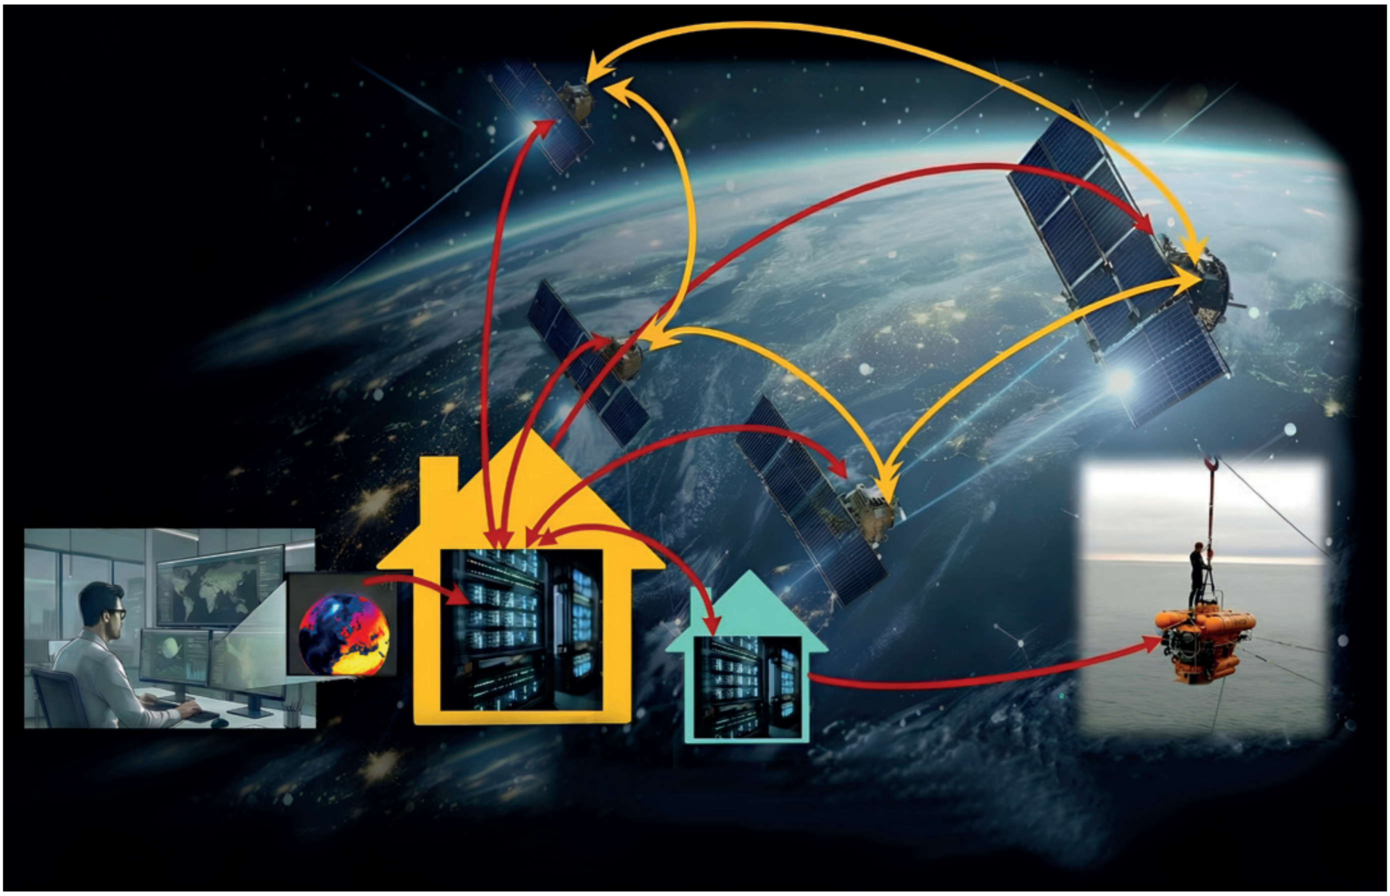

source: rasdaman

Distributed data fusion results in a unified, homogenized information space.

Here Comes a Raspberry Flying …

The SkyFed project demonstrates how datacube analytics is transforming Earth observation by combining distributed data fusion and AI, as shown using 64 Raspberry Pi computers.

author: Heike Hoenig

Millions of gigabytes are pouring in. Day after day, Earth observation satellites deliver new information: images, time series, weather and climate data, and elevation models with ever-higher resolution. What once seemed like a sign of progress has long since become an infrastructure problem. Archives are growing faster than traditional analysis processes can cope with. Those who want to gain timely insights can no longer simply collect data; they must analyse it where it is generated.

From data streams to data intelligence

This is precisely the area in which SkyFed steps in: This project, which is co-funded by the EU's ERDF and the state of Bremen, brings together three partners: rasdaman GmbH, Constructor University and the University of Applied Sciences Bremen. Rather than sending vast amounts of raw data back to Earth, the satellites themselves will be able to answer questions, understand space and time, combine data sources, integrate into federated systems and perform AI-supported conclusions directly on board. The result: no more uncontrolled data streams, but rather targeted information, just a few kilobytes in size.

SkyFed: A satellite swarm federation

This vision is called the "Location-Transparent Earth/Sky Datacube Federation." It is an integrated data space where data centers, cloud systems, drones, aircraft, and satellites function as a single information unit.

The foundation is rasdaman, a mature big data technology that is already being used worldwide to federate geodata as part of the EarthServer initiative.

The term “federation” plays a central role here. In a federated system, many distributed nodes—in this case, 64 Raspberry Pis—work together as a unified whole. Each node contains a portion of the data. When a request is made, the distributed data fusion process ensures that each node provides the necessary data to fulfill the request. Users do not need to know which part of the data is stored on which device—the system automates distribution and load balancing. This provides users with a unified view of the data, even though many components are involved in the background. It is precisely this principle—location transparency—that makes the demonstrator so significant.

A team under pressure

The schedule was tight. There were only a few weeks between the project launch and the planned presentation at the Space Tech Expo in Bremen. Building and configuring the hardware was a gruelling task. The rasdaman server had to be ported and optimized for the Raspberry Pi platform. Configuring the federation and visualizing the datacubes were the final steps. Although each employee had a clearly defined area of responsibility, everything depended on everything else: If the power supply was unstable, even the best software was useless. Without a functioning federation, the cluster remained nothing more than a collection of individual circuit boards.

Logistics as a Challenge: Hardware meets Deadline

It all began with logistics. Sixty-four of the latest generation of Raspberry Pis had to be procured at short notice, along with storage media, power supplies, USB cables, mounting materials, Wi-Fi infrastructure and a custom-made acrylic enclosure. Planning was particularly challenging for this enclosure in particular as it was specially manufactured and, due to cost and time constraints, there was effectively only one attempt. Spacing, drill holes, cable routing, air circulation and stability — with 64 computers, improvisation quickly becomes costly. Then the first parts arrived. Each Raspberry Pi required a microSD card which had to be set up, labeled and tested. The individual devices were first assembled into towers of eight — small vertical stacks that were later placed together into the acrylic enclosure. With the arrival of the enclosure, the eight towers became an installation. The transparent construction made visible what often remains hidden in computing infrastructures: the individual. nodes, the order of the layers, the repetition of a principle. Sixty-four small computers, each unimpressive on its own, together formed an object that immediately attracted attention when placed together. At the same time, a mounting plate was created for the electrical infrastructure, which was hidden beneath the presentation table but served as the demonstrator's foundation. This housed the power distribution system, power supplies and the Wi-Fi router. Communication between the Raspberry Pis took place entirely via Wi-Fi; the visible USB cables were solely for power supply. This decision was not just pragmatic: a satellite swarm is not connected via

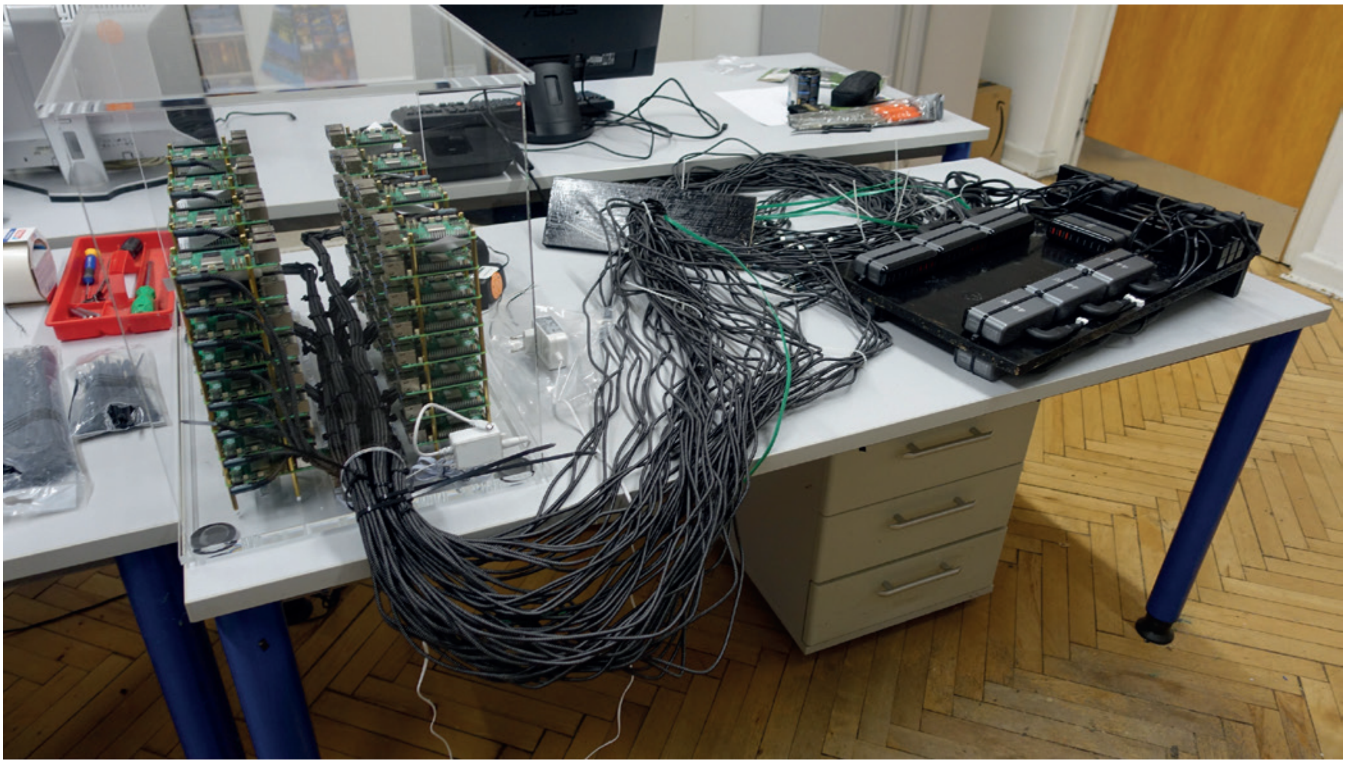

The Raspberry cluster is powered by 64 separate cables connected to 13 power supplies, each with 5 outputs.

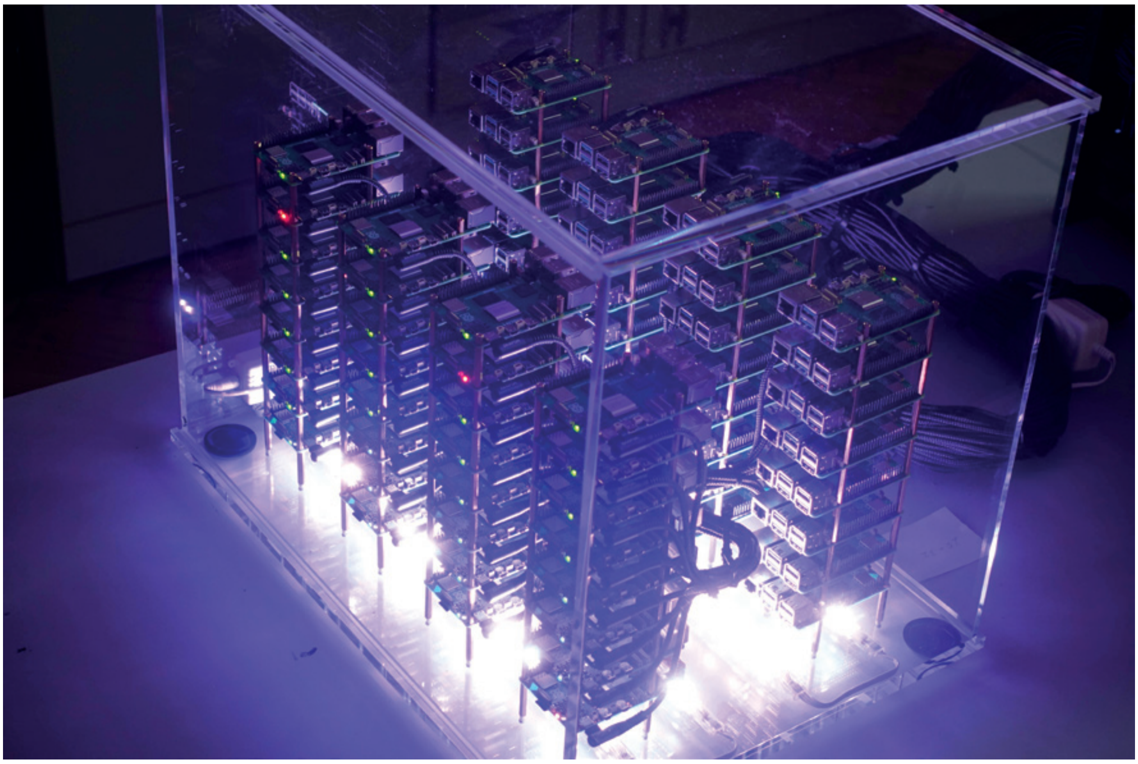

The finished Raspberry cluster demonstrates how edge devices can be federated.

Ethernet cables. The cluster was designed to reflect this principle.

On the software side, the task was to create a homogenized data space across 64 computers. An ERA5 climate datacube was distributed across the nodes and combined into a single virtual cube. This allowed requests in the geo datacube query language WCPS (Web Coverage Processing Service) to be processed in a distributed manner—highly dynamically, without any programming effort on the user's part. In a subsequent step, existing archives were integrated: the CoperniCUBE datacube service and the aviation weather service from the NATO SPS project Cube4EnvSec. This meant that the cluster was not an isolated showpiece, but part of a larger data ecosystem.

Then came the big day: Space Tech Expo Europe in Bremen, announced as "The Largest B2B Space Event in Europe." At the booth, the cluster became an eye-catcher. Visitors stopped because 64 Raspberry Pis in an acrylic structure immediately sparked curiosity. But the real effect emerged in conversation: what looked like a spectacular maker project turned out to be a demonstrator for a new data architecture. Even small computers can be part of big-data federations. Satellites and edge devices, too, can perform data cube analytics. And even massive Earth observation archives can be connected in such a way that technical boundaries no longer remain a challenge.

Outlook: A New Era of Geoinformation

The transparent cluster symbolizes more than its 64 circuit boards reveal at first glance. It stands for speed — in just a few weeks, planning challenges such as hardware logistics, software porting, and trade fair presentation turned into a functioning demonstrator. It stands for pragmatism — off-the-shelf Raspberry Pis show what becomes possible with smart architecture. And it stands for a new stage in Earth observation: sensors that not only deliver data, but become part of an intelligent, federated information space. Big ideas eventually have to fit into small screws — and that is exactly what succeeded here.

links:

- https://skyfed.space
- https://dataspace.copernicus.eu/ecosystem/services/copernicube
- https://weather.cube4envsec.org/rasdaman-dashboard/
- https://cube4envsec
- https://earthserver.eu

video for this post:
https://www.youtube.com/watch?v=oe6WqWMHE6o&t=9s

contact
Heike Hoenig
science journalist
www.heike-hoenig.de