

Dreaming of a CS lab

systems and networking perspective

Concepts and needed capabilities
a 'fieldlab demonstrator'?



Maastricht University | Department of Advanced Computing Sciences

*David Groep,
November 2024*

Why the CS lab?

Initiator and 'basic' and short-term 'urgent' needs

- new Bachelor Computer Science, with strong engineering component (esp. in Y2)
- Large-scale IT and robotics also in DSAI Year-3

But: it meshes with {master, PhD} thesis work and cross-discipline and cross-organisation research as well:

- HPC simulation and modelling , ML models, LLM applications, and collaborating with research output 'as a service'
- science and engineering beyond DACS: ET pathfinder & LHCb for GWFP, the MSP programme, ...
- data science in other faculties: need for 'student sandboxes'
- large multi-modal 4D imaging (e.g. for NWO ENW XL FASTER with Ron Heeren), and EU Horizon and Euregio collaborations

Period 1	Period 2	Period 3	Period 4	Period 5	Period 6
Computer Networks;	Software Engineering & Architectures;	P R O J E C T	Embedded Programming;	IT Management & Privacy	P R O J E C T
Introduction to Artificial Intelligence;	Principles of Programming Languages	O J E C T	Computer Security;	Numerical Methods	
Intelligent User Interfaces	1 out of the two elective modules:		Parallel Programming	1 out of the two elective modules	
	- Intelligent Interaction: course Image & Video Processing + Project 2-1: Human-Computer Interaction - Artificial ...			- High-Performance Computing: course High-Performance Computing + Project 2-2: High-Performance Computing - Cybersecurity and ...	

Year 3

Period 1*	Period 2*	Period 3	Period 4	Period 5	Period 6
Digital Society;	Computer Security;	P R O J E C T	Operations Research Case Studies **	BACHELOR'S THESIS	
Game Theory;	Software and Systems Verification;		Intelligent Systems;		
Semantic Web;	Logic for Artificial Intelligence;		Data Analysis		
Recommender Systems;	Parallel Programming;				
Robotics and Embedded Systems;	Large Scale IT and Cloud Computing;				
Introduction to Quantum Computing	Introduction to Bio-Informatics				
PROJECT			BACHELOR'S THESIS		

T: choose the course between information security or ubiquitous computing
Project 2-2: Cybersecurity IoT

'Wereldwijd verbonden vanuit een stevige basis in onze Euregio'

Research and education are inherently collaborative, and 'services exist' and are being used .. whether we know about them or not:

- European Open Science Cloud, 'ESFRI' research infrastructures, Erasmus Without Papers, EU Identity Wallets, SURF-NWO Rekentijd & Snellius, EOSC EU Node 'free VMs for faculty'

'Services await us'
both in thematic and in horizontal e-Infrastructures



ELIUR RI and Life Sciences AAI (bfr), ESCAPE Data Lake by Ricardo Di Maria (CEIM), CSIMESH4ESC - Science Mesh and Services <https://csimesh4esc.eu/science-mesh>

Maastricht University | DACS

Scalable HPC strategy: from local "T2" to European "T0"



Nikhef "Stoombot" Analysis Facility

SURF National Infrastructure solid foundation and essential stepping stone

Samenwerken in onderzoek - een Nederlandse e-Infra aanpak

Given our strategy strives for an attractive research climate "Met hoogwaardige onderzoeksfaciliteiten stellen we hen in staat om excellent onderzoek te doen" - which includes ICT!

Experimental services



Service documentation
Research Data Management
Other services
ObjectID Registry
URN Registry
SSO Identity Management
What is my IP?
Liveness testing at Nikhef
ShareID vintageloc notes
Audioconferencing PIN
Orchestra Experimental Cloud
NWO-I Compute Tool
About Experimental Services

<https://www.nikhef.nl/pdp/doc/experiments>

Sharing more than resources: from data to software & expertise

And some EOSC is coming 'whether you want it or not' ...



AIR – infrastructure for with re-usable software



Vanzelfsprekend hebben we daarbij aandacht voor de impact van digitalisering op onze instelling en de samenleving. Bijvoorbeeld de impact op de manier waarop we leren, werken, contacten onderhouden en onze vrije tijd besteden. Vanuit een continue alertheid spelen we in op de kansen en uitdagingen van digitale technologie, zowel in ons onderwijs, onderzoek als de bedrijfsvoering. Hierbij houden we rekening met het transformatieve karakter van nieuwe technologieën voor de kansen en uitdagingen van de digitale maatschappij.

Images: UM Strategic Programme 2022-2026, SURF ETP
slide captures: www.nikhef.nl/~davidg/p/ for UM-CI office & LSIT course

Translating ‘wereldwijd verbonden’ education & research as a technical systems and network design

1. ‘we can *use* existing services from outside’: AAI interoperability, network connectivity and bandwidth, devolution of responsibility and subsidiarity of concern
2. ‘we can *contribute* in collaborations in education and research’: we can offer services, support pan-European student mobility in our programme, and provide innovative lab environments that attract students and make our programmes unique
3. ‘we teach and train our students to understand, design, build, study, work with, and experiment with interconnected services and systems that are globally connected’

Of course we should do that in a safe, scalable, and transparent way, *and* at the same time build and use our own science and engineering R&E concepts in an agile way

With a view towards the future ... a decadal vision of common reference platform model, usable for DACS & FSE, ... for FASTER 4D-imaging, ... for SBE, FASoS, and UM-wide?

Technical inspiration

- DACS CS-Lab concept *for today* focusses on systems and networks
 - those act a foundation for more specific lab settings: robotics, XR, security
 - focus on education labs, but equally applicable to research networks
- design around micro-segmented networks and inter-subnet controls
 - so you could build multiple labs on a partly-shared infrastructure
 - leveraging the ESnet (US DoE) 'ScienceDMZ' concept, and
 - on the UvA Security and Network Engineering master lab ('every student a machine on the public internet and they are *entirely responsible* for this box')
 - with inspiration from SURF Experimental Technologies Platform and Nikhef open innovation lab
- for CS and engineering education (extensible to research), and **not** targeting enterprise services, **nor** for now sensitive data services



'ScienceDMZ'

Predicable performance
and data access for research

*'where research services,
data, and researchers meet'*

- latency hiding through caching
- security zoning/segmentation protects specific data sets
- outside any enterprise perimeter

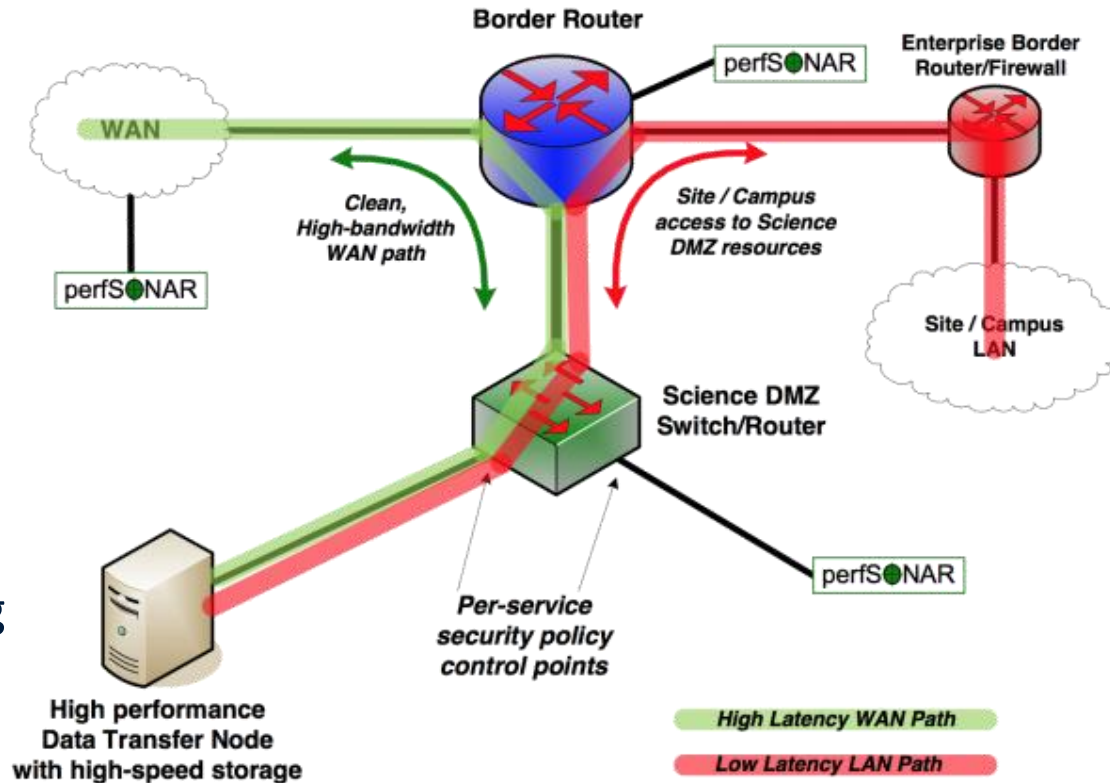
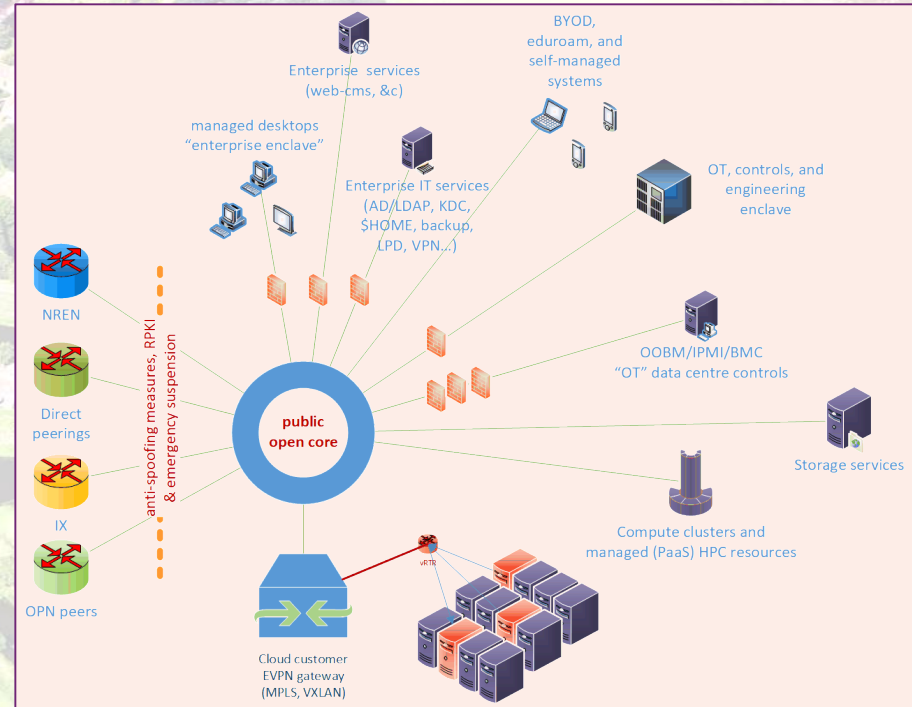
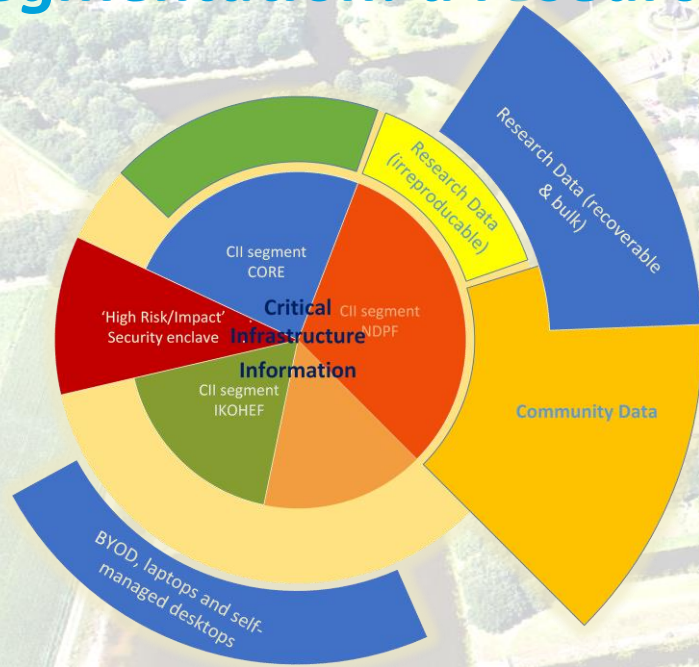


Image and 'ScienceDMZ' concept promulgated by ESnet (see fasterdata.es.net)

Network structure design: logical and topological view *segmentation: a research network with office enclaves*



achtergrondbeeld: Vesting Bourtange;
segmentation impression represents Nikhef Security Technical and Organisational Measures

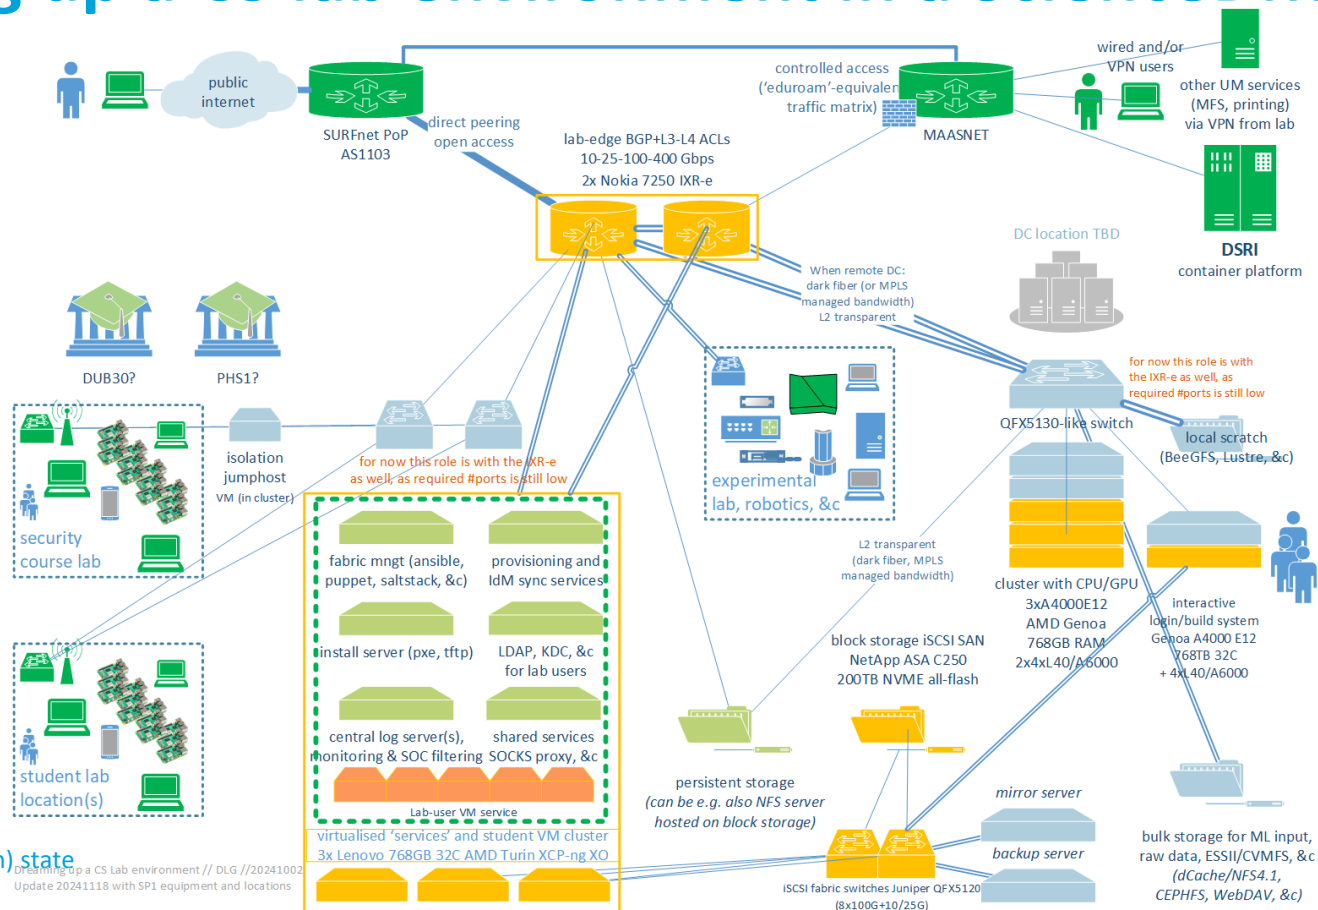
The Value Of Routine Performance

- It's important to get to where high performance is normal
- No magic, no arcana, things just normally work – for petabytes of data
- DOE HPC facilities now easily shuffle around hundreds of terabytes
 - Some people have smaller data sets too
 - But the point is that it's normal and routine
- What follows is one specific example, chosen because of some specific features



From Eli Dart (ESnet), “The Strategic Future of the Science DMZ”, TNC23, <https://indico.geant.org/event/2/contributions/186/attachments/168/>

Dreaming up a CS lab environment in a ScienceDMZ



SOLL and IST (yellow, green) state
of the CSLab foundation

Dreaming up a CS Lab environment // DLG // 20241102
Update 20241118 with SP1 equipment and locations

CS Lab – the 2024 foundation (<20 December 2024)

Virtualised base infrastructure: fabric management, logging, IAM bridging (to SURF SRAM or locally with a SAML SP), proxy and gateway services for sub-labs, ~300+ student sandbox VMs, XCP-ng & XenOrchestra managed with SAML or OIDC auth

- 3x SR635v3 (each: 786GB RAM, 2x25G data network, 2x25G iSCSI, 64 logical cores 9355P)

Storage back-end service: block storage for VMs (supporting NFSv4 via VM with linux)

- NetApp C250 with 200 TByte net NVME all-flash capacity, 4x100G iSCSI

HPC starter cluster: AI/ML courses, reasoning, HPC modelling, interactive & batch

- 4x ASUS ESC4000A (each: 3.25GHz 64 logical core EPYC 9355P, 768 GB RAM, 2x16TB NVME, 2x25G data network), 3 of these systems have 4 Nvidia L40 GPUs each (so 12 GPUs in total)
- one node as the 'interactive' validation and compilation node (scheduling will be in a VM)

Network: access gateway for standard edge connectivity (BGP, MPLS, 802.1q, EVPN, L3/L4 ingress and egress ACLs and policers) and iSCSI SAN fabric switches (25/100G)

- pair of Nokia 7250 IXR-e, each: 2xQSFP28 100G, 8xSFP28 25G, 24xSFP+ 10G (or 1G)
- pair of Juniper QFX 5120's, each: 8xQSFP28 100G, 28xSFP28 25/10/1G

More discussion time ...

David Groep

david.groep@maastrichtuniversity.nl

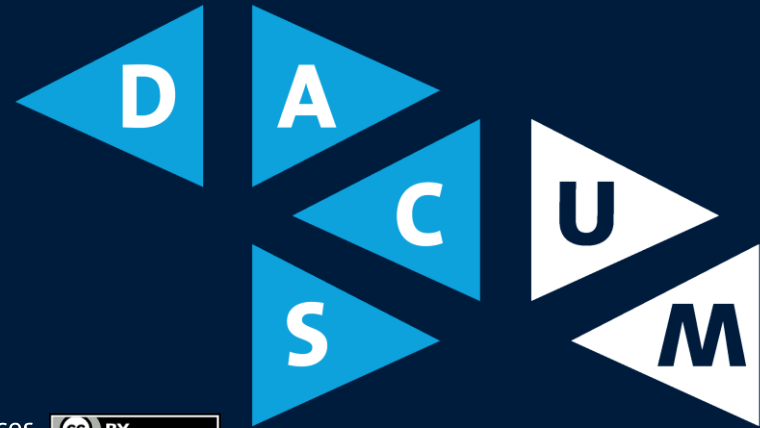
<https://www.nikhef.nl/~davidg/presentations/>

 <https://orcid.org/0000-0003-1026-6606>

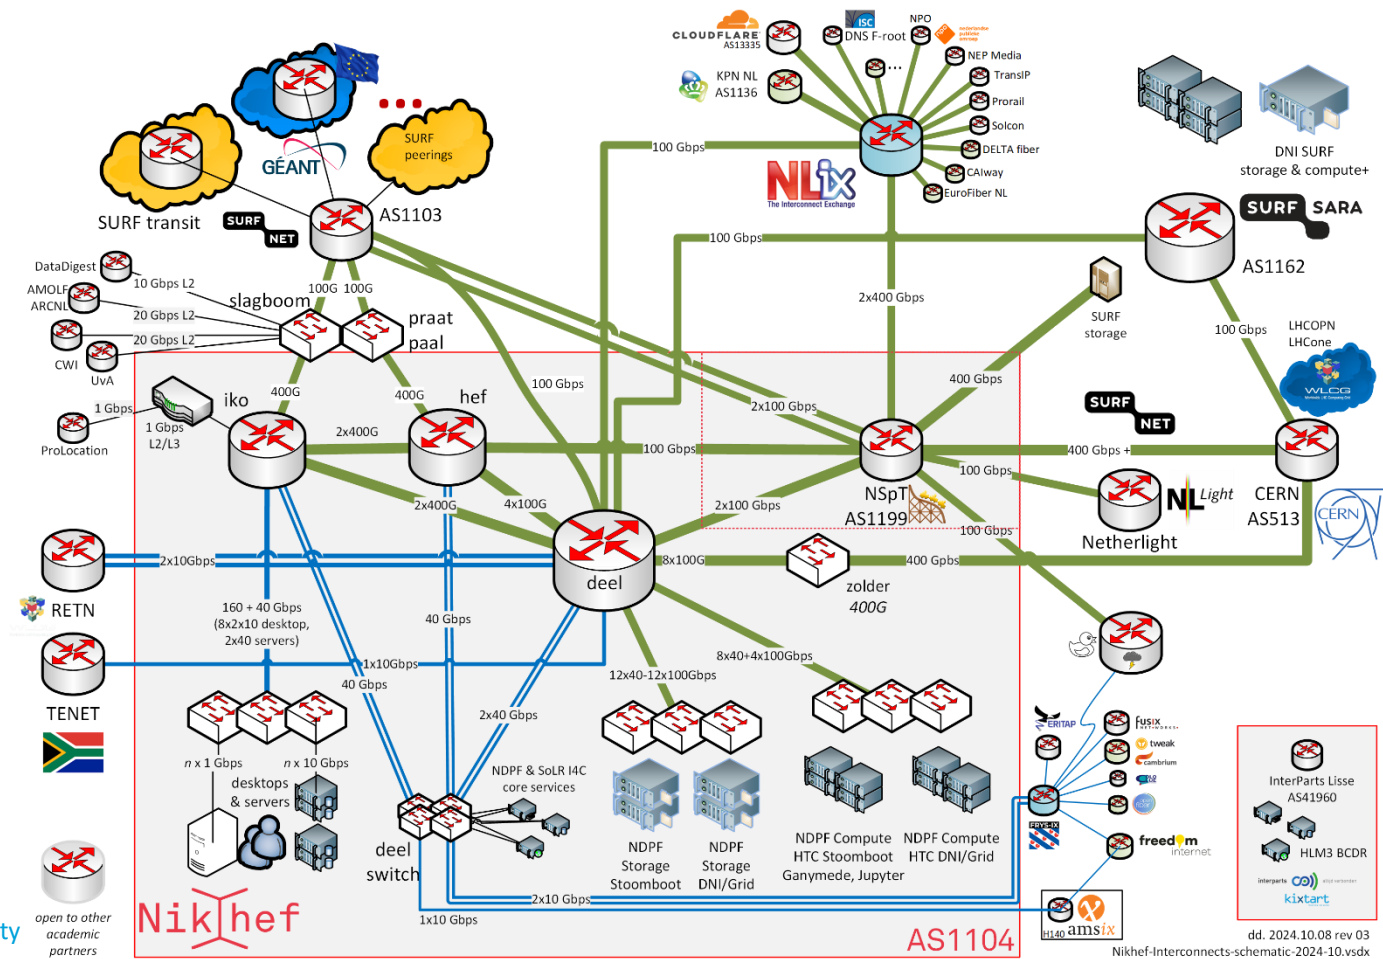
Nikhef



Maastricht University | Department of Advanced Computing Sciences



AS1104



<https://www.nikhef.nl/pdp/doc/facility>