Data transfer orchestration

Neptune as the glue between different systems

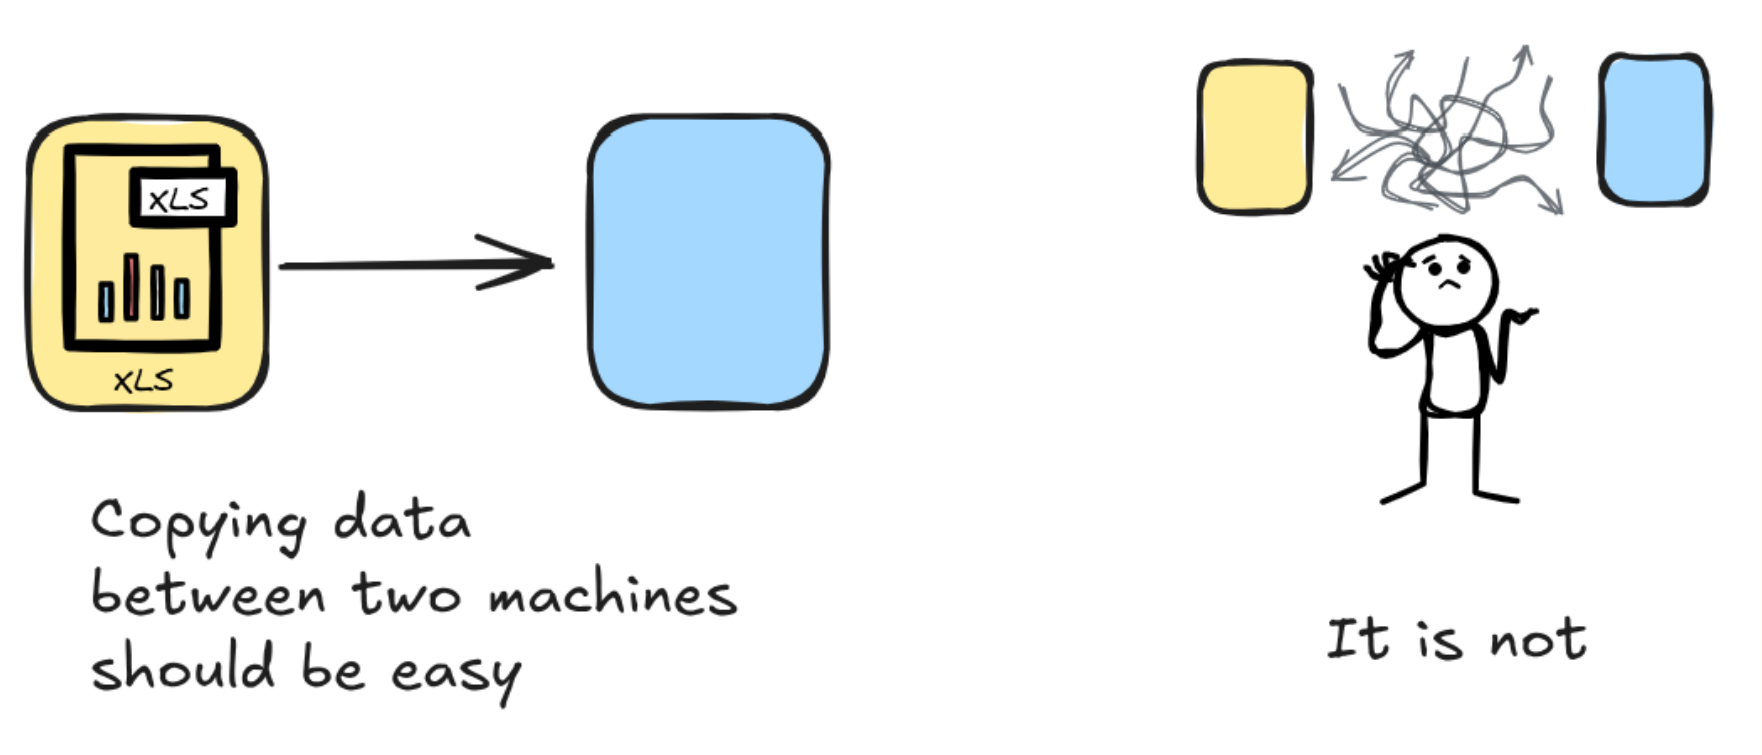
Claudio Cacciari, Eduard Klapwijk

23 November, 2026



Copying data should be easy

WHY



In theory a single arrow. In practice a tangle.

SURF

What is missing: a layer of interoperability

WHY

Systems don’t talk to each other	Data services are not built to communicate on their own. Another system or a person always sits in between.
Collaboration is hard	Every institution has its own services, logins and rules. Moving data across those borders takes a lot of effort.
Sovereignty and security	Once data leaves a controlled environment, you lose sight of where it is and who can reach it.

SURF

This is a Europe-wide problem

WHY

~90%

of users handle **10 GB to 10 TB** of data, yet have no dependable, resumable or automated way to move it.

Each institute reinvents the wheel, adding overhead and uneven security.

The bottleneck isn't bandwidth: it's **usability and interoperability**.

SURF

Source: GÉANT GN5-2 WP4, "Data Movement: Problem Statement & Objectives", TNC'26, June 2026

Three ways in, one engine

HOW

01 · DataConnect



Web portal

For researchers and data stewards: log in, connect services, schedule a transfer.

02 · Command line



CLI

For power users and scripts: transfers as part of your own workflow.

03 · Waves



API

For services and platforms: build Neptune into your own system.

All three run on the same engine: the Waves API and the Triton agent.

SURF

How it works under the hood

HOW

1 · DataConnect

Log in

The user logs in to the web portal and picks the services.

2 · Vault



Connect

Credentials are encrypted and stored in the vault.

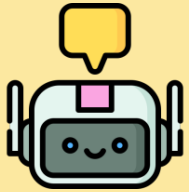
3 · Waves



Schedule

The API stores the transfer request in the database.

4 · Triton



Stream

The agent picks up the request and streams the data between the two services.

5 · Triton

Monitor

The agent monitors the transfer and reports back.

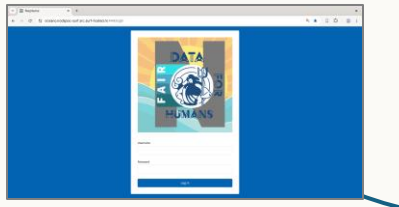
The data never touches the user's laptop.

SURF

Overview

HOW

DataConnect (Web UI)



Surfie (CLI)



Access control

Vault





Waves

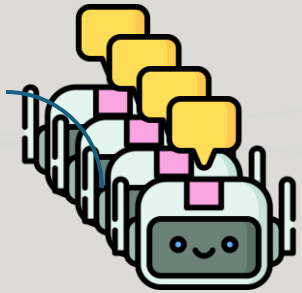


DB

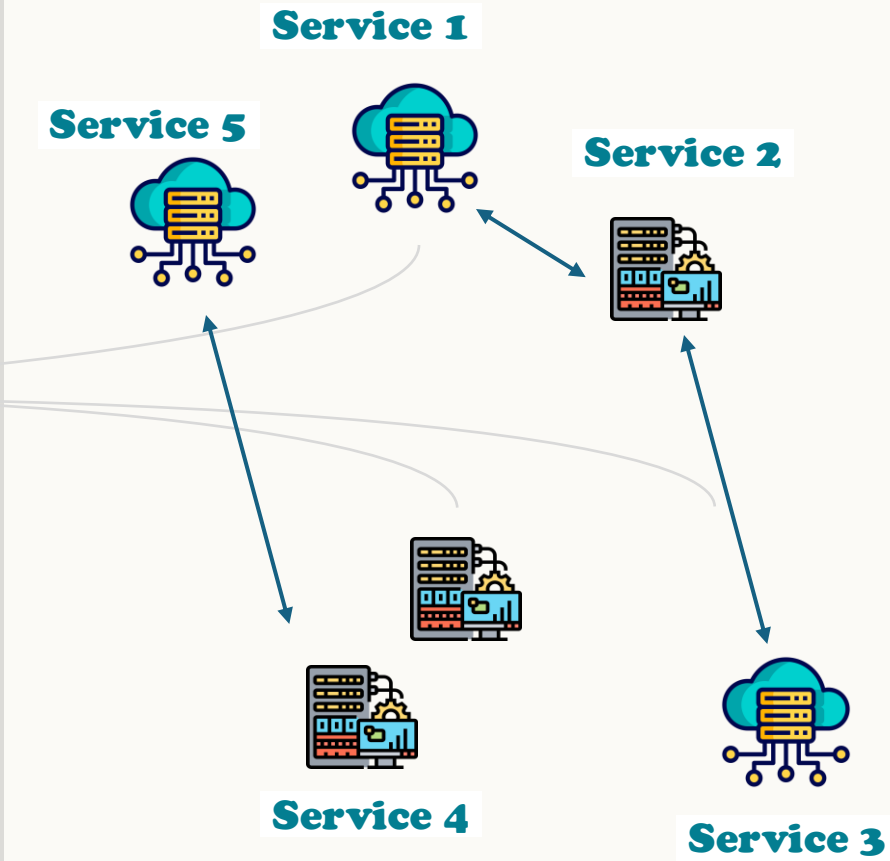
This is where you **scale** *the number of users*

Interoperability

Triton (agent)



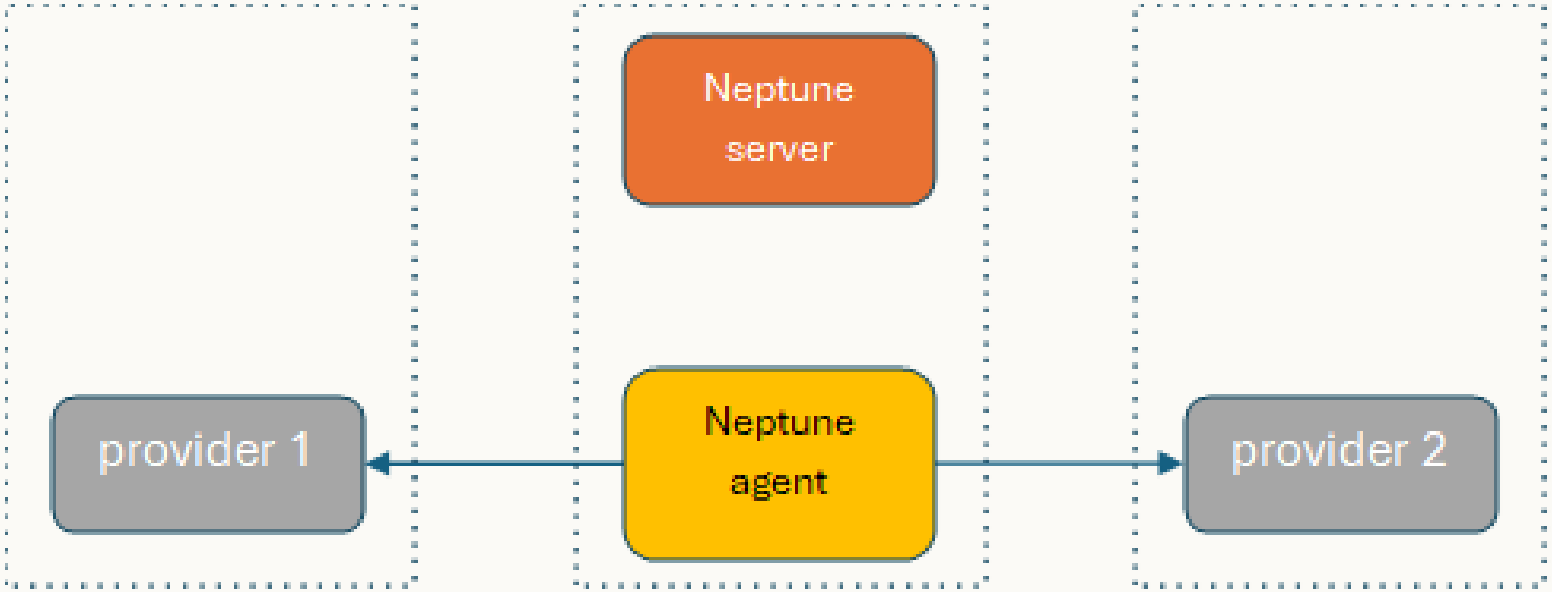
This is where you **scale** *the number of transfers*



SURF

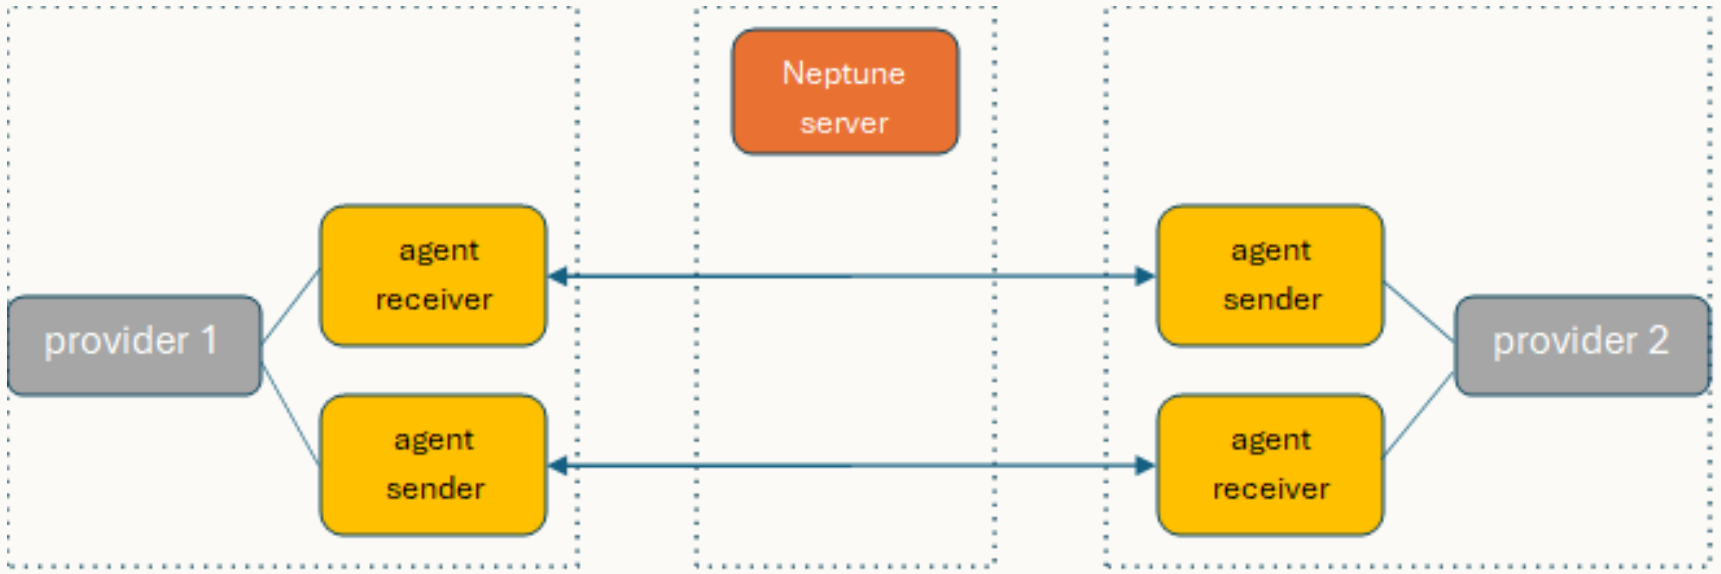
Linux Terminal Vectors by Vecteezy Agent icons created by Freepik - Flaticon Api icons created by SBTS2018 - Flaticon DB icons created by Design Circle - Flaticon Secret icons created by Freepik - Flaticon

Copy with a single agent



The agent streams data directly from the source to the destination

Copy with a sender and a receiver



One agent serves the data, another one pulls the data

Share


XLS
XLS
everyone
a group
a colleague
XLS
XLS
select
your data
select your
recipients
share

Share? How?



share? how?



email



a link



web portal

select the method





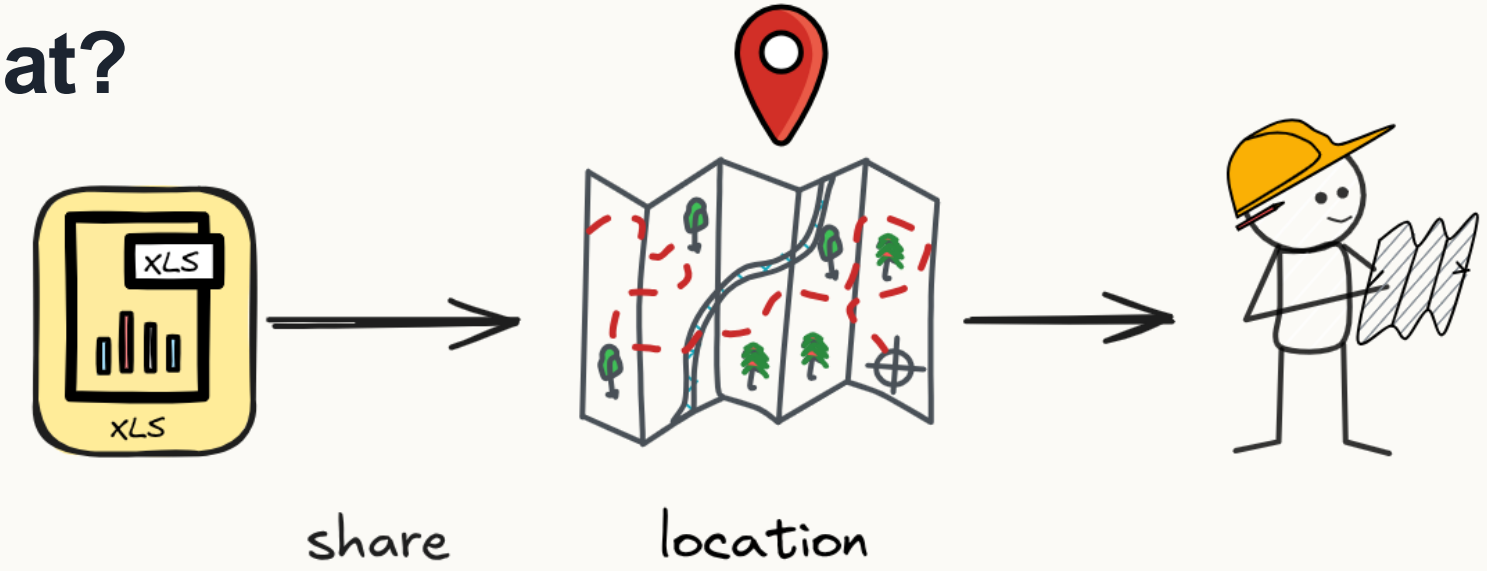
directly accessible

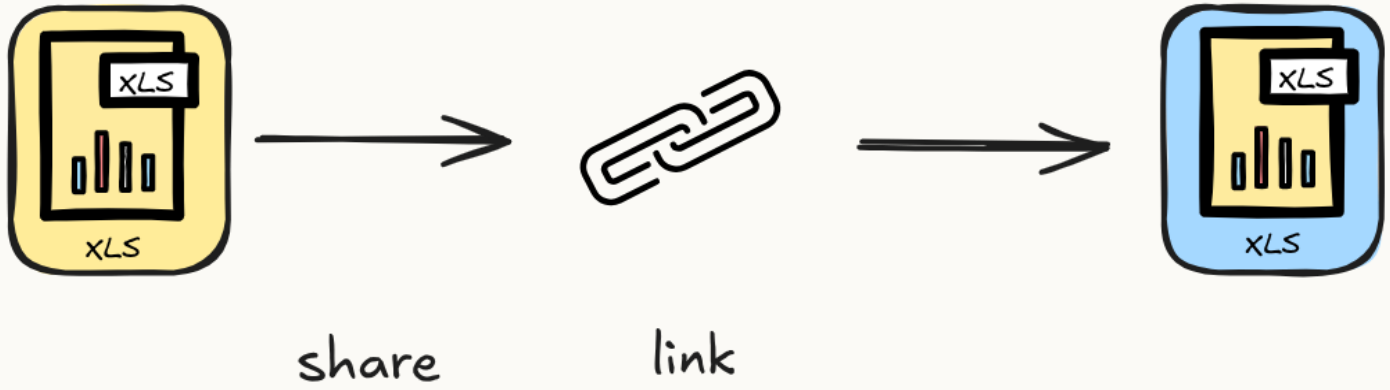


only with my approval

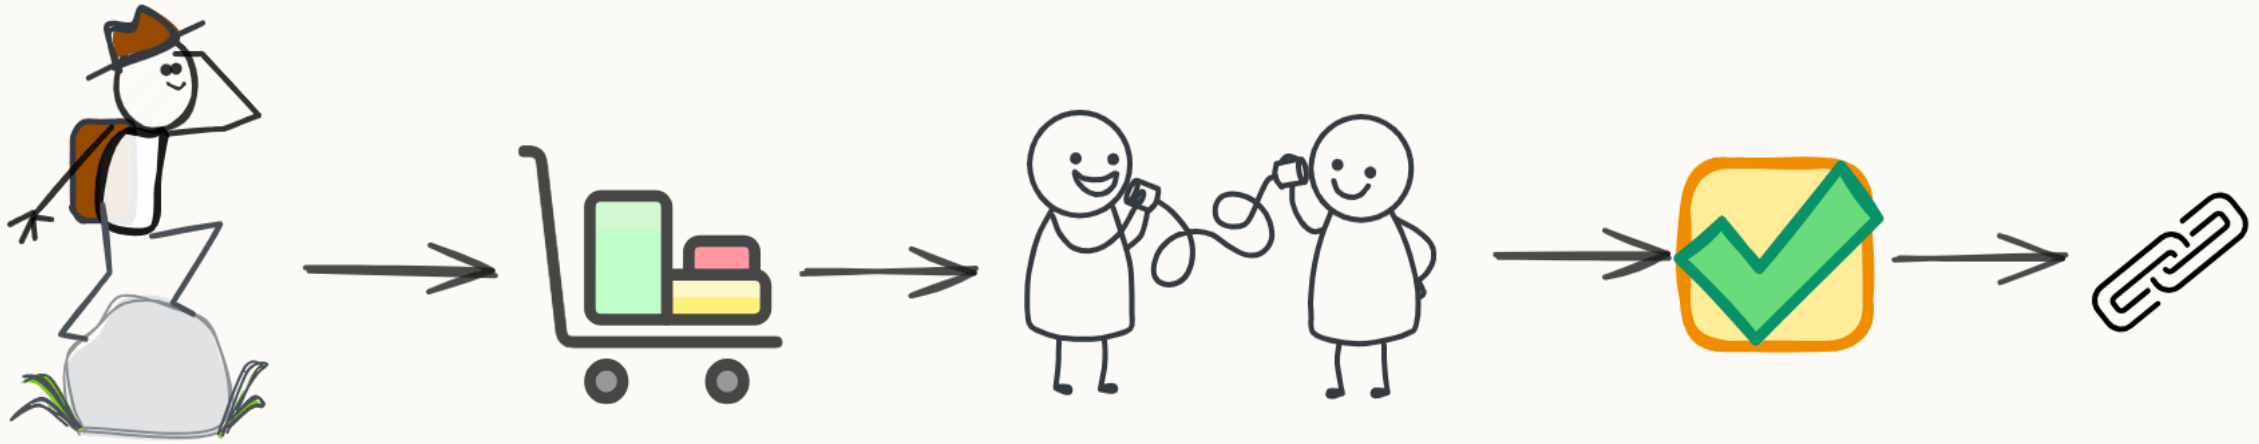
decide how much control you want

Share what?





Request



look for data

select what you want

request it

get approval

get a link

SURF

Conclusions

- We have presented a new tool for data orchestration that offers a friendly way to copy, share and request data.

- We plan to further develop it to make it both more reliable and more interoperable.

- We plan to support
  - Oauth2 Token Exchange protocol
  - Copy and Share metadata
  - Projects as a group of users and services
  - Additional protocols and back-end providers

Questions

Neptune code repository:

https://gitlab.com/surf-dms/neptune_project

edu.nl/3ybpj

