Topic 1c - Part 1: Supporting enterprise and innovative solutions - Data processing and supercomputers

This building is the European centre for medium range weather forecast in Reading. And I'm here because it hosts an organization called Copernicus Atmosphere Monitoring Service. And that does what it says on the tin. It collects all kinds of atmospheric data from satellites and in-situ data, but that's a lot of data. And so in this building, there are plenty of experts who are building tools to help individuals, businesses, policymakers, and researchers make sense of that data and make use of it.

Open data is one of the key principles of Copernicus and of Cams. In creating Copernicus 20 years ago, there was a realization that Europe needed a program to observe the earth-- it was just after Kyoto-- to see whether the international treaties and the decisions were being effective.

But also that the information generated by knowing more about the environment, air quality, radiation, et cetera could generate an economy. And for this economy to grow, it was decided very early that the data should be fully open. And one of the key success criteria for Copernicus is the uptake of the data in particular by the private sector and the jobs created on this economy.

We're talking about these vast amounts of data as though they're easy to process and store and analyze, but you must need a huge amount of computing power to do this job. What's your computer capability here?

Yes, here at ECMWF, we have the largest database for weather information, but also one of the largest computers in Europe, which we use for weather forecasts but also for the CAMS forecast. It is among the 20 largest computer in the world, and the few largest in Europe. And this is required because we process so many observations. We have to represent the earth, the entire earth, with a very fine mesh, so we really need supercomputing to do that.

Was all that data stored on site, or do you have ways of sharing that data in the cloud, for example.

Our computer delivers the data locally, so we get these-- we have this massive archive of data. And currently the way our users work is that they come-- they get the data, and they download the vast amount of data. And we are arriving at the limit of the systems-- that's why we are embracing cloud computing technology. The idea is rather than sending our data to the users, we facilitate the fact that the users can come to our data and do their own processing. And it's a very important development in Copernicus to have this capability so that the users will be able to tap in these vast amount of data without having more infrastructure than a laptop.

It sounds incredible having access to a planets worth of data just from your laptop, but that is in the future.

Yes. Yes, that's exactly what we want to do, because we realize that there is a limit to the data that can be transferred. And there is a limit in developing data center and the facility foreseeing our number of users increase further. In fact, we've inherited lots of our procedures and infrastructure from numerical weather prediction, from weather forecasts. What we do is to acquire a very large number of data streams.

Typically we are speaking of 70 different satellite instruments. Several hundreds or thousands of ground based measurements. We acquired that. We want to do forecast, so we need to do that within a few hours only. Then we have our modeling procedures that can blend these observations with what we estimate from a previous forecast with the state of the atmosphere.

We blend that. We have a new accurate picture for the starting point of the forecast, and we do the forecast. Which means that a user will not have to do all this work of acquiring all the observations that come in different formats, different time, et cetera. What they get is numerical fields, three dimensional with a time dimension. And then they can pick maybe just the time series in the very place that they are interested. So in a way, we are preparing the work for end users so that this is done once for good, for Europe at least, and this is very cost saving for us.

So here we are at the ECMWF console room, which is where we monitor and operate the forecasts every day. Here we have Antonio, who is one of our console operators, and can tell you about the system.

Hey. Pleased to meet you.

Nice thing to meet you.

So you are in control of all of this. And this is-- what does this-- what is being controlled here?

Well there is an awful lot. Should I show you around the-- if you follow me here, this is actually a summary of our main duties. Because the way it is displayed, this is the things that we be aware. If anything happens and it requires an urgent action, they should be displayed here. So for instance, here we have ECFS, which is one of the applications where we store lots of information that are used by the jobs that we use, or by users. And this displays the movers. So at a glance, we can actually see very quickly if there is anything that requires immediate action.

So your role here is that there's lots of information coming into this point, into this room, and you're in charge of making sure it goes out to the appropriate places.

That's a reasonable summary. That's a reasonable summary. But there is also a lot of equipment that we monitor. Because all those flows of information can only be smooth if everything, all the little bits and pieces, all the hardware is actually functioning. So also for instance, this tool here allows you, through callers, to see if there are issues. For instance, we've got some tape libraries in the computer room.

Proper tape, even today,

Yes. Yes. They are like robotized, so there are robots. So are much more advanced than when I started here where we had to do it manually. So now robots will pick up the tape and mount it in the drive.

And how much data comes through this console room in a day?

Well not right to our console, but there are certainly terabytes of data going every day in and out. Depending on the jobs, It can vary a lot of course from one day to the other, but definitely a lot of data goes through. And it depends on what application you're talking about. It depends if you're talking about leaving the Weather Center and coming here, or just if you're talking about a local thing. So it depends on what you are talking about.

But if you're talking about over months, we could talk about petabytes for sure. But we also get lots of satellite data through the GTS. We get data from all over the world. Every country in the world, when the World Meteorological Organization will feed information. And we'll get some of that information, which then we use some of it in our model to produce forecasts. And then we disseminate that information to member states and other organizations interested to know about it.

It's fine having this data, but you're also monitoring here some of the processing of it. These plots here, tell me about those.

Yeah. Well those plots are very useful, because they give us a very good idea of our supercomputers are working. As you can imagine, supercomputing is a very expensive commodity. So we tend to use our supercomputers as well as we can. So we have loads of jobs running almost at the same time. There are queues, and we try to prioritize, for instance, our models for operational purposes-- get the priority. And then research work, we'll have less priority.

So that means that the parts of the model work, there has to be an output every day at the same time. Those get priority.

That's correct.

And then the scientists who have an experiment today, maybe a different experiment tomorrow, they get the gaps in between.

That's exactly-- you just summarized it all very well, yes.

So all of this-- so how much-- you've got two computers, two supercomputers.

Yeah. We've got two classes, correct.

And so these are all the different tasks that are going on, all the different colors.

You've got a here-- it's useful to compare it with the previous day. So if there is a variation, often it tells us that something is not quite right. So the variation should not be too great, unless we know about it, of course.

So it's very impressive that all of the things your supercomputer must be doing can be condensed onto the two plots here. And the supercomputers just next door, is it?

Yes. Just in the machine room. We've got two machine rooms in case there are crises. So they are in two different rooms, so if there is a problem-- fire, flood, whatever in one room, we still have another cluster working, and vise versa.

So you've got a backup built in.

Some sort of redundancy mechanism, if you know what I mean.

OK, can I go and have a look?

Yes, of course.

So we're here now at the ECMWF computer hall. This is where we keep our supercomputers, and this is where we run our forecasts every day. So this is one of our supercomputers here. We have two of these Cray XE40 machines, and each one is 20 cabinets. And each cabinet has about 7,000 processors, and these are the supercomputers which do our forecasting, which is 10 to 15 days for the medium range weather forecast. And for the CAMS, atmospheric composition forecasts, those are two five day forecasts which run everyday at midnight and mid-day UTC.

So what's inside here can basically simulate world weather for 15 days.

Exactly. So each of these cabinets contains all the computing power that we need for computing, for calculating the weather, and forecasts for 10 to 15 days, and even seasonal scale at about 10 kilometers, 9 kilometers, spatial resolution globally with hourly time resolution in the outputs.

And you can tell it's working causes these are cooling fans we can hear. All this noise is cooling.

So the infrastructure in here is very loud, because there's a lot of infrastructure for cooling, keeping the supercomputers cool, and keeping the power consumption is quite high as well.

Right. And then how about the data that goes with all of this-- where does that get stored.

So the data is all stored in archives, which is easily accessible. And there's also a tape archive which stores many, many years worth of mythological and atmospheric composition data, and they're kept through the back end.

So a tape archive.

Tape archive, which is magnetic tape, yes. So here you can see, these are some of the magnetic tapes, which are currently not in the archive, but this is where they are brought sometimes to be stored-- And here is-- this is the tape archive.

So how does this work? You've got little robots.

Yeah. So we have all the tapes are stored in here. And when a request for a particular data set is made, the robot will go and find the tape that the data is stored on. Will bring it here to a staging area so that then you can download it to your computer.

So you're physically taking out magnetic tapes from the pile and then putting them in that thing.

That's correct.

And that's getting copied off. And then they put the tape back.

And then they'll go and put the tape back where it belongs. Currently we have about 15,000 user requests per day, and that amounts to about 200 terabytes of data being interchanged in and out of the tape archive here.

And' why would you do this on magnetic tape instead of having a hard drive?

So we do it on tape for the cost. It's much cheaper to store things on the magnetic tape. Computer storage is a bit more expensive, and this is a more secure and longer term means of archiving the data.

EUMETSAT

And it's definitely a better filing system the mine at home.

imperative space