## CS lab: a fieldlab for fieldlabs Massinda University

The DACS FSE Computing Sciences Labada as a Fieldlab Demonstrator

Concepts, capabilities, connections



Department of Advanced Computing Sciences

January 2025

#### 'Wereldwijd verbonden vanuit stevige basis in 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 all faculty'



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

🔀 Maastricht University | DACS

voor de kansen en uitdagingen van de digitale maatschappij.

# 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, and devolution of concerns and work to research infrastructures (ESFRI&GWI)
- 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, GWFP/LHCb, ET, ... for FASTER 4D-imaging, ... for other faculties and UM-wide?



## Why the CS lab?

#### Initially triggered by 'basic' short-term needs

- new Bachelor Computer Science, with strong engineering component (esp. in Y2 & Y3)
- Large-scale IT and robotics, also in DSAI Year-3

## But with wide use for {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

re ering & ctures; les of mming lges of the ective es: lgent	P R J E C T	Embedded Programmin Computer Security; Parallel Programmin	ng; IT m N N N	Manage- nent & Privacy umerical Nethods	P R O J
les of mming ages if the ective es: igent	O J E C T	Computer Security; Parallel Programmi	N N 1	umerical Aethods	к О Ј
f the ective es: igent	C T	Parallel Programmi	1 nor tu	out of the	O J E
igent	-		m	out of the vo elective odules	C T
e Image eo ssing + ct 2-1: an- outer action cial			-	High- Performance Computing: course High- Performance Computing + Project 2-2: High- Performance Computing ¬yber-	
Period 3	Period A	Period 5	Period 6	curity and	
_	Operations Research Ca Studies ** Intelligent Systems;	BACHEL THES	OR'S IS	the course tween formation curity or biquitous	
ntic Web; Logic for P Artificial R Intelligence; O		is		omputing IoT + oject 2-2:	
E C T				/bersecurity IoT	
	Igent action: ie Image leo issing + ct 2-1: an- puter action clal Period 3 Period 3 Period 3 Comparison R O J E C T T	Igent Action: ie Image leo Ssing + ct 2-1: an- puter action cial Period 3 Period 4 Operations Research Ca Studies ** Intelligent Systems: P P Data Analysi C C T	Igent Action: ie Image leo ssing + ct 2-1: an- puter action clal Period 3 Period 4 Period 5 ( Operations Research Case studies ** Intelligent Systems; P R O J E C T Data Analysis	Igent Action: ie Image leo Ssing + ct 2-1: an- puter action cial Period 3 Period 4 Period 5 Period 6 Period 5 Period 6 Period 5 Period 6 BACHELOR'S THESIS BACHELOR'S THESIS	Igent Action: le Image leo ssing + ct 2-1: an- puter action clal Period 3 Period 4 Period 5 Period 6 Period 7 Period 5 Period 6 Period 7 Period 7 Period 7 Period 7 Period 7 Period 7 Period 8 Period 8 Period 6 Period 9 Period 7 Period 6 Period 9 Period 9 Period 9 Period 9 Period 9 Period 9 Period 9

Maastricht University | DACS

https://curriculum.maastrichtuniversity.nl/sites/default/files/dacs\_-\_study\_guide\_2024-2025.pdf DACS FSE Computing Sciences Lab as a Fieldlab Demonstrator

## Suddenly ... CSlab

Maastricht University

#### And then there was Sectorplan 1 ... whose funds needed to be **spent in 2024**





### **Technical inspiration**

- DACS CS-Lab concept for today focusses on systems, networks, and course/project labs
  - 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
  - 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 Security and Networ
   Inot targeting enterprise services, nor for now sensitive data services

Maastricht University | DACS

### 'ScienceDMZ'

## Predicable performance and data access for research

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

- lossless transfers and latency hiding
- elements horizontally scalable
- security zoning/segmentation protects services and data
- outside enterprise perimeter

#### Science DMZ Design Pattern (Abstract)



for a different type of traffic: science data transfers rather than eyeball networks

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

🔀 Maastricht University | DACS

#### **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/

Maastricht University | DACS

DACS FSE Computing Sciences Lab as a Fieldlab Demonstrator

# A small amount of packet loss makes a huge difference in TCP performance

10000 9000 8000 With loss, high performance Local hroughput (Mbits/sec) 7000 (LAN) beyond metro distances is 6000 essentially impossible International 5000 Metro Area Regional 4000 Continental 3000 2000 1000 ............. ........... 0 10 20 30 40 50 60 70 80 90 100 Round Trip Time (milliseconds) Measured (TCP Reno) Measured (HTCP) Theoretical (TCP Reno) Measured (no loss) .........

Throughput vs. Increasing Latency with .0046% Packet Loss

Eli Dart, The Science DMZ Design Pattern, NASA talk, <u>https://edu.nl/a7efv</u>



6 - ESnet Science Engagement (engage@es.net) - 7/13/17

#### Labs as a design pattern ...



See <u>https://doi.org/10.3233/SPR-140382</u> and Eli Dart's NASA talk: <u>https://edu.nl/a7efv</u> for a quick read or glance

Right: components in a medical sciences DMZ

from https://crd.lbl.gov/news-and-publications/news/2017/esnets-science-dmz-design-could-help-transfer-protect-medical-research-data/

Enterprise Border Router/Firewall Border Router 10x10GE 100G WAN Clean, High-Bandwidth WAN path 10x10GE External access to Site / Campus 8x10GE access to Biomedical Science **Biomedical Science DMZ** through portal and head DMZ resources through portal nodes and head nodes Biomedical Science DMZ Switch / Router (stateless firewall) Site/Campus LAN multiple 10GE High Latency WAN Path Low Latency LAN Path high volume data flows security services and permitted by monitoring Router/Firewall for portal authenticated users Per-service head nodes security policy accessed via PKI control points Portal Per-service security policy control points compute and storage nodes that accessed by authenticated and authorized users via the head node using PKI

Maastricht University | DACS

### And once you have multiple FieldLabs

#### Multiple Science DMZs – Dark Fiber



https://smd-prod.s3.amazonaws.com/science-red/s3fs-public/atoms/files/3-2-Science%20DMZ%20Design%20Pattern%20-%20ESNet%20-%20Eli%20Dart.pdf

Maastricht University | DACS

Slides: Eli Dart, ESnet (2017)

DACS FSE Computing Sciences Lab as a Fieldlab Demonstrator 11

#### Dreaming up a CS lab environment in a ScienceDMZ

of the CSlab foundation



DACS FSE Computing Sciences Lab as a Fieldlab Demonstrator 12

#### 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

Maastricht University | DACS

## A lab for fieldlabs

It is also a technology lab & computing sciences engineering innovation lab

- networking from commonly used protocols like evpn-vpls/mpls to high-speed udp science data transport or alternative tcp stacks, or segment routing and connecting to LHCone at 400Gbps to 1.6Tbps?
- storage, compute, acceleration, cluster computing, virtualisation, orchestration, automation, AuthN & AuthZ infrastructures and IAM, ...
- power usage, application efficiency, divergence of 2019/424 to actual usage?
- 'Lab' aspect need not to stop at layer-7
- ITSRM, SMS & FitSM/ISO20k/ITIL, ... to federated services and Tier-2/1/0 scaleout
- red-teaming in resilience exercises with SURF/NMDDoS or for ourselves?

#### A collaborative demonstrator

The Fieldlab & ScienceDMZ concepts are readily generalizable

- both a technology lab (networking, computing, storage, IoT, security)
- as well as a infrastructure and systems lab for SMS and ITSRM concepts

So let's try and use this for an integrated approach already now for the CSLab

- linked to current and evolving ICT at UM as a whole
- ITSRM for the CSLab in line with UM IT security risk management e.g. Lab ISO27k5 ITSRM doc leverages and supports the UM risk assessment

look at this also as a playground for new IT service mngt ideas and experiments – this is a 'harmless' lab with forgiving users and local flexibility 😳

### The most important ingredients ...

- Creativity
- Flexibility
- Desire to experiment

CSLab equipment and configuration is a usual high availability design, but

- can, may, and should be used for experiments it is not our national critical infra ©
- a common playground for students, researchers, and ICT engineers
- diversity in equipment, components, client devices, IoT stuff, and tools is *intentional* for experimentation
- overprovisioning is intentional, as lab can consolidate 'loose' things and 'lone' systems

#### Other labs may have a different scope and purpose, but all share the top-3 elements ...

### **Collaboration ... because it's absolutely needed**

Although FSE will take care of hardware, systems management, networking, &c, *in our interactions with 'the world', UM must appear as a coherent whole* 

- interaction with SURF on connectivity, Internet number resources, SURF Access even though FSE DACS can initiate and prepare SURF staff for the requests, SURF really wants these to come from our institutional contact persons
- incident response coordination & Sirtfi
- network transport across campus building locations
- housing and secured access, power and PDUs, compliance BaL 3.7.3 & 5.4.1 (5.4.1a)



### **Collaboration ... because it's just better together**

- Identity and Access Management, for linking lab services to UM identity & eduGAIN an experiment: CSLab end-users shall not have any lab-specific re-usable credentials, use only SAML, OIDC, or Wallet VCs to authenticate, or ssh pubkeys for login/access
- release of attributes and provisioning interfaces for some lab services *like: which students are in which course, and hence need to be granted access to which parts of the lab?*
- prevent duplication of SMS support that the CSLab needs but can be usefully shared
  - Can we re-use elements, structure, and templates of an SMS that is already in place?
  - If need to construct (part of) our own SMS (which takes time) can we do it such that it is generally useful?



### **Collaboration ... because we're all human**

- coordinating shared data center facilities
- alignment on identifiers because it prevents confusion in human communication and for systems communications as well in some cases...
- joint access to documentation and SMS information
- process for Q's sent by accident to central helpdesk that should have gone to the lab?
- mechanism for periodic technical meetings & chat, since working together is nice, ... and because we all forget to ask for the right stuff sometimes!



### **Collaboration - practically**

								Chat	-
< > Q. Search (Ctrl+E)									
All teams Ontwerp Posts Files CSLab commissionin v									
_									
FI		+ Add new item 🜐 Edit in grid view 🏷 Undo 🖄 Share 🏥 Export - 📰 Forms (New) P\$ Automate - 🕀 Integrate							
E-CS lab		CSLab commissioning and dep		0					
lain Channels			T Prin v	T Status v	T Due v	Comp ~	Description and notes	OneDrive	
General hardware lab 🚓		High lovel design degume	Normal	le Drograce	2024 11 27		Included in ITCPM document as see	Wined	
LabBeheer 👌		High level design docume	Worman	Consolate	2024-11-27		DEB1 alcore for source or ideaset		
Ontwerp		Posed and comment on IT	High	La Diseasant	2024-12-20	× ×	Paulaw select and compared on the		
		Bravicioning of actuarts ci	Normal	Not Started	2025-01-14		Page at 2x100G LP4 link to SLIPE up	+	
		Interconnect 2x OE2 fibro	Normal	Not Started			Request 2x 1000 Live link to Sold in	Apps	
		Pequert number recourse	High	Not Started			Paquet		
		11 transport on campus	Low	Not Started			nequest		
		Er transport on campus	2011	Not Started			2x OS2 fiber pair from DEB1 to DUB		
		L2 transport for CSIab net	High	Not Started	2025-02-XX		Transport of vlans from a physical (1		
		Coordination of RFC1918	High	Not Started	2025-01-XX		Ensure no numbering collisions betw		
		Rack and PDUs in DEB1	High	in progress	2025-02-XX				
		jumphost for commissioni	Normal	Not Started	2025-02-XX		Jumphost with linux (Debian or Alm This prevents commisioning/admin		
		Ability to create per-stude	Normal	Not Started	2025-05-XX				
		Re-centralise SURF SRAM	Normal	Not Started	2025-05-XX		Move administration and config of 9		
		Define incident response	Normal	Not Started	2025-02-XX		Define an incident response and har		
		Define process to connect	Normal	Not Started	2025-03-XX		Define a process for connecting CSI		
		Define which collaboratio	Normal	Not Started	2025-02-XX				
		Coordinate cable numberi	High	Not Started	2025-01-XX				Page 1 of
		Coordinate dot1q number	High	Not Started	2025-01-XX		Reserve 50 consecutive numbers fro	m the dot1g rar	nge (2.40
		Define process for central	Low	Not Started	2025-09-01		Users (esp. students) will likely file su	upport tickets or	r incident:
		Arrange periodic technica	High	Not Started	2025-01-XX		Organise periodic progress and tech	nical meetings	
		Arrange for a project man	High	In Progress	2025-01-09		Get a project manager to guide the	implementation	and com



#### Maastricht University | DACS

Tii

₽\$

Molowe

P Activity

Chat

Calendar

ij

C

Calls

...

 $\pm$ 

Apps

FS ....

#### DACS FSE Computing Sciences Lab as a Fieldlab Demonstrator 20

₽

P Activity

 $\sim$ 

5~

## More discussion time ...

David Groep david.groep@maastrichtuniversity.nl https://www.nikhef.nl/~davidg/presentations/ bttps://orcid.org/0000-0003-1026-6606



Maastricht University | Department of Advanced Computing Sciences [ 😋 📴



#### AS1104

#### a pretty standard research lab network



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

#### Maastricht University | DACS

DACS FSE Computing Sciences Lab as a Fieldlab Demonstrator 22

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



Storage services

OT. controls, and

engineering

enclave

#### Labs and more labs!





## Subset of Eli Dart's ScienceDMZ slides

Eli Dart's NASA talk (2016): <u>https://edu.nl/a7efv</u> See also <u>https://doi.org/10.3233/SPR-140382</u>