.

6 Hybrid Cloud Architecture

Combining dedicated servers, rented or on-premise, with modern cloud services is what is called a "Hybrid Cloud" architecture. This approach is very common in order to get the best of both worlds in terms of scalability and cost-effectiveness.

We already described why using the cloud for spiky on-demand workloads is a perfect fit. What is less obvious is why a large dedicated storage server is for our use case more effective than a full cloud setup. The reason is that the big cloud providers usually offer easy-to-use services for small-medium data, and then very elaborate services for scaling up databases which are accessed by millions of users daily. Our scientific access pattern falls a bit out of line here: We have considerable amounts of data but only few users. Still, the relatively small number of users routinely needs to look at a lot of historical data, which is exactly the opposite of i.e. millions of social media users looking up their current feed.

This means that there is no appropriate plug-and-play service in the cloud which could host our storage backend in a serverless manner. As an example, a managed CloudSQL instance in GCP cannot scale big enough for our needs, while the other managed services like Cloud Spanner or BigTable could handle our data easily but at the same time cost way too much because they are meant to service large amounts of users.

Therefore we have to implement our storage and database backend ourselves anyways, and hosting it in the cloud just introduces more cost and egress fees. Consequently as of 28.08.2023 we have rented a large storage server (200+ TB disk space, 196GB RAM) from the provider OVHCloud.com.

7 Migration to centralized server

Since August 2023, our team has been diligently working on transitioning all services and data from the legacy lbX servers (STRATO) to a new, centralized server infrastructure (OVH). This migration significantly enhances our ability to manage and maintain services while improving data accessibility across the board.

In conjunction with this transition, we've leveraged the opportunity to reorganize our services, incorporating more efficient and scalable solutions, and adopting the latest technological advancements:

1. **Containerized Applications**: A majority of our services have been dockerized. This shift ensures that each application operates in its isolated environment, making them more robust, easier to deploy, and simpler to maintain. This change also paves the way for greater scalability and flexibility in our operations.
2. **Advanced Orchestration Tool**: We have upgraded our primary orchestration tool from Dagster to Temporal IO. This new tool offers enhanced capabilities, allowing us to process multiple days of data concurrently and handle on-demand tasks efficiently. It also enables us to manage server load internally, eliminating the need for outsourcing the processing to the cloud, though we retain the flexibility to do so if necessary, thus optimizing costs without sacrificing performance.
3. **Enhanced Data Accessibility**: We've introduced new applications aimed at improving data accessibility for users. Notably, the addition of FileBrowser, an online tool similar to Google Drive, allows users to easily view and download processed data, thereby improving the user experience and expanding access to our resources.

As of May 2024, the migration of data and services has been successfully completed, and all PGN services are now operational on the new centralized OVH server, ensuring smoother and more reliable service delivery. Out of a total capacity of 200TB, the OVH server currently requires only 38TB to store the (current) entire PGN data fleet (across all levels). With the current data growth rate of 12GB per day (including both raw and processed data), the existing infrastructure is projected to accommodate this growth without requiring intervention for approximately 13.500 days, or 36 years.

In the figure "Storage on STRATO vs OVH" the OVH server displays only 100TB out of the available 200TB. This amount represents what we have allocated for our system so far. We can expand this to up to 200TB if needed.

Figure "Storage on STRATO vs OVH"

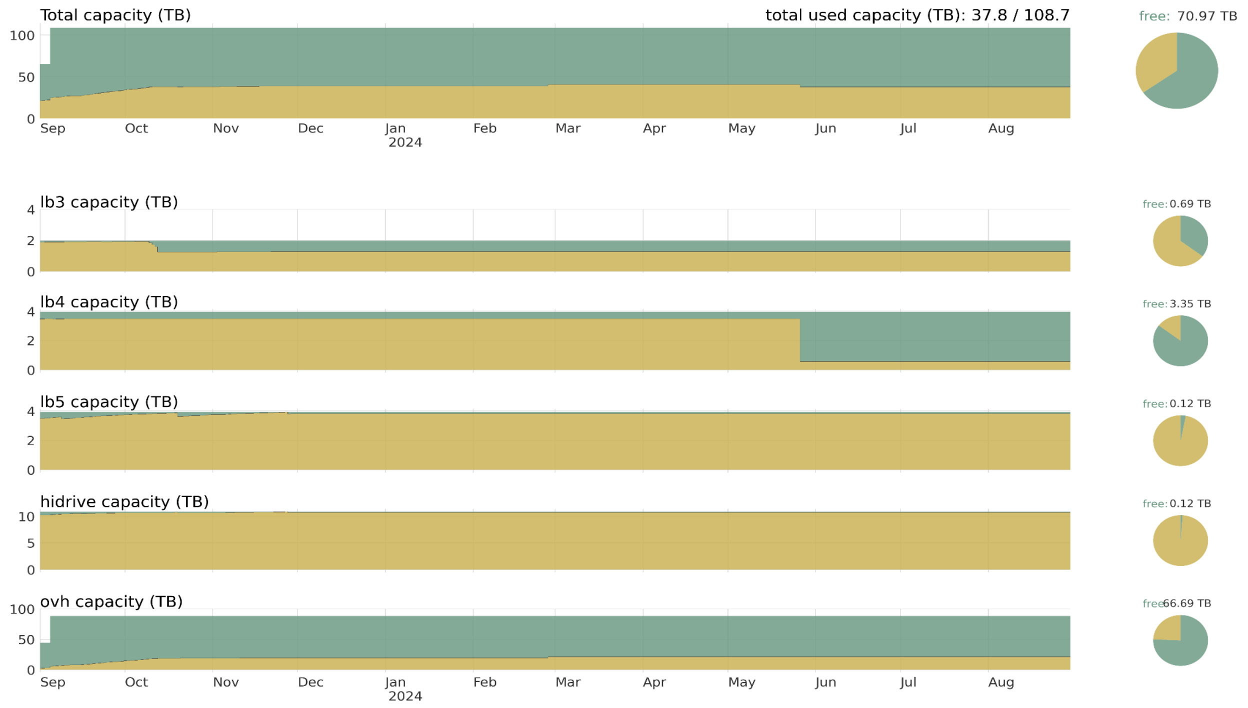

Figure “PGN Architecture Overview”

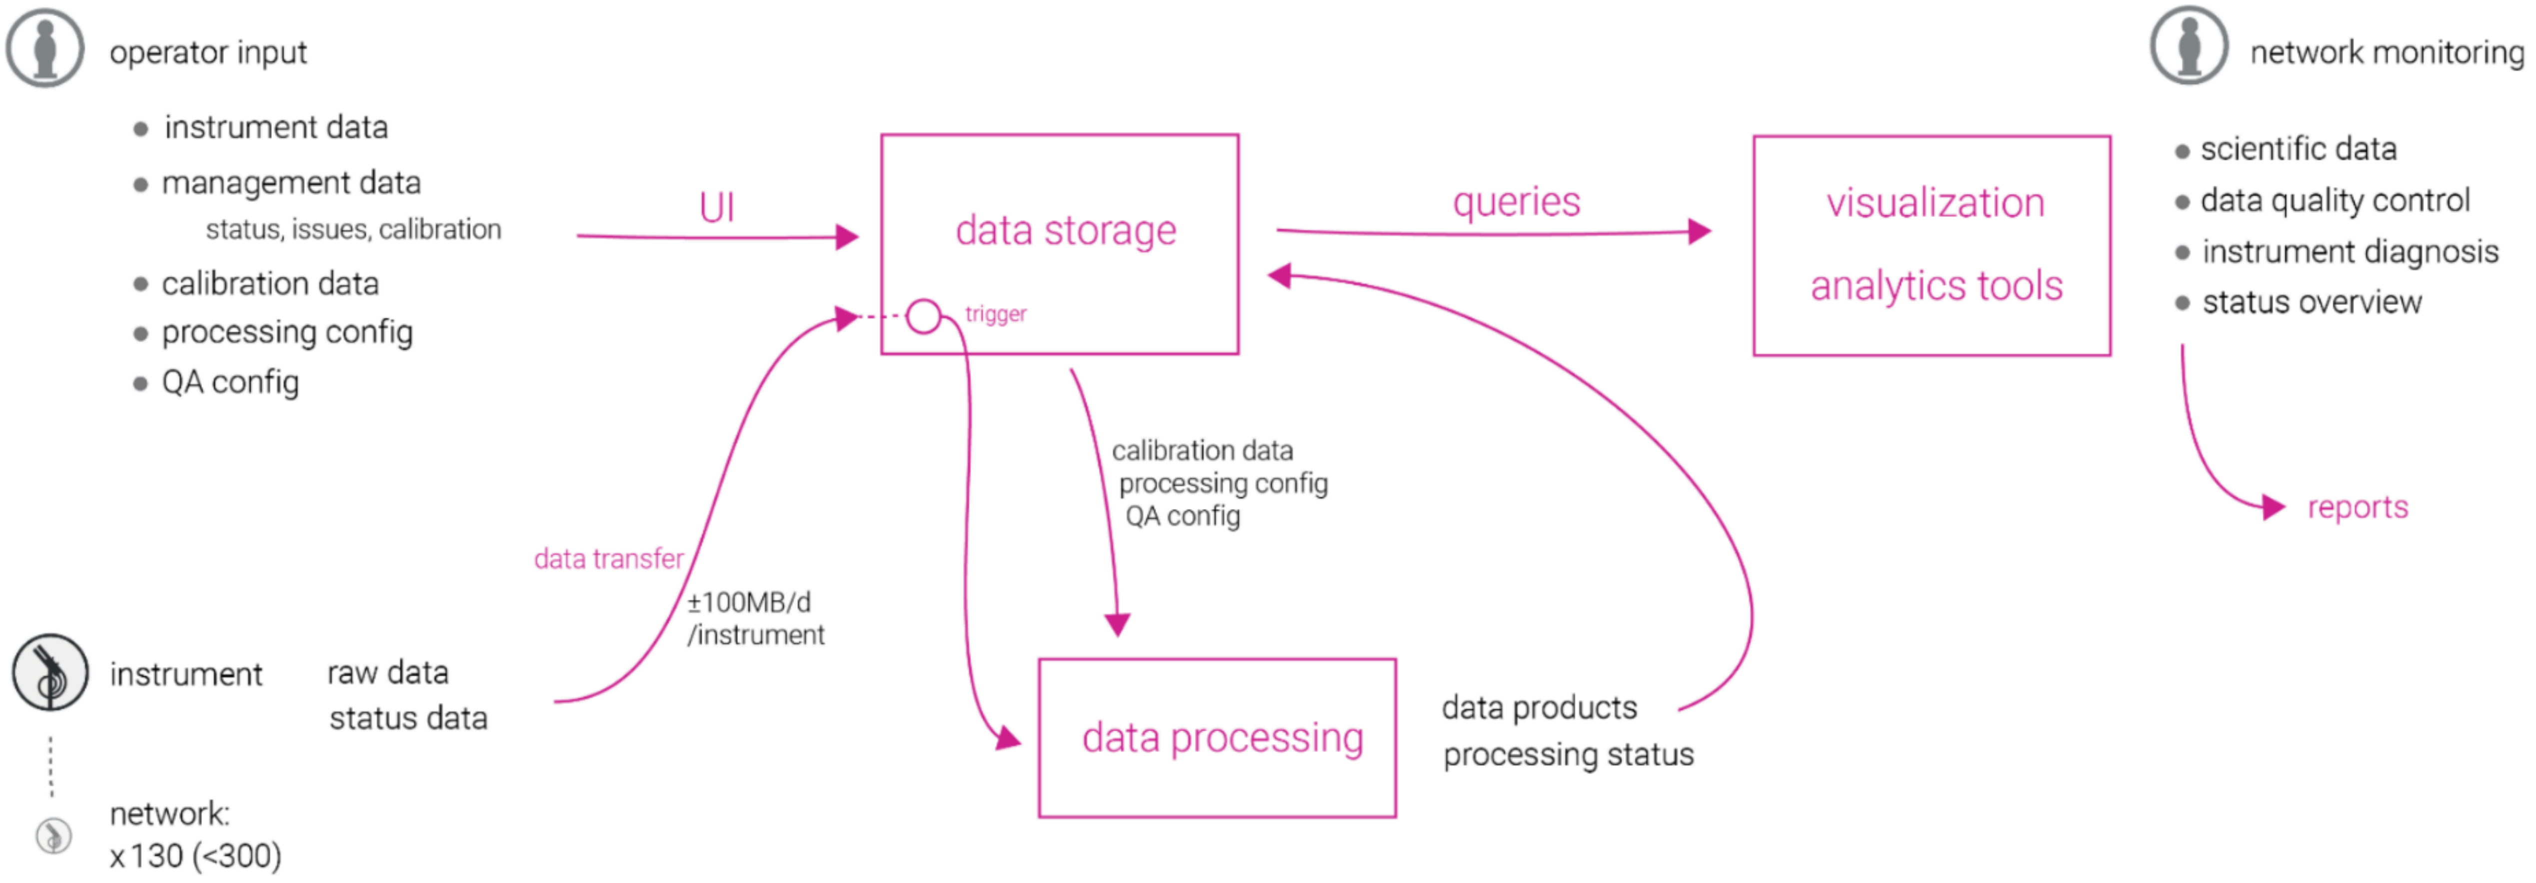

8 Summary and Future Plans

The “PGN Architecture Overview” figure provides an abstract schematic of PGN’s architecture, independent of the specific infrastructure on which it is deployed. Currently, all PGN services are centralized on our OVH server, including:

- Data storage
- Data Processing
- Web applications like BlickM (for instrument monitoring and operator input data) and BlickV (for data visualization)
- Internal visualization, analytics, and reporting tools

Regarding data processing, we still retain the option to outsource to Google Cloud for large or urgent tasks, thereby maintaining a Hybrid Cloud Architecture. Figure "PGN IT Infrastructure" displays our current situation. We have terminated our contract with Strato for all machines except LB 4, which we will retain as a testing and development server to ensure that new features function properly before they are deployed to production. Due to contract terms the contract for the other machines will end in spring 2025.

Centralizing all the above mentioned services and tools on a single server optimizes data access across different applications, reduces network traffic, and simplifies server maintenance.

However, the migration to the centralized server has introduced a few challenges:

- **Real-Time Data Display**: We lost the ability to display near real-time data in BlickV. Currently, it temporarily shows daily data instead of updates every 10 minutes.
- **Live Data Monitoring**: We also lost the live data monitoring functionality in BlickM, which indicates when the last data from each instrument was received. This feature is particularly valuable for Network Operators to detect when an instrument stops producing data.

We are prioritizing the restoration of these functionalities as quickly as possible.

Looking ahead, our future plans will focus on optimizing the processing chain and developing tools to facilitate the download of large data series. We also plan to automate server deployments, using LB4 as a testing platform.

Figure "PGN IT Infrastructure"