

SURF(net) NG, moving towards intelligent networks

4th GRP, 9th & 10th of October, 2023
Peter Boers

SURF

| Introduction



Peter Boers

IT specialist/software architect


- UVA/SNE
- 7 years @SURF(net)
- Automation & Orchestration
- SCINET
- ML/AI
- Workflow Orchestrator Programme Technical Lead
- **not to be confused with Pieter de Boer** 😊

A brief overview of the network

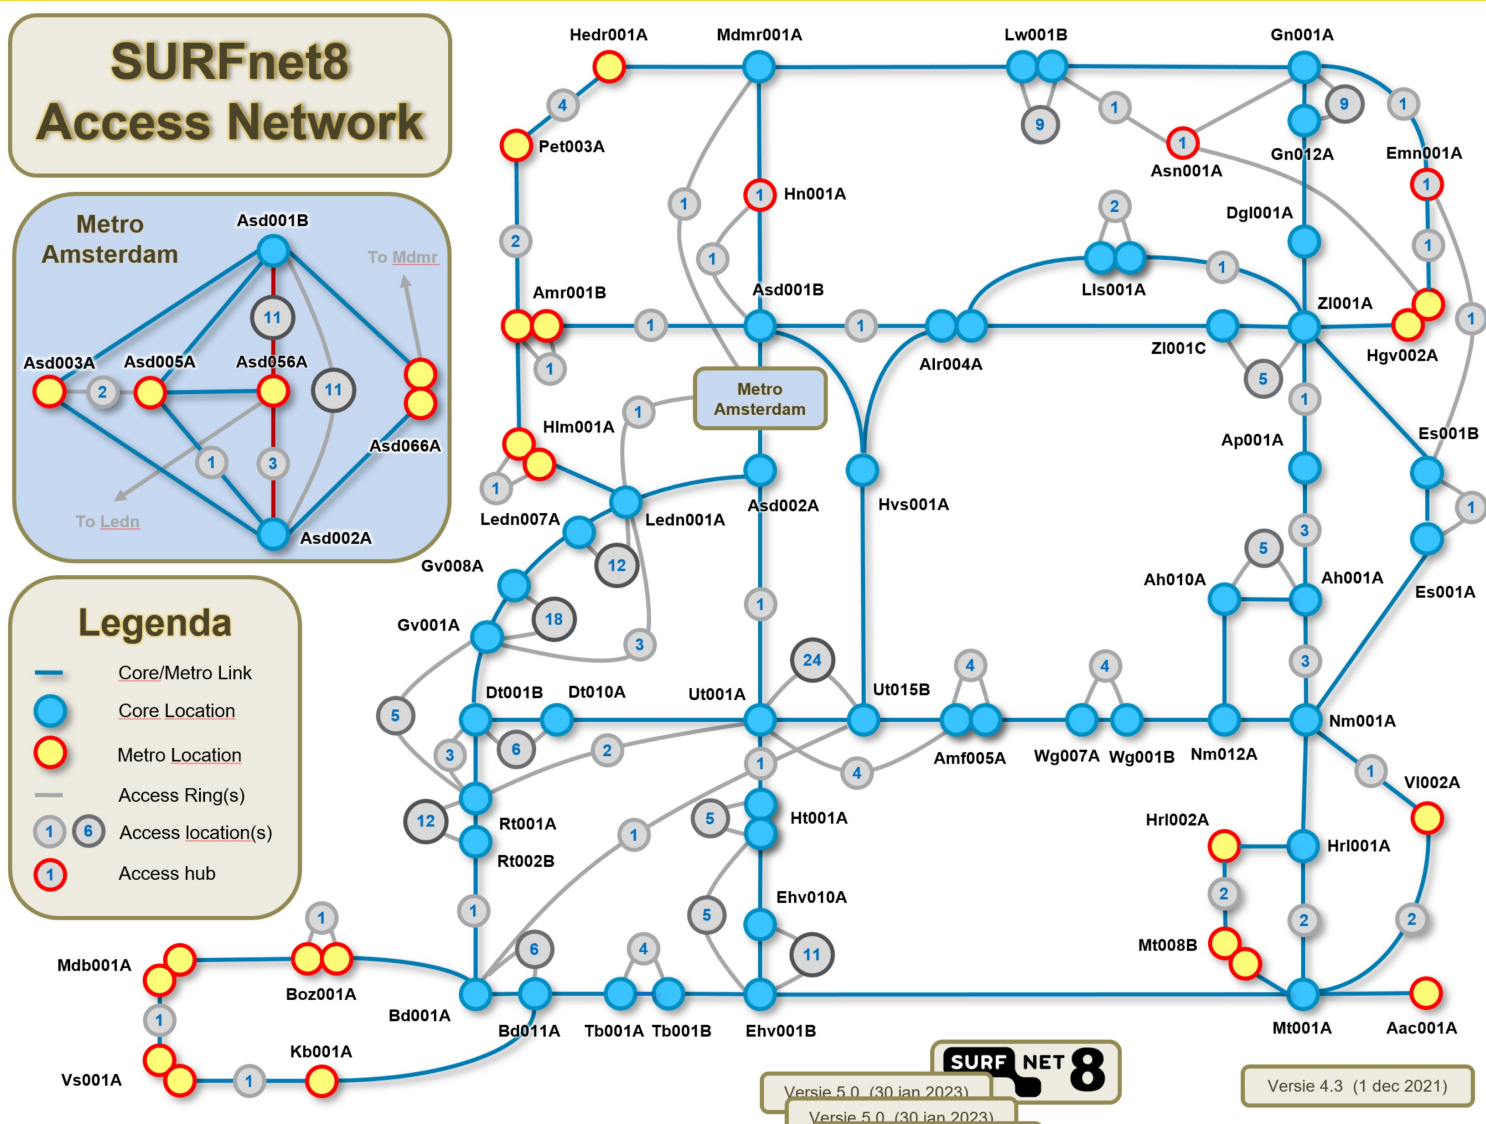
SURFnet8

The eighth refresh of the network since the start of SURF(net):

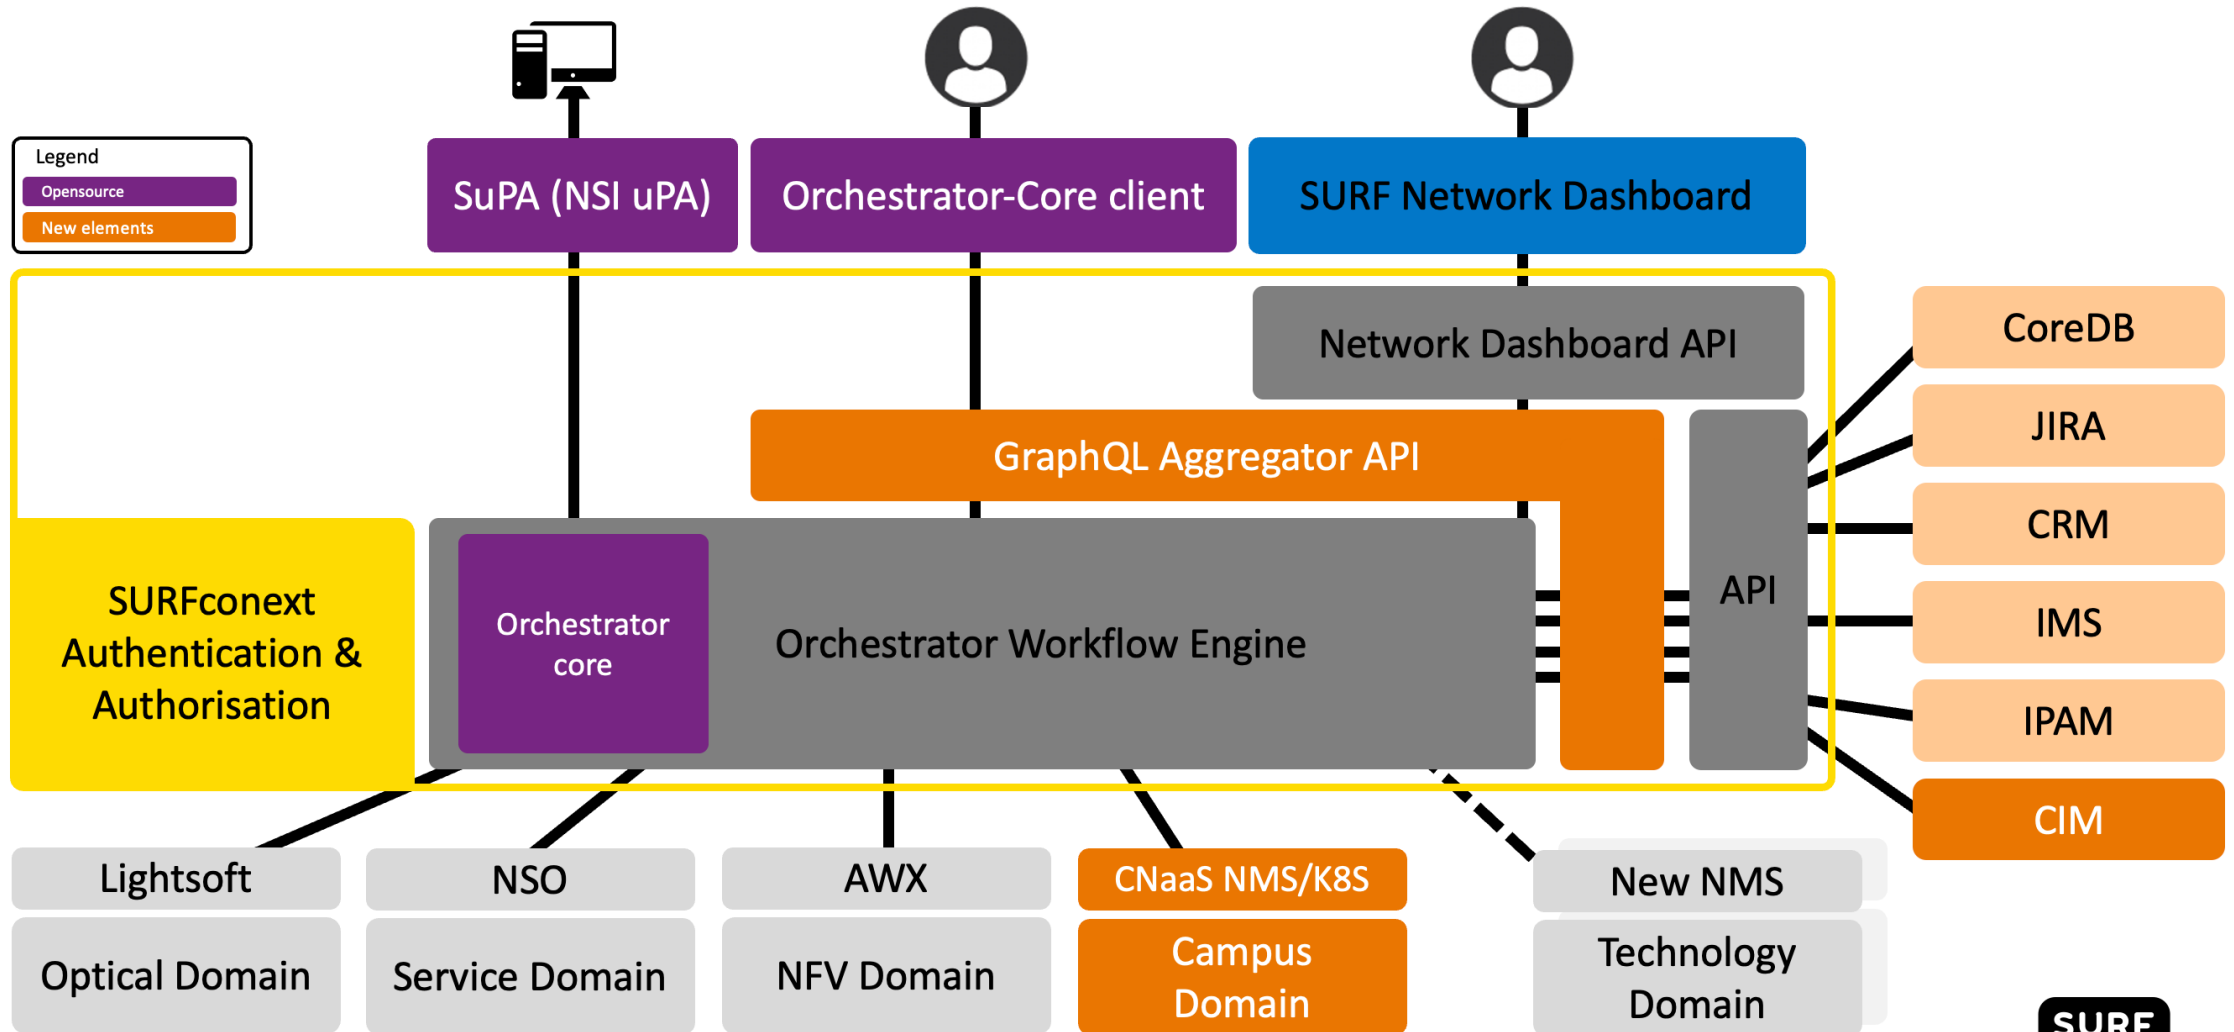
- Optical 400G/100G DWDM, 10G CWDM - 13,000 KM of darkfibre
- 800G upgrade on the CERN link
- 400GE on NetherLight, 200GE (Aggregates) on Backbone
- Daily traffic peak almost 1 Tb/s:
- +/- 420 Gb/s commodity Internet
- +/- 550 Gb/s EVPN, L3VPN: LHC(ONE|OPN), Public clouds, etc
- SR/MPLS
- Juniper MX (Trio)
- 350 + PoP's in the Netherlands and Europe
- 8-9 % of the Dutch population on our network during the day



Network Architecture



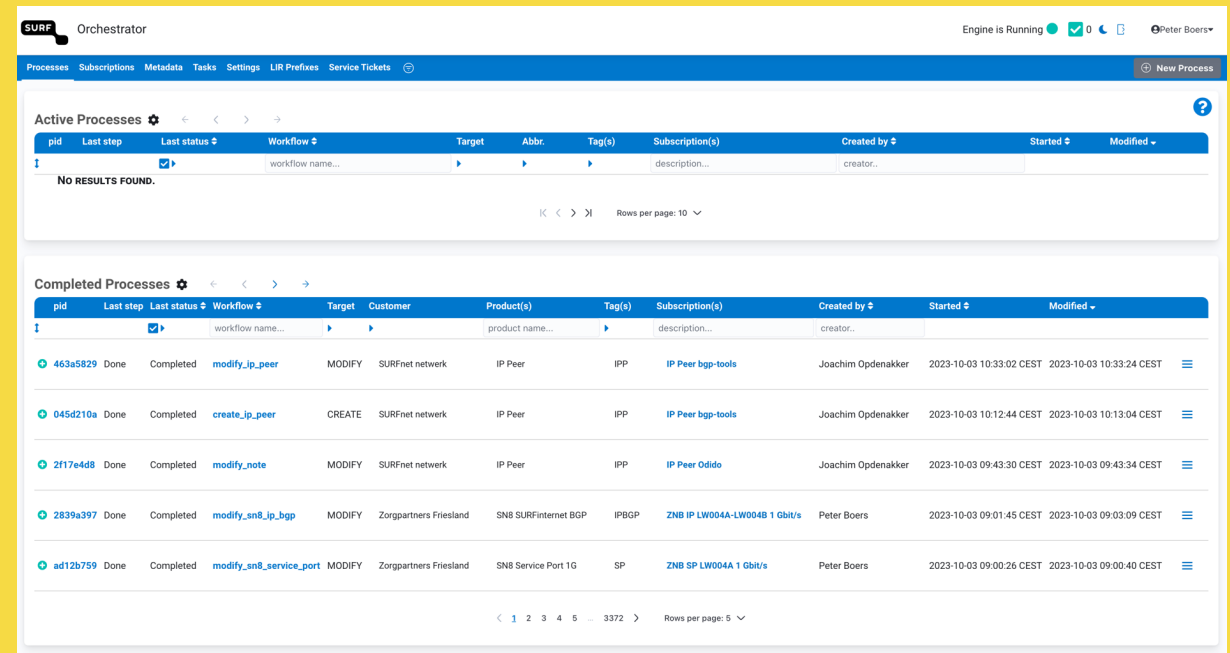
Software architecture



Wide variety of services (1)

Service portfolio

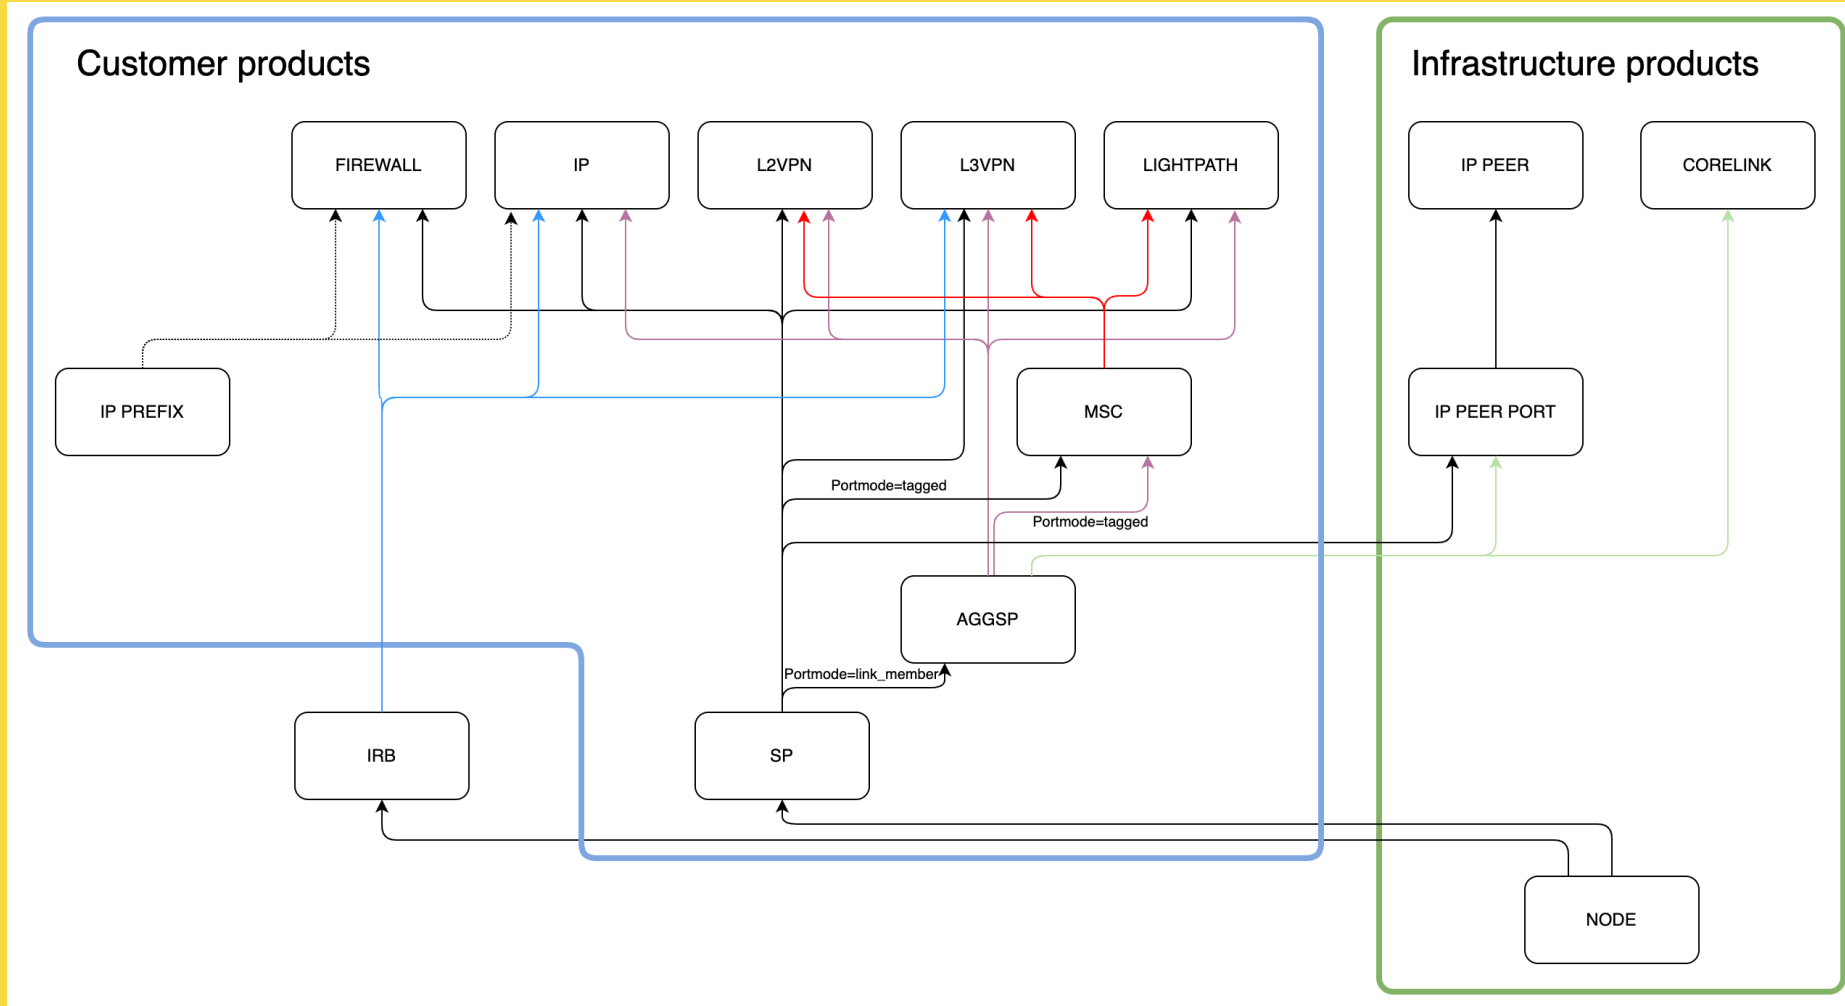
- Diverse customer base (300 institutions): Research, Universities, Polytechnic, Vocational, Primary and Secondary education
- Customer facing services:
- Connectivity services: 1GE – 400GE
- EVPN p2p and p2mp, L3VPN, Internet, Firewall
- LIR
- Resource facing services:
- Nodes, Backbone links, Peering
- Fully Automated and Orchestrated provisioning
- Since 2018 more than 24000 orchestrated changes



The screenshot displays the SURF Orchestrator web interface. At the top, the 'Active Processes' section is empty, showing 'NO RESULTS FOUND.' Below it, the 'Completed Processes' section contains a table with the following data:

pid	Last step	Last status	Workflow	Target	Customer	Product(s)	Tag(s)	Subscription(s)	Created by	Started	Modified
463a5829	Done	Completed	modify_ip_peer	MODIFY	SURFnet network	IP Peer	IPP	IP Peer bgp-tools	Joachim Opendakker	2023-10-03 10:33:02 CEST	2023-10-03 10:33:24 CEST
045d210a	Done	Completed	create_ip_peer	CREATE	SURFnet network	IP Peer	IPP	IP Peer bgp-tools	Joachim Opendakker	2023-10-03 10:12:44 CEST	2023-10-03 10:13:04 CEST
2f17e4d8	Done	Completed	modify_note	MODIFY	SURFnet network	IP Peer	IPP	IP Peer Odido	Joachim Opendakker	2023-10-03 09:43:30 CEST	2023-10-03 09:43:34 CEST
2839a397	Done	Completed	modify_sn8_ip_bgp	MODIFY	Zorgpartners Friesland	SNB SURFInternet BGP	IPBGP	ZNB IP LW004A-LW004B 1 Gbit/s	Peter Boers	2023-10-03 09:01:45 CEST	2023-10-03 09:03:09 CEST
ad12b759	Done	Completed	modify_sn8_service_port	MODIFY	Zorgpartners Friesland	SNB Service Port 1G	SP	ZNB SP LW004A 1 Gbit/s	Peter Boers	2023-10-03 09:00:26 CEST	2023-10-03 09:00:40 CEST

Wide variety of services (2)



| Things are running well, so why migrate...

SURFnet8 is showing its age:

- European procurement laws
- Contracts running out
- Some chassis have limited upgrade potential
- We need to be able to handle the diverse requirements of our customers
- Network architecture is not scalable enough



| SURFnet NG

Why not SURFnet 9?

We are very happy with the stability and quality of the network. No need for a major hardware/software/architecture replacement, just a refresh and evolution.

MX204 shenanigans (2021)

Is an excellent CPE device. Planned usage at SURF until at least 2028, perhaps even into the 2030s. Supports the 1GE – 100GE connectivity services and is 85% of our install base. Not suited for high bandwidth services. Staggered migration to a new CPE device, only when necessary.

Core devices upgrade and architecture change

Disaggregation P vs PE devices, current network is PE only. Geographical redundancy (not just Amsterdam). Smaller devices, high capacity, high density, no more modular chassis(?)

Planning and Logistics

Currently in the RFI phase, planning is to do the RFP in 2024 H1. First purchase in 2024 Q4 and install in 2025 H1

| Homogenous vs Heterogenous

Homogenous network architecture

- The current architecture supports **any** service and interface **anywhere** on the network
- Easy to operate and maintain
- Key to the introduction of Orchestration
- No longer scalable unless we do a full hardware refresh. Can you do 400GE and 1GE on a single box? (No, and you probably wouldn't want to either)
- Rigid architecture and contract. Large potential for wastage.

Heterogenous network architecture

- Multi OS, Multivendor network, and P vs PE
- Investigating Broker model in RFI; we would like to be able to pick and choose vendors and devices
- Be more flexible in price and service offering: low bandwidth and cheap connectivity vs high bandwidth with QoS
- Leverage the power of NSO. The northbound api should remain stable
- Continuous evolution and upgrade, no more big bang upgrades

| Operational challenges

Running a multivendor Network

- Braindrain, **less** people have to do **more** work
- More engineering and integration between vendors
- How do you keep a single pane view on problems
- Different procedures for different boxes
- Network planning and capacity management

How can we reduce the workload of our engineers and leverage the standardisation of our Automation and Orchestration platform, to overcome multivendor network challenges?



| Creating an intelligent network

Relies on a sane data architecture and reliable sources of truth

- SURF validates all data across the network and OSS/BSS
- Precise definitions of services consistently provisioned across the network
- Data labelling is very precise and easy to correlate
- Large amounts of hi quality data
- **Fully orchestrated provisioning**

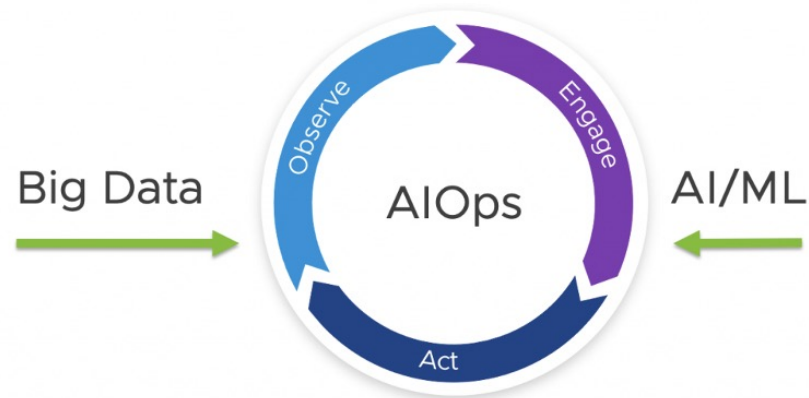
Relies on clear use cases and well-trained algorithms

- Intelligence is relative – How far should you go?
- Training the algorithms towards the wrong outcomes will not result in “intelligence”
- What checks and balances do you need?

| Intelligent network starting point

AIOps

- *An AIOps platform combines big data and machine learning functionality to support all primary IT operations functions through the scalable ingestion and analysis of the ever-increasing volume, variety and velocity of data generated by IT. The platform enables the concurrent use of multiple data sources, data collection methods, and analytical and presentation technologies. - [Gartner](#)*
- Highly supervised human in the loop



| Intelligent networks - long term evolution

Creating a comprehensive network model

- Principle Component Analysis of the network
- Create a virtual network model that enables engineers and architects to plan ahead for topology changes
- Enable network simulations to model network failures
- Use network simulations to drive just in time bandwidth planning and allocation
- Research resource optimisation through the network model that is able to optimise on multiple inputs: Capacity, costs, energy efficiency

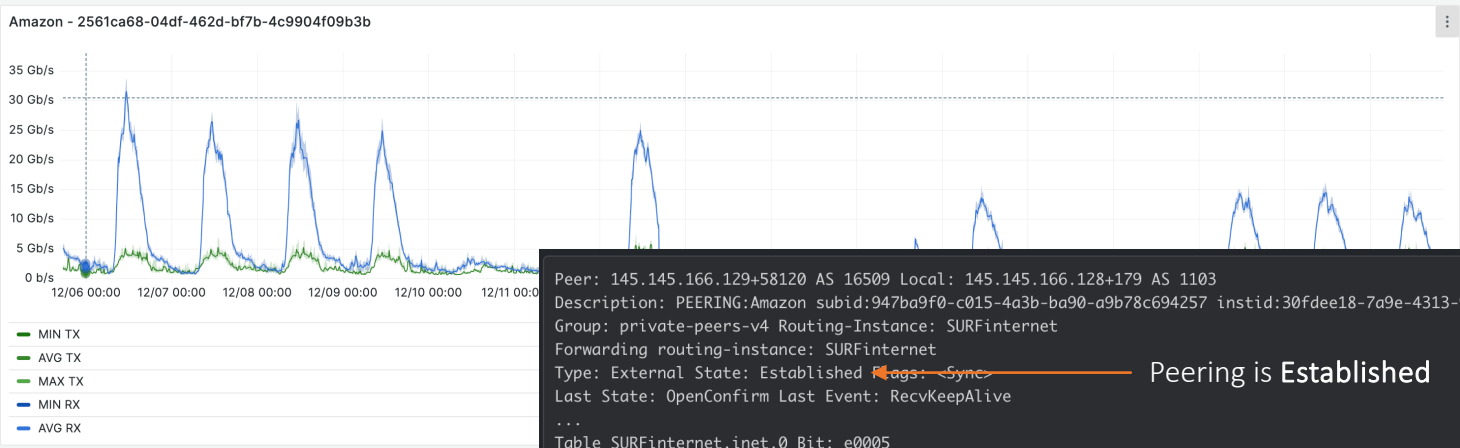
| Where do we start - use cases

- Re-routing of traffic on interface/link failures
- Prefix filter updates based on PeeringDB
- Capacity management
- AI-chatbot in our networkdashboard to capture intent and translate it to self-service
- Troubleshooting suggestions and root cause analysis
- Energy efficient networking
- **Event correlation and anomaly detection**
- **Trend monitoring**

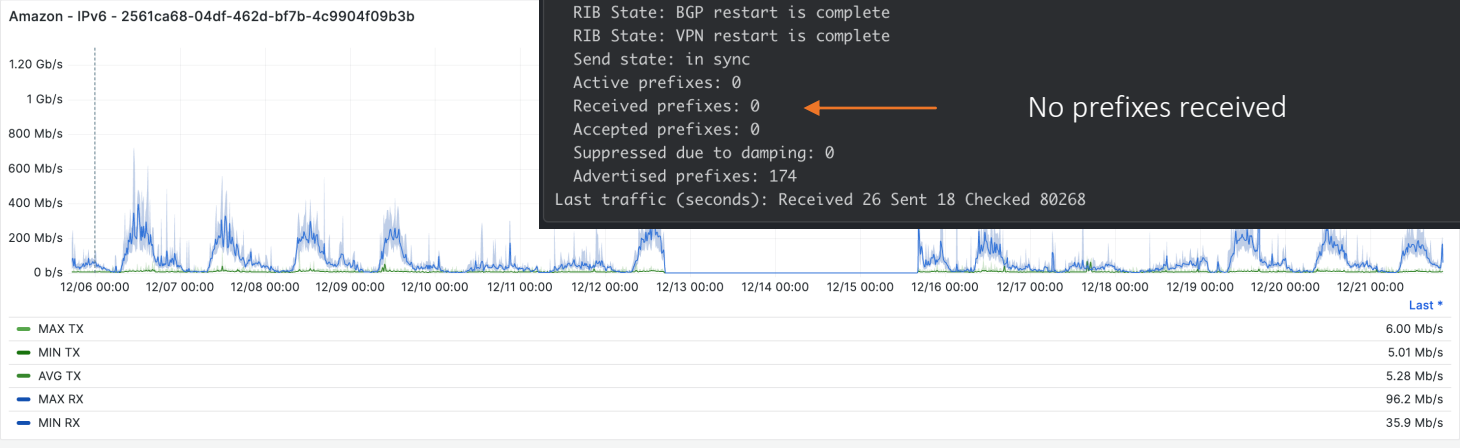


Amazon scenario – event correlation

Incident or not?



```
Peer: 145.145.166.129+58120 AS 16509 Local: 145.145.166.128+179 AS 1103
Description: PEERING:Amazon subid:947ba9f0-c015-4a3b-ba90-a9b78c694257 instid:30fdee18-7a9e-4313-9a56-f3365ac8f152 nsoid:8a13dec2-ad8a-4248-9caf-0b55b7302820
Group: private-peers-v4 Routing-Instance: SURFinternet
Forwarding routing-instance: SURFinternet
Type: External State: Established Flags: Sync ← Peering is Established
Last State: OpenConfirm Last Event: RecvKeepAlive
...
Table SURFinternet.inet.0 Bit: e0005
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: in sync
Active prefixes: 0
Received prefixes: 0 ← No prefixes received
Accepted prefixes: 0
Suppressed due to damping: 0
Advertised prefixes: 174
Last traffic (seconds): Received 26 Sent 18 Checked 80268
```



Microsoft scenario – Trend monitoring

Incident or not?



| Trend and event monitoring

A tangible place to start

- Traditional monitoring has a single dimension
- Is the BGP session up?
- Is the interface up?

A healthy service is described in multiple dimensions

- Are you receiving prefixes?
- Is traffic flowing at the correct levels?
- What is the baseline?
- When is a service healthy?
- What is a healthy service?



| How are we going to start?

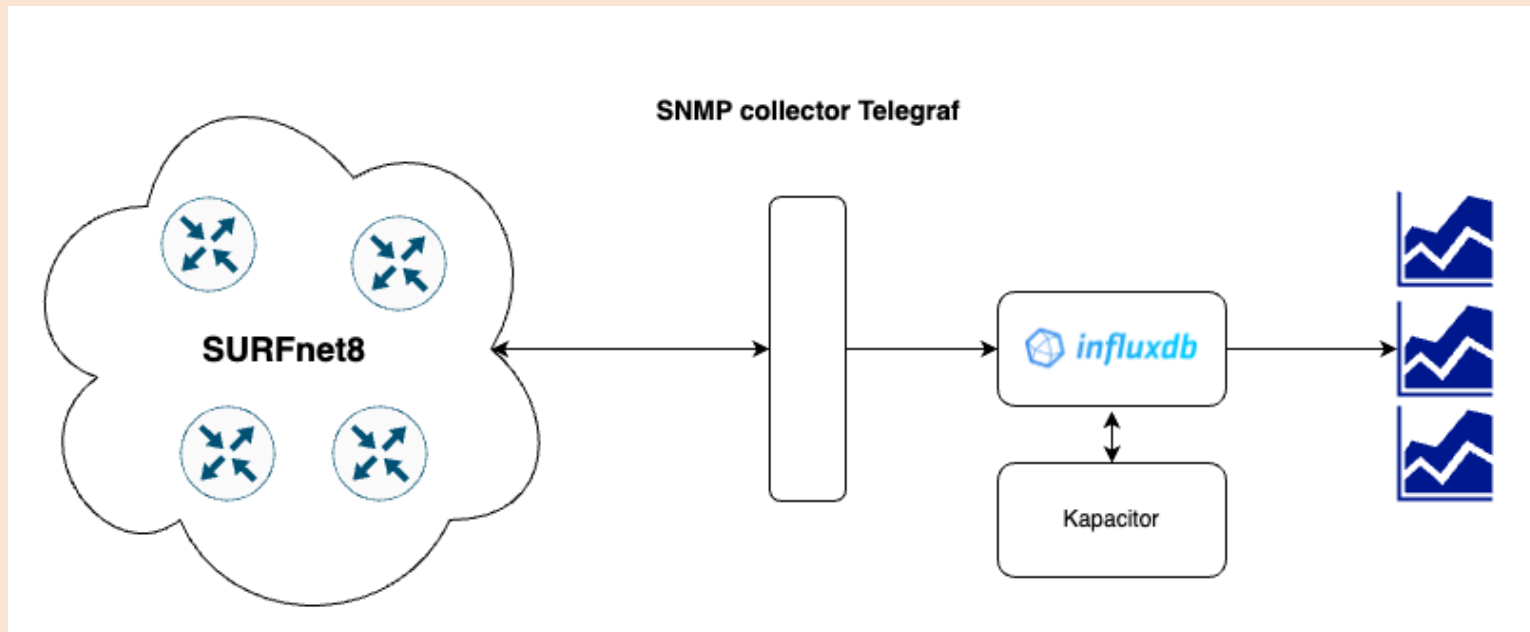
Getting the basics right

- Redesign of our telemetry platform
- Make sure we get the right data with the right resolution
- Start “playing” with the data and creating simple ML models
- LSTM network traffic predictions are relatively accurate (0.96) – ideal for the peering use case



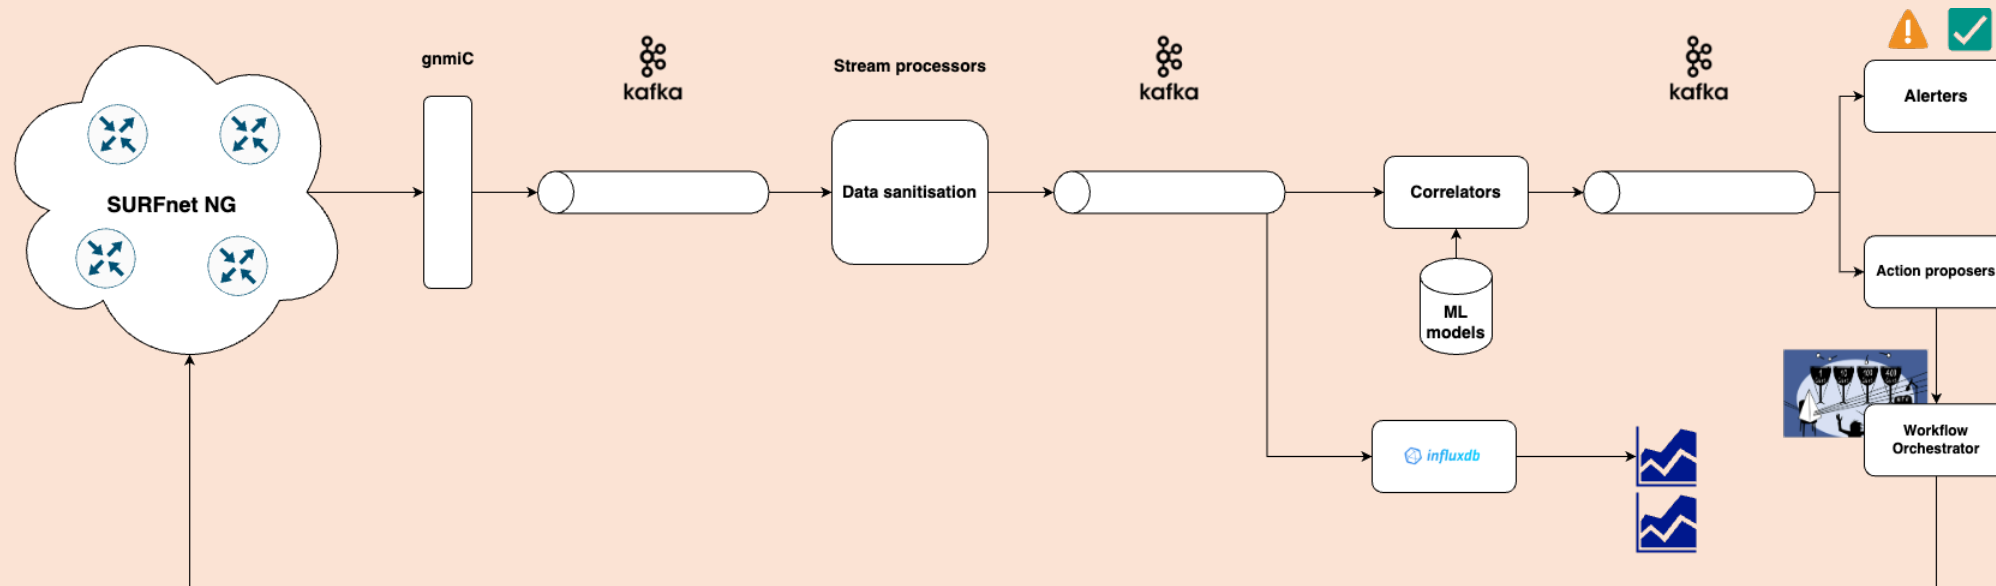
Current telemetry platform

Legacy method of querying devices



| An engineering solution....

We need a way to process real time data from routers in a vendor agnostic way to enhance our monitoring capabilities



| Start small – think big

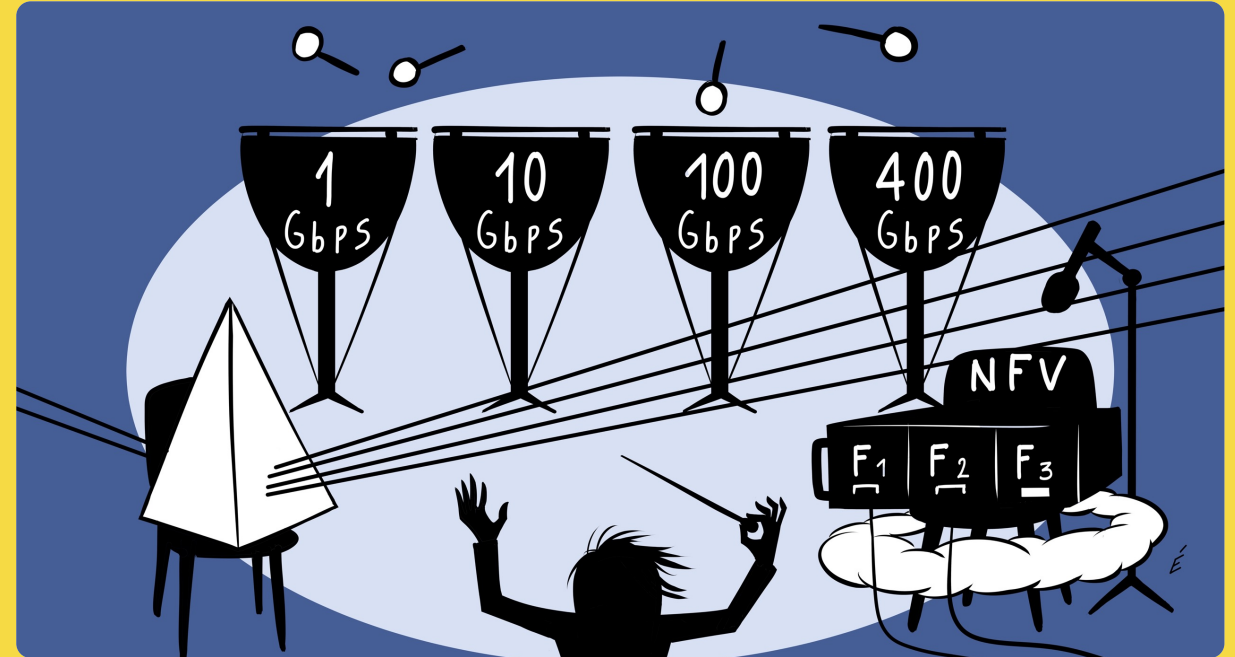
AIOps and ML roadmap:

- Build a scalable streaming platform late 2023
- Early 2024 implement Peering event and trend monitoring to generate alerts
- Build out the operational use cases late 2024
- Start doing the Principal Component Analysis
- 2025 and beyond:
 - automatic reconciliation
 - energy efficiency
 - network simulation – digital twin

Workflow Orchestrator

Essential to our version of an Intelligent network

- Opensource: <https://workfloworchestrator.org>
- Apache 2.0
- Governance in the Commons Conservancy
- 3 tiered membership model
- Used in production by: SURF, ESnet, Géant
- Being prototyped by: HEAnet
- Being investigated by: aarnet, Canarie, Belnet
- We welcome new members!






Any questions?

Thanks!

 Peter Boers

 peter.boers@surf.nl

SURF