

# Distributed Supercomputing for climate modeling?

Dr. Jason Maassen  
Technology Lead Efficient Computing  
Netherlands eScience Center

netherlands

eScience center

by SURF & NWO



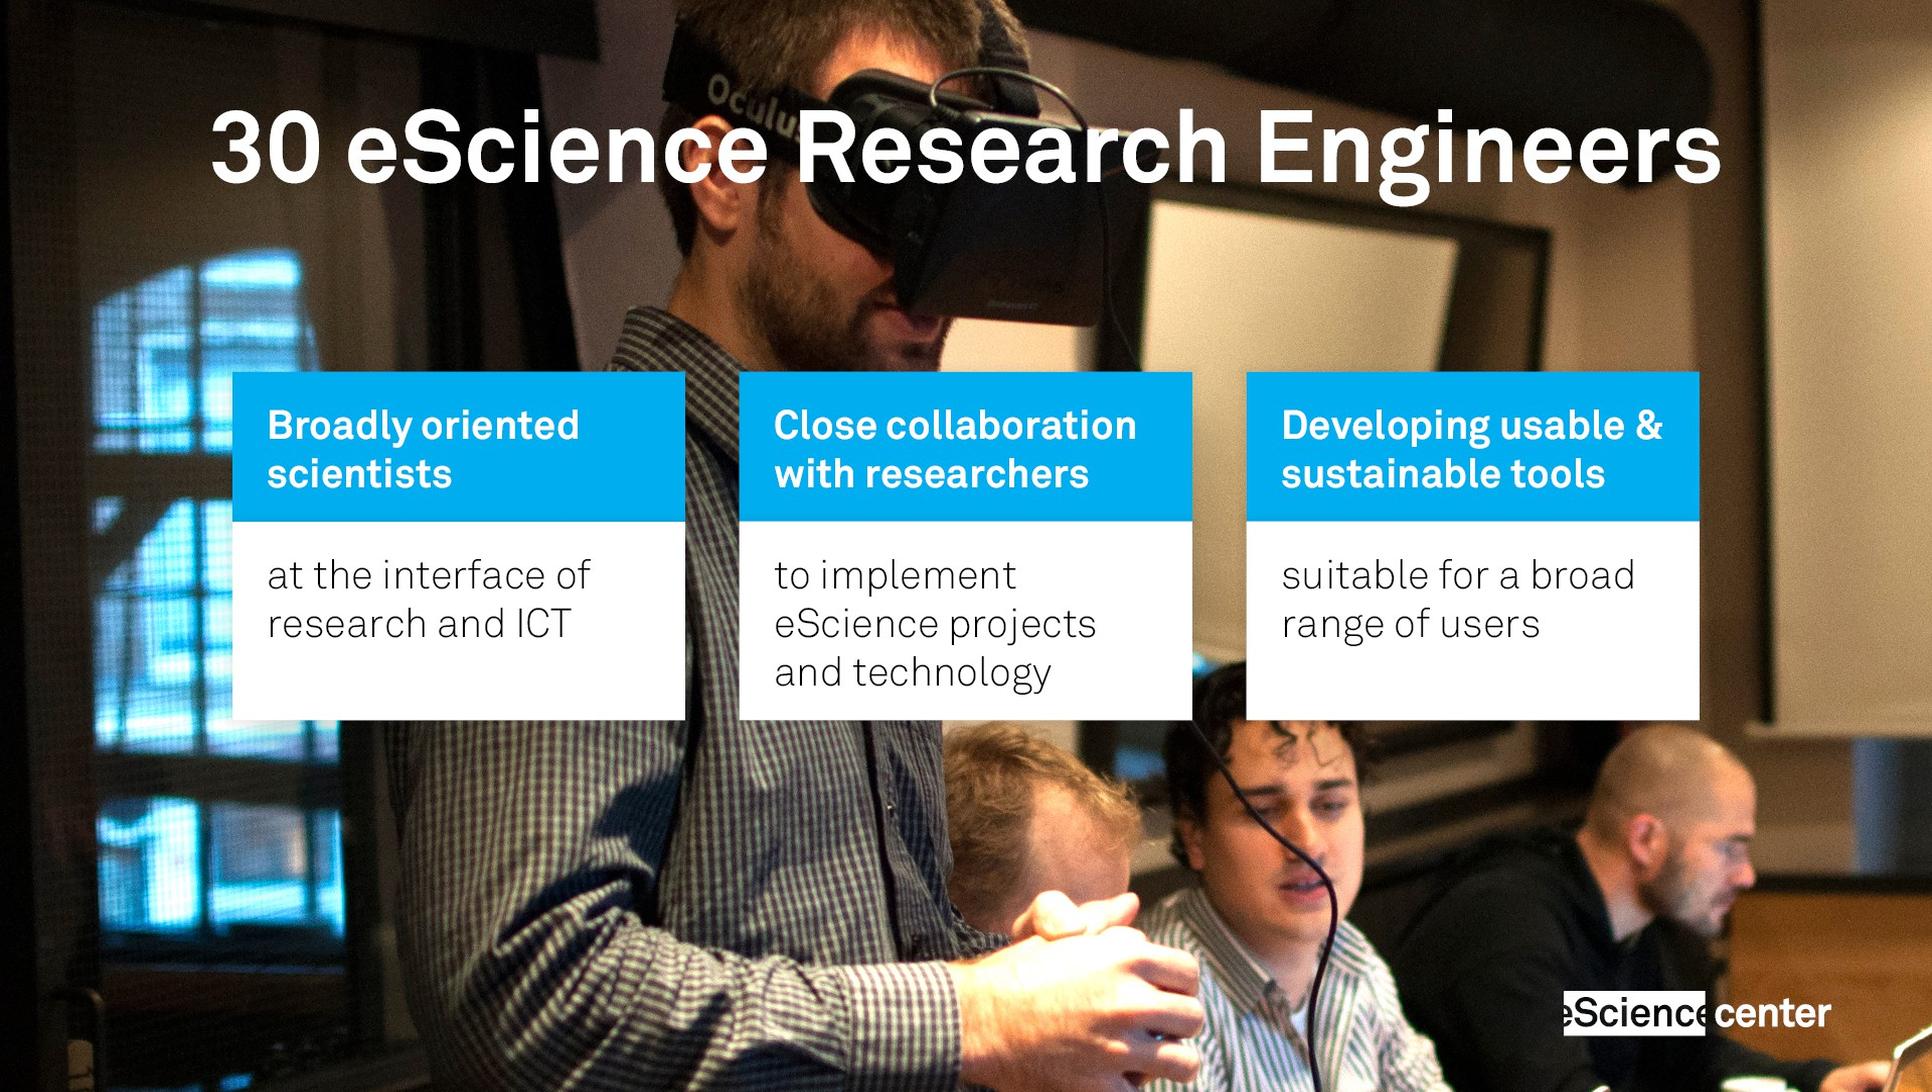
# Netherlands eScience Center



By Frits Ahlefeldt

**Bridging the gap between research and e-infrastructure**

# 30 eScience Research Engineers



**Broadly oriented  
scientists**

at the interface of  
research and ICT

**Close collaboration  
with researchers**

to implement  
eScience projects  
and technology

**Developing usable &  
sustainable tools**

suitable for a broad  
range of users

# Background: eSalsa Project

(2012-2015)

Predict **local** sea-level changes  
(caused by climate change) with  
an unprecedented level of detail.

*Utrecht University (IMAU)*

Prof. Henk Dijkstra (Ocean Modeling)

Dr. Sandra Brunabend

*VU University (CS-HPDC)*

Prof. Bal (GPU computing)

Dr. Ben van Werkhoven

*Netherlands eScience Center*

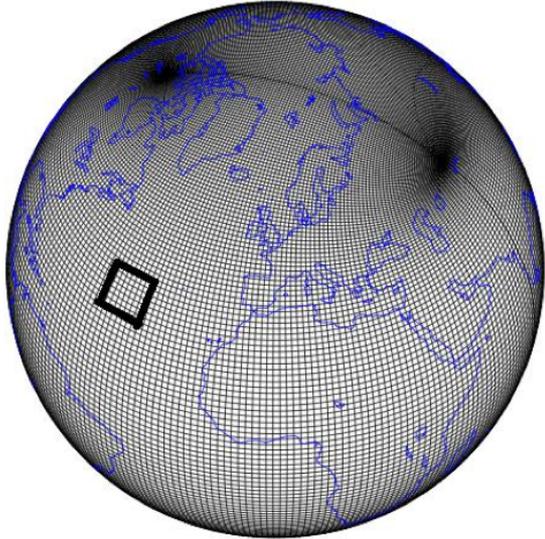
Dr. Jason Maassen (Distributed Computing)

Drs. Maarten van Meersbergen (Visualization)



Source: Actueel Hoogtebestand Nederland

# Typical Supercomputer application



— Source: Los Alamos National Laboratory

POP at  $0.1^\circ$  resolution (10x10 km)

$3600 \times 2400 \times 42 \approx 360$  million grid points



Catresius Supercomputer (source: SURFSara)

Single 100 year simulation takes **weeks** on 2~4K cores. Need ensemble of 10's.



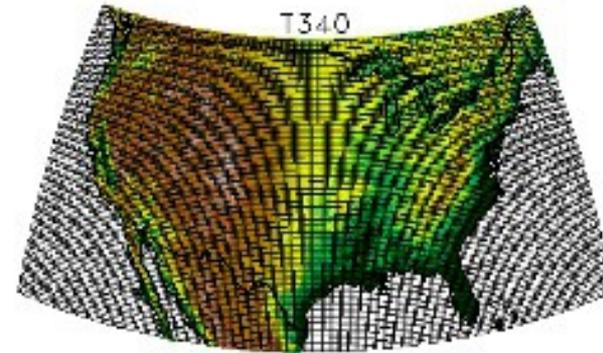
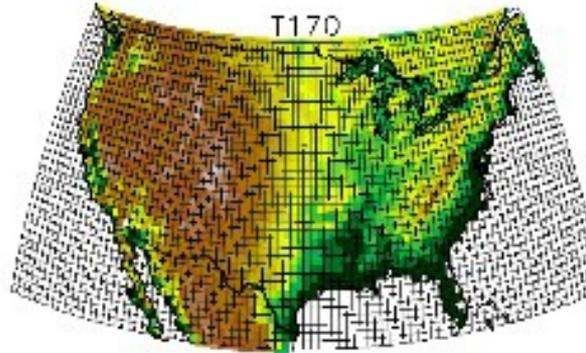
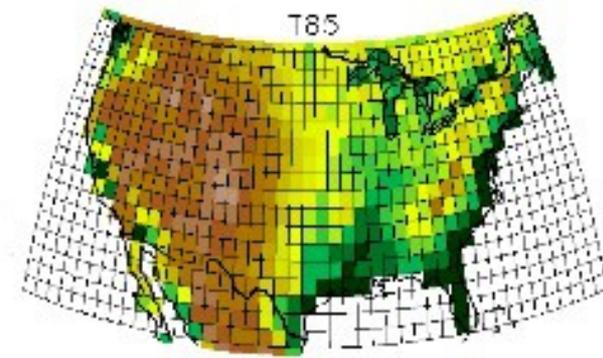
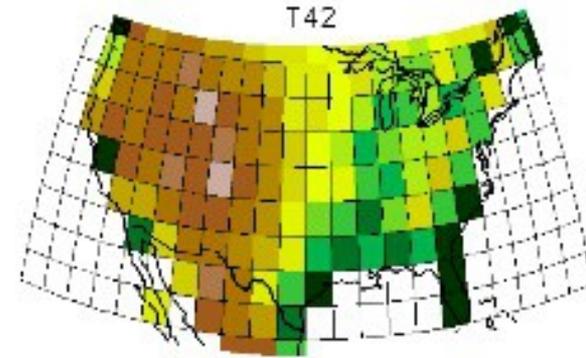
# Quest for Higher Resolution

Scientists want higher model resolution (this gives more detailed results).

Ultimate goal:

0.01° resolution (1x1 km)  
**fully eddy resolving**

100~250x increase in  
compute time!



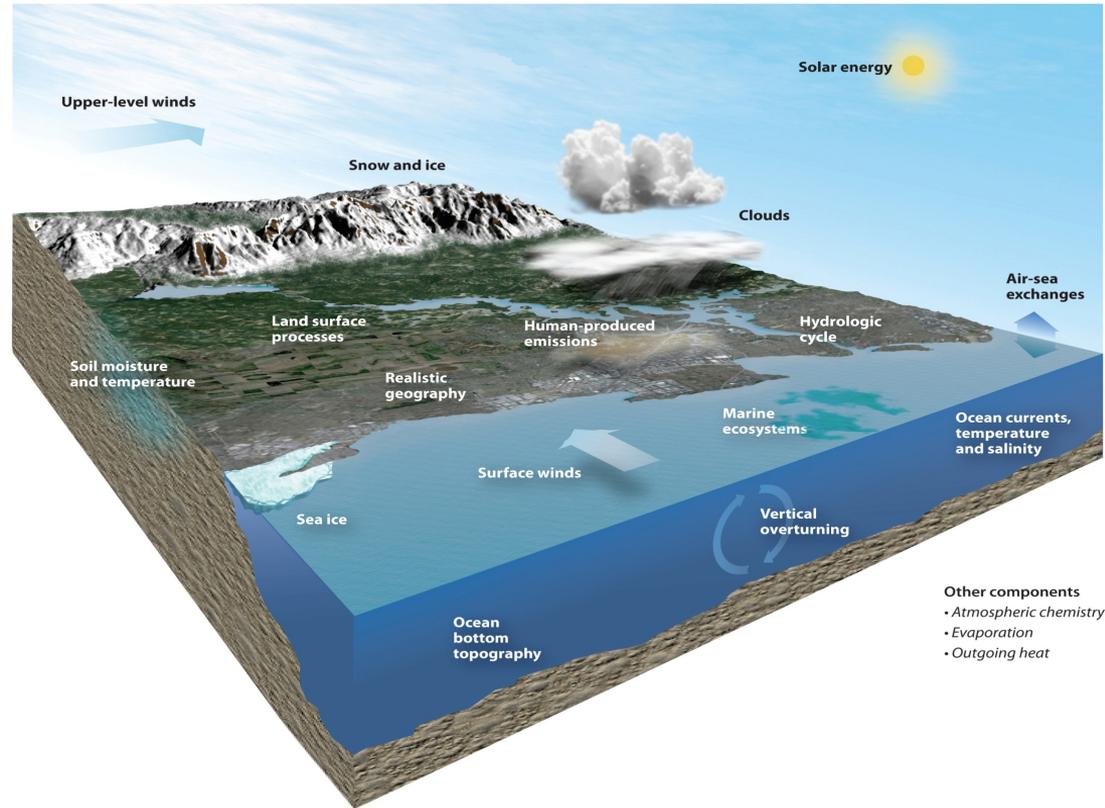
Source: Illustration courtesy Warren Washington, NCAR

# Earth System Modeling

Climate is more than just ocean!

Scientists also want to **combine** ocean, land, ice and atmosphere models (to simulate the interactions)

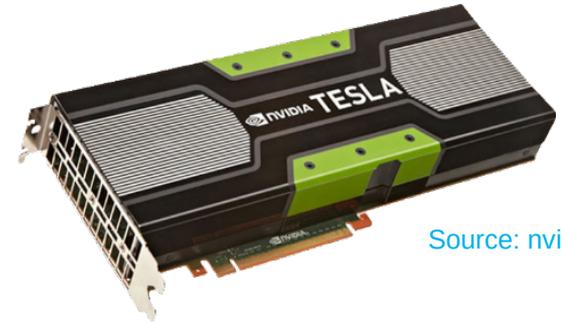
2~3x increase in compute time.



# Where do we get the compute power?

## Accelerators

- for more compute power per node
- success, but slow progress!
- model implementations are diverging  
(CPU, GPU, Xeon Phi)



Source: nvidia

## Distributed Computing

- for more overall compute power
- allows use of heterogeneous system  
**(focus of this presentation)**



Source: intel





Spinoff of the eSalsa project at the Netherlands eScience Center

Our a submission to **Enlighten Your Research global** competition in 2013

.... and we where one of the winners!

# Enlighten Your Research Global 2013/14

Top 500 #10  
2.9 PFlop/s



**SUPERMUC (GER)**

10G

Top 500 #184  
223 Tflop/s



**CARTESIUS (NLD)**

10G

114 TFlop/s



**EMERALD (UK)**

GPUs

GPUs

10G

Extremely ambitious goal:  
combine 2 or 3 supercomputers  
from Europe and the US to run  
Earth System Model at 0.02°  
(2x2 km) resolution!

Top 500 #7  
5.1 PFlop/s

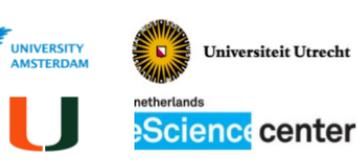


**STAMPEDE (USA)**

Xeon Phi



# Requires a lot of international cooperation!



19 project members:

*5 climate researchers*

*9 computer scientist*

*5 system and network specialist*

4 countries (NL, US, UK, DE)

4 research institutes

6 universities

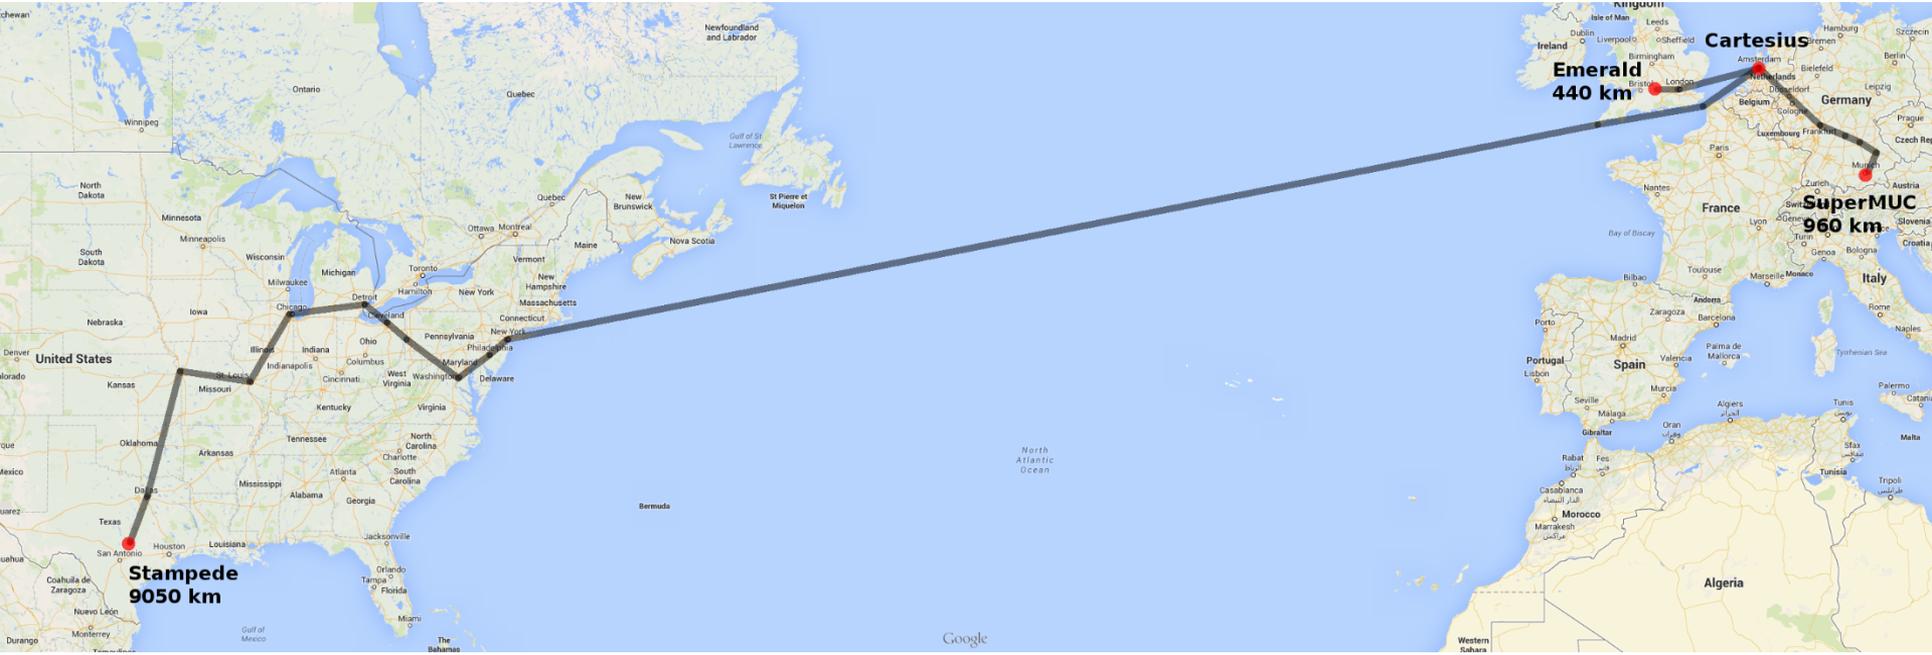
4 supercomputing centers

PRACE & XSEDE

... and this does not even include the NRENs yet!



# World wide 10G network links



FUNET

janet

INTERNET<sup>2</sup>

SURF NET

NRENs provide the **essential** 10G network links between the sites.

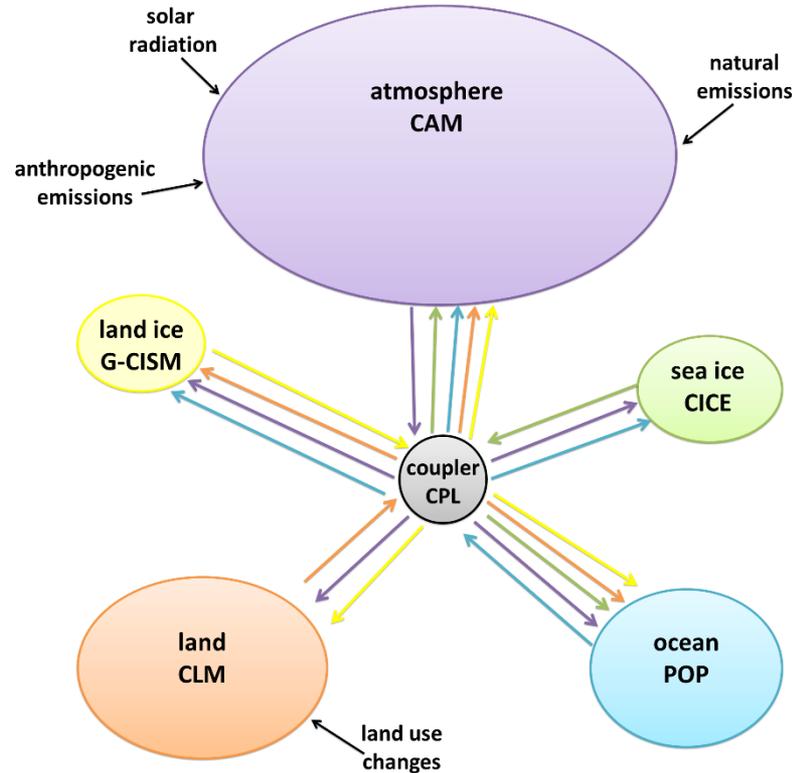


# CESM & model coupling

We use the **Community Earth System Model (CESM)**.

CESM combines 5 simulations each simulating a different aspect of the climate (ocean, atmosphere, land, etc).

Models regularly **exchange data** to model the influence on each other.



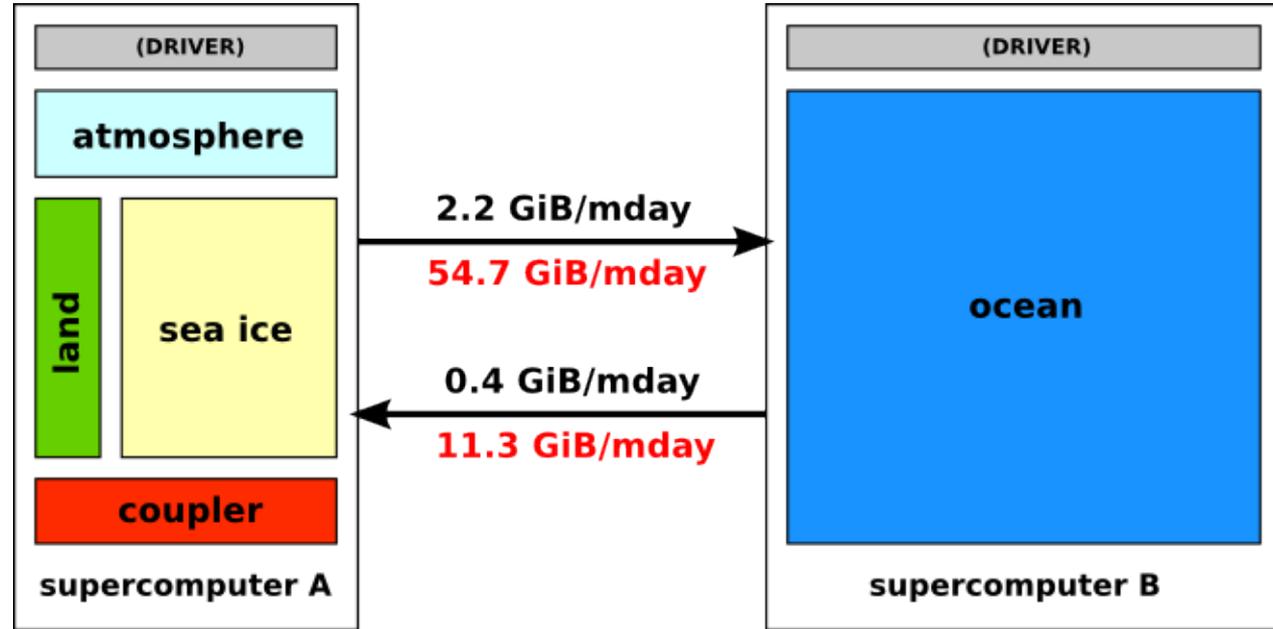
Source: [climatesight.org](http://climatesight.org)



# Distributed CESM

We can run the ocean model **separately!**

Coupling occurs 4x per simulated day (once every **~30 sec**) and **data volume** is limited.

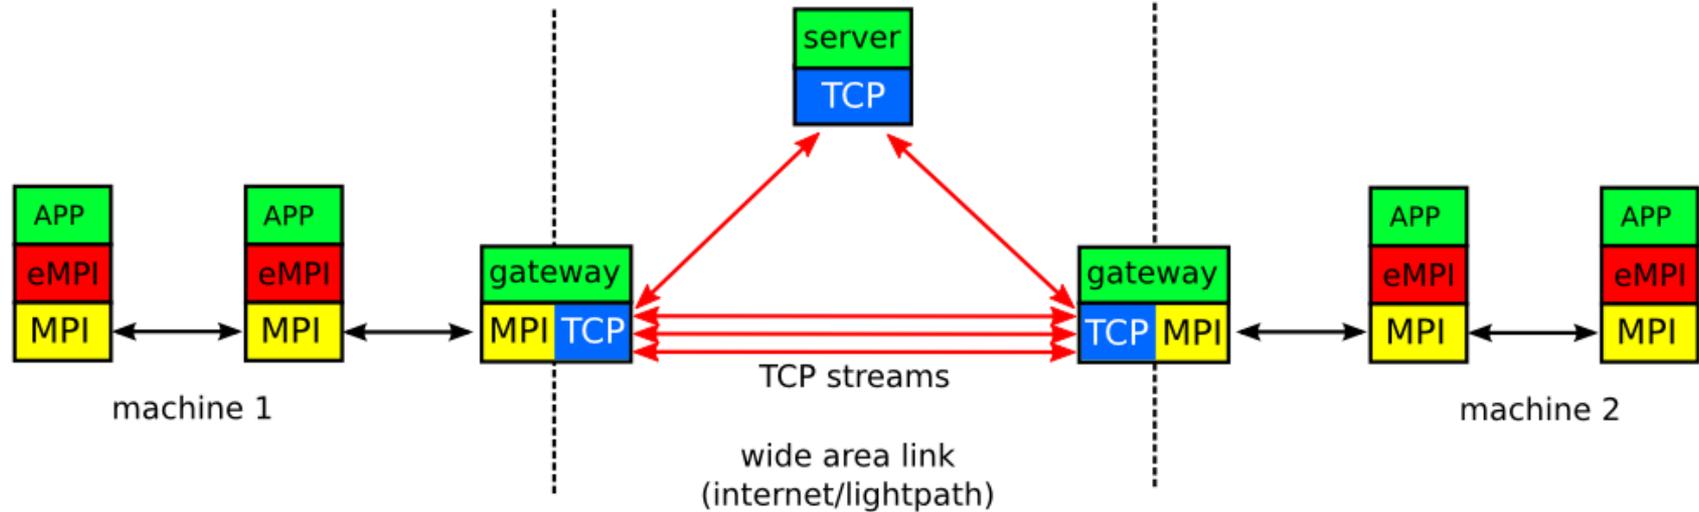


2.6 Gigabyte per exchange for 10 km resolution  
66 Gigabyte per exchange for 2 km resolution

**bursty traffic, not latency sensitive!**

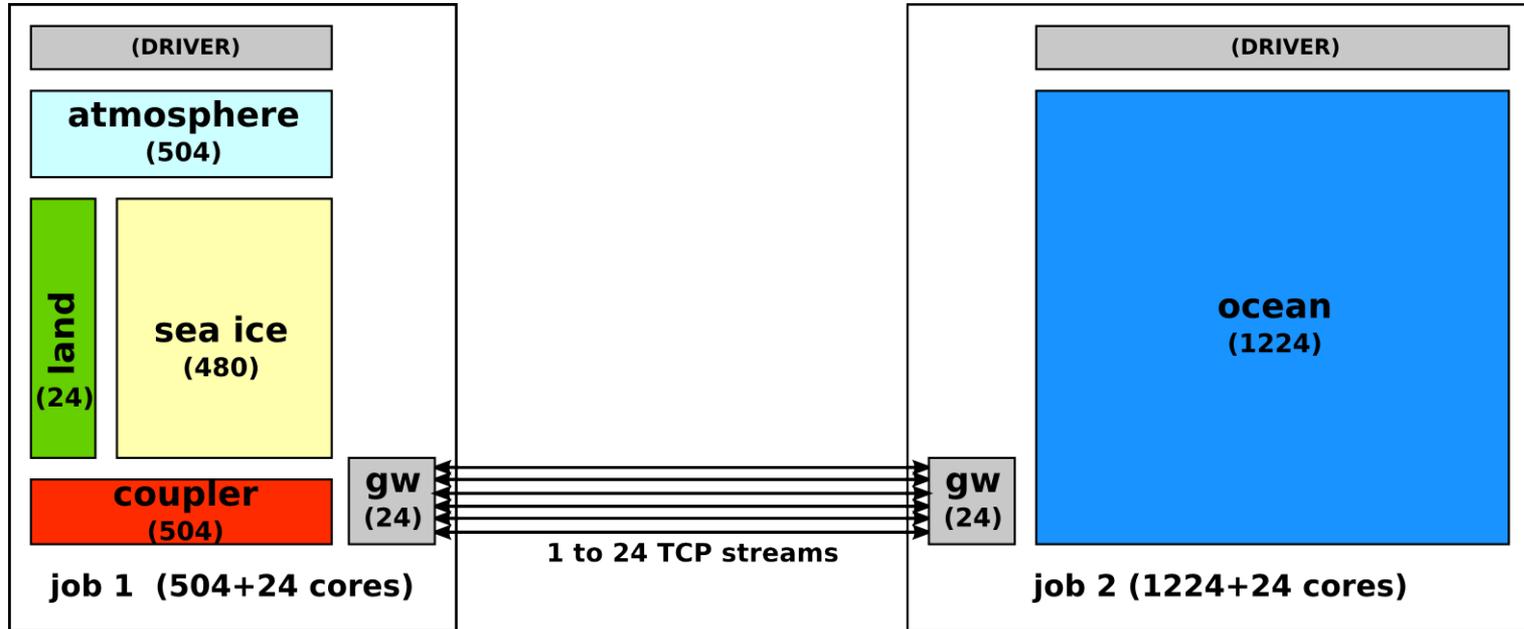


# Coupling sites using eSalsa-MPI



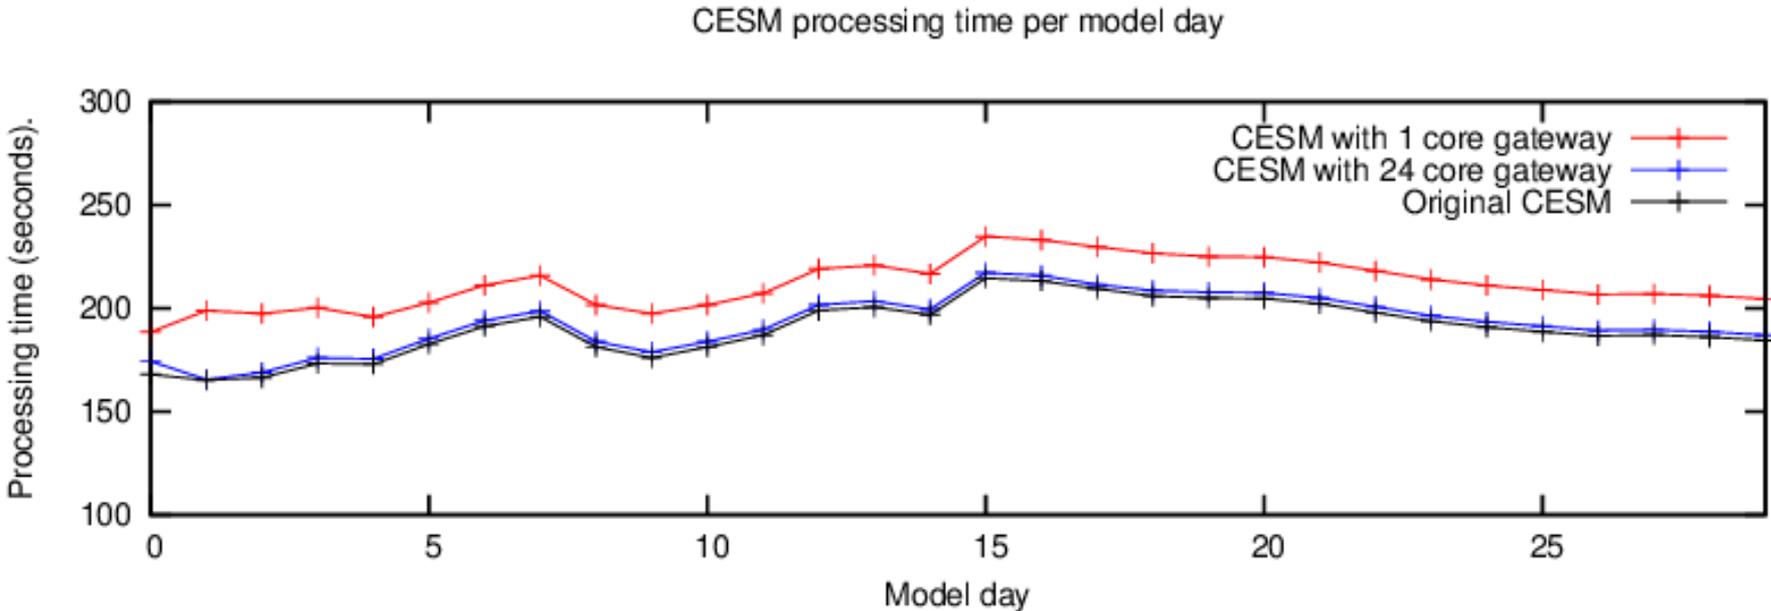
eSalsa-MPI is an **MPI wrapper library** that can combine two MPI runs on different supercomputers, and presents it as a **single set of processes** to the application. An extra **gateway node** is required at each site.

# Single site test: Two jobs on Cartesius



We combine two MPI jobs (504 and 1224 cores) by using two **gateway nodes**. All traffic between jobs goes through these gateways.

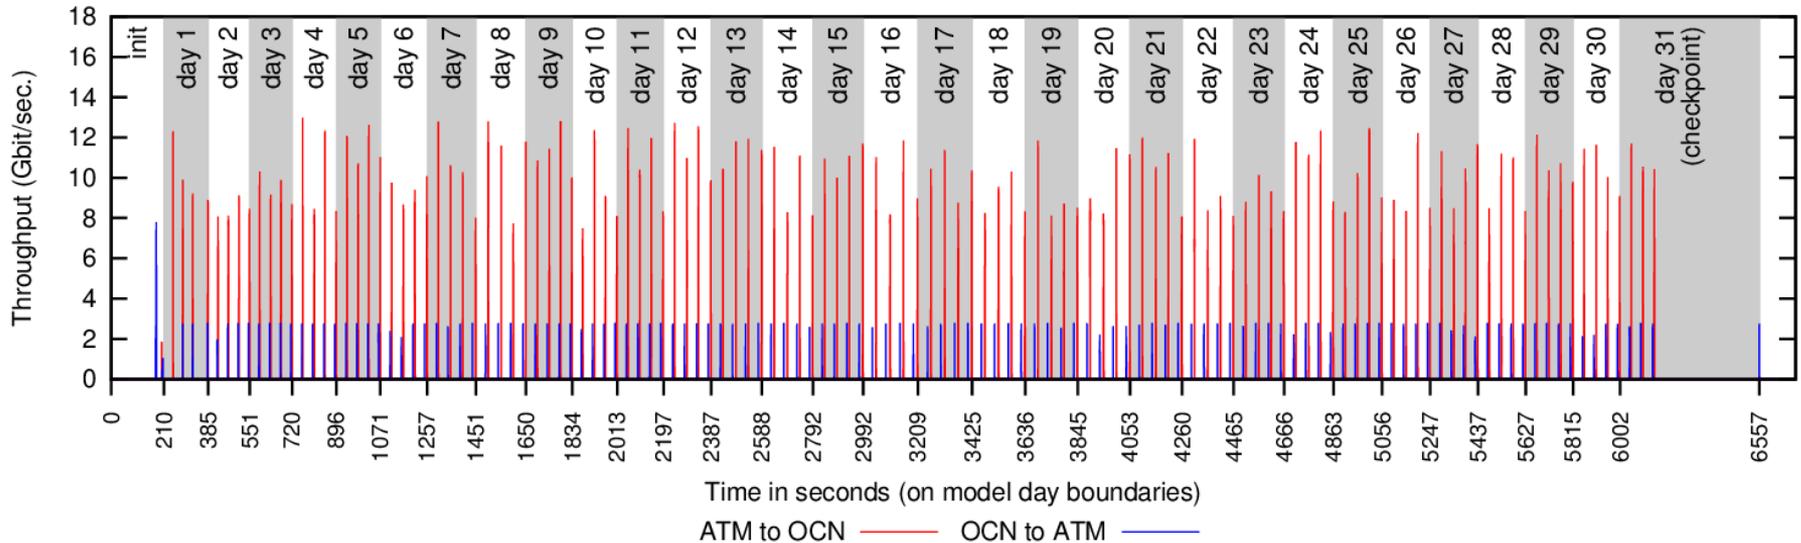
# Single site test: time per model day



Two separate jobs (Cartesius 504+1224 cores) connected via eSalsa-MPI  
Overhead 10%-1.5% dep. on number of TCP streams between gateways

# Single site test: gateway throughput

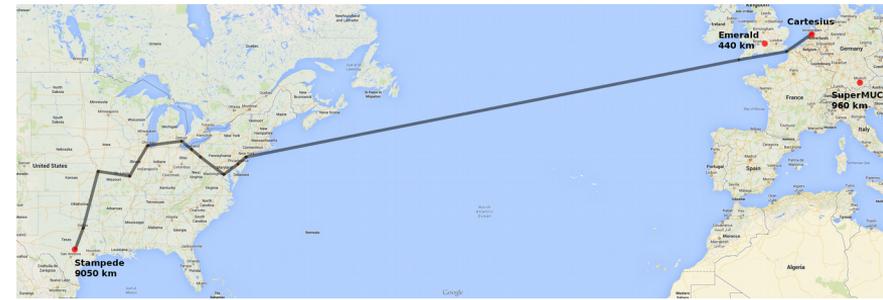
Gateway transfer rates during 1 model month of CESM f05 t12 on 504 + 1224 application cores + 2x24 gateways



Bursty traffic pattern clearly visible. ATM to OCN uses about 8-12 Gbit/s, OCN to ATM uses about 2.5 Gbit/s (sends less data).

# Lightpath setup

Cartesius (NL) – Stampede (US)



Atlantic lightpath was up pretty quickly

Last mile took longer, as TACC wasn't connected at 10G yet

Setup of gateway node was quite difficult!

Unclear who to contact, unclear how to set up hardware and network

Colliding IP ranges complicate matters

Finally a Stampede service node was configured to use a Cartesius IP adres!

Initial run worked nicely, but then we had a **break in the cable!**

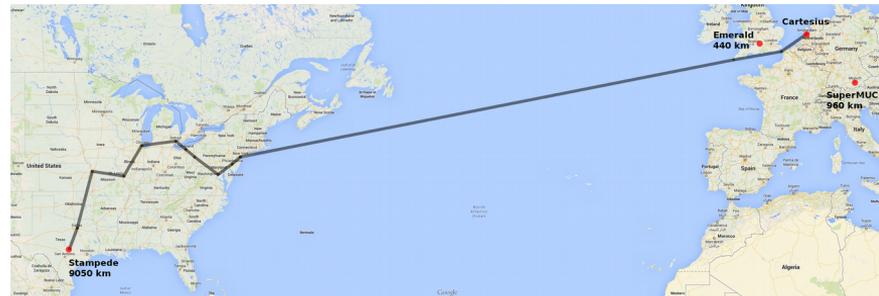
Lost the Stampede node while repairs were taking place

Replaced node by VM, which had unreliable performance



# Lightpath benchmarks

Cartesius (NL) – Stampede (US)



Low level performance benchmarks:

RTT: 126 ms

TCP throughput:

32MB buffer, 1 stream : 1.57 Gbit/s

32MB buffer, 3 or more streams : 4.52 Gbit/s

UDP throughput:

32K packets : 9.7 Gbit/s (almost zero packet loss)

UDT throughput:

9K packets, 150 MB buffer : 9.5 Gbit/s (no packet loss)



# Lightpath setup

Cartesius (NL) – SuperMUC (DE)

Lightpath already available

Re-use the existing **PRACE** link

Setup of gateway node again quite difficult!

Took a bit of convincing to get support

Network hardware issues limited the performance initially

Required a change of the network setup at SuperMUC

Jumbo frames do not seem to work on this link – still unclear why...

But we did end up with a working setup!



# Lightpath benchmarks

Cartesius (NL) – SuperMUC (DE)



Low level performance benchmarks:

RTT: 14.9 ms

TCP throughput:

1 stream : 5.3 Gbit/s

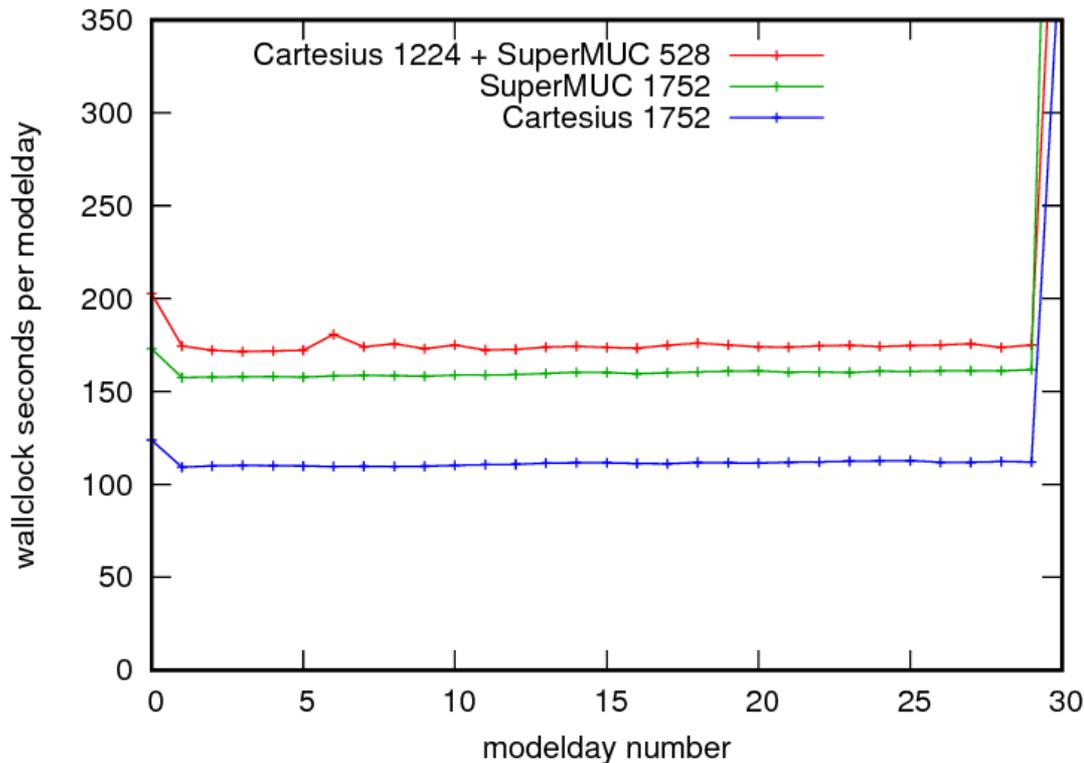
Multiple TCP streams does not help due to relatively low latency

UDP / UDT not usable due to lack of jumbo frames



# Application run

## Cartesius (NL) – SuperMUC (DE)



**It runs!**

Performance tuning needed:

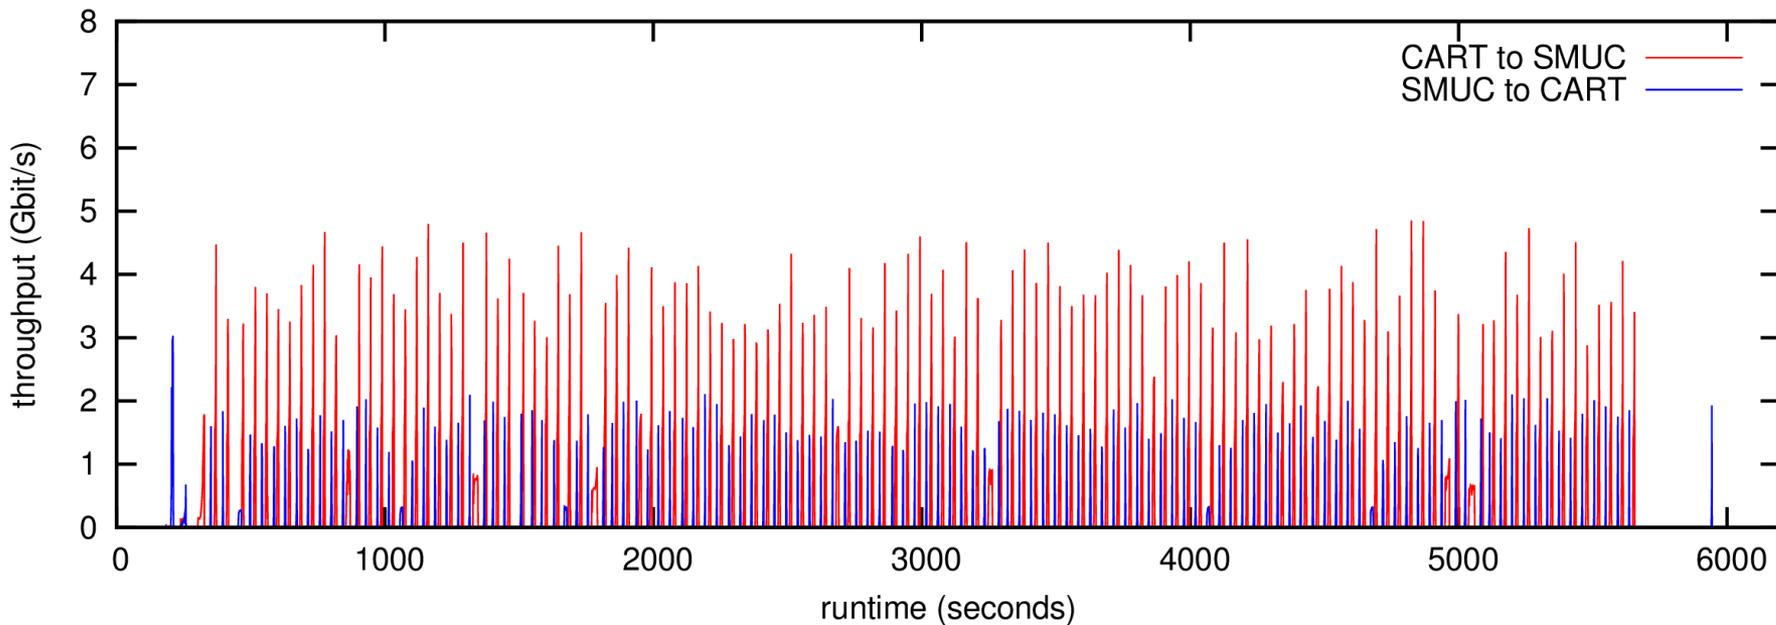
Already 44% performance difference between Cartesius-only and SuperMUC-only.

Distributed adds extra 10%



# Gateway throughput

Cartesius (NL) – SuperMUC (DE)

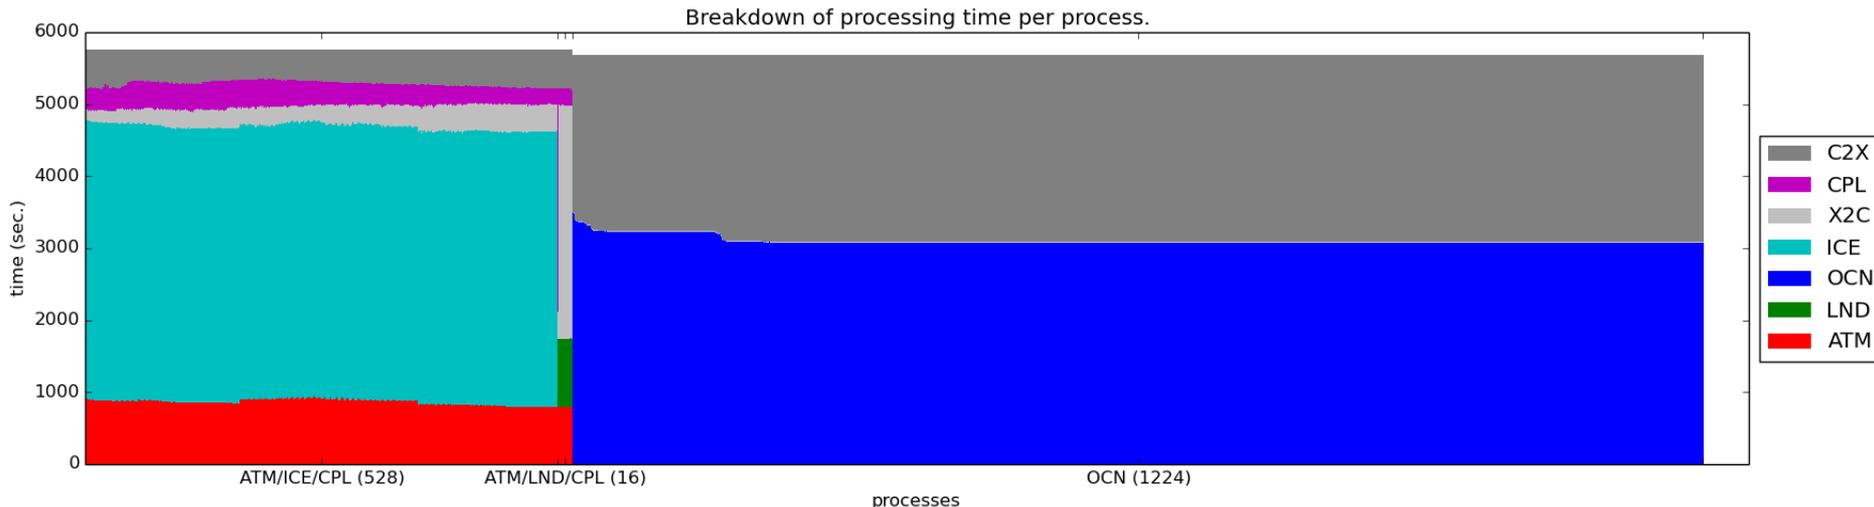


Wide area network behavior seem fine...



# Load balance

Cartesius (NL) – SuperMUC (DE)

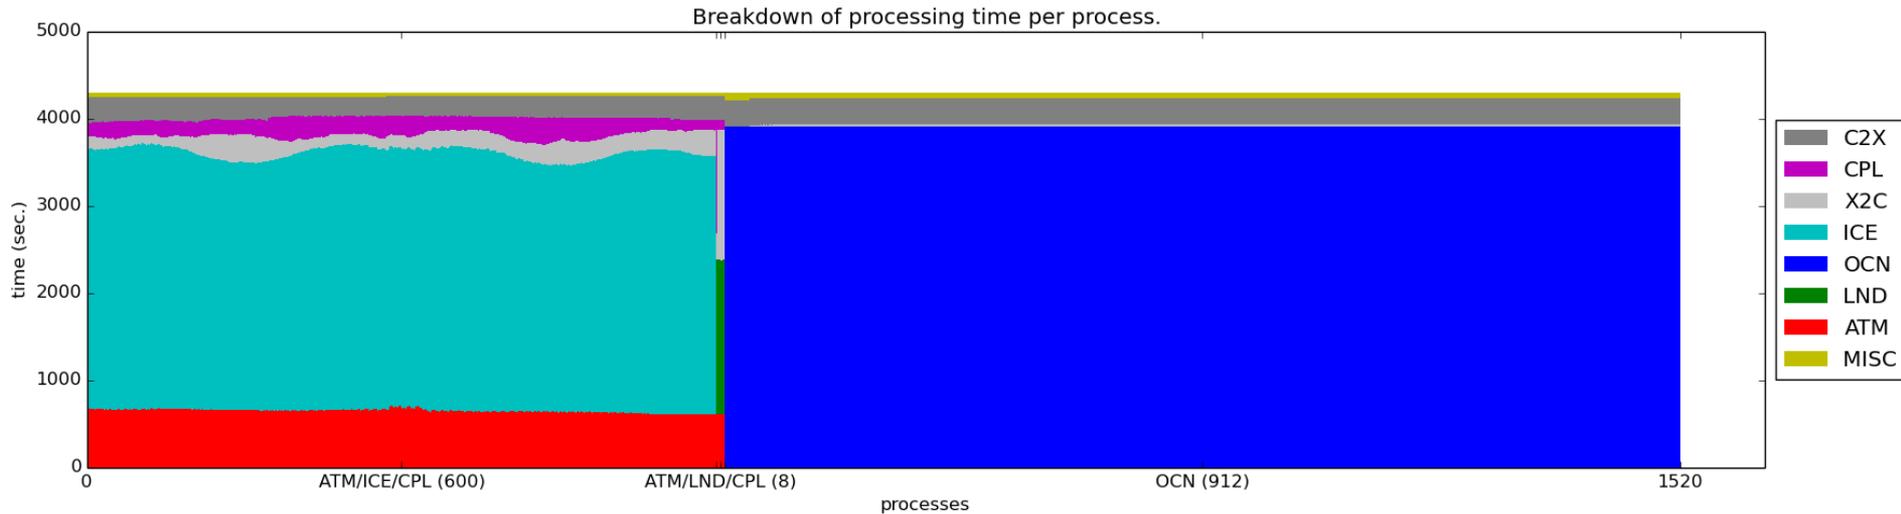


Desparately needs better load balancing!  
Ocean is running too fast (too many cores)



# Fixed load balance

Cartesius (NL) – SuperMUC (DE)

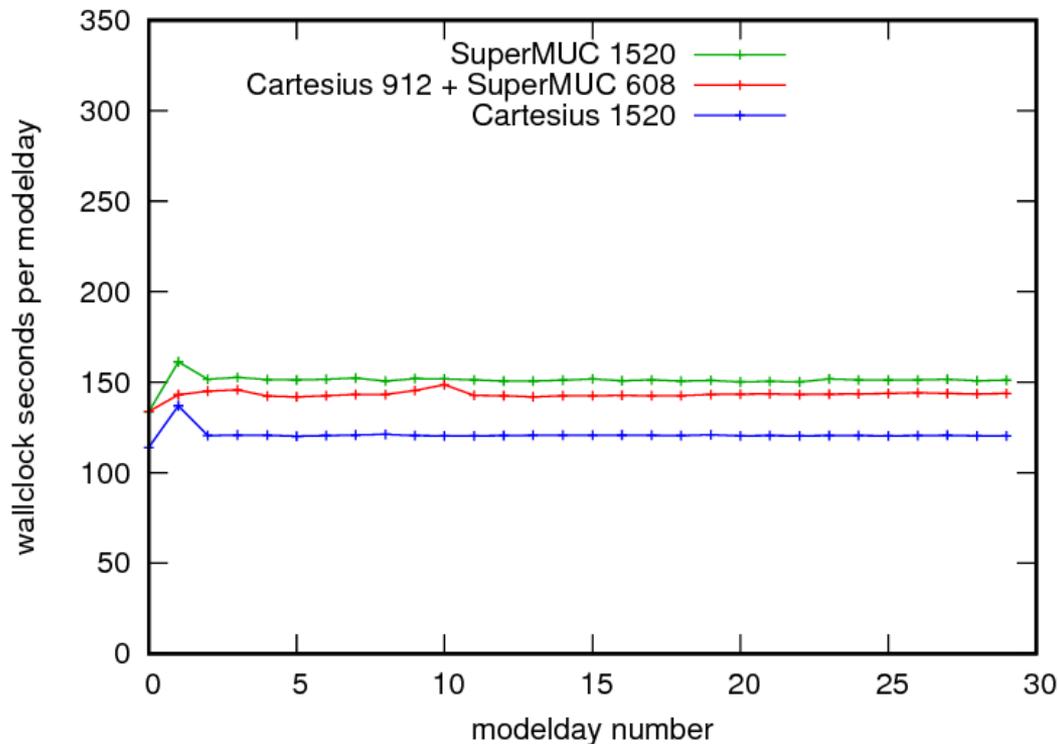


Less cores on Cartesius (-312), more on SuperMUC (+80)  
Much better load balance between the systems!



# Application run (retry)

Cartesius (NL) – SuperMUC (DE)



25% performance difference between Cartesius-only and SuperMUC-only.

Distributed in between the two!!



# .... and then we ran out of time

Working on this (on-and-off) for two years

Parent project had finished long before

Core hours started disappearing

People started disappearing

Machines started disappearing (Emerald was gone)



# What *did* we do

**Distributed MPI:** we managed to get the eSalsa MPI up and running with reasonable performance

**Small scale experiments:** we managed to run small scale experiments on combinations of 2 machines O(2K) cores at 10km resolution.

**Published a paper** on our experiences:

*On the complexities of utilizing large-scale lightpath-connected distributed cyberinfrastructure* (DOI: 10.1002/cpe.3853)



# What *didn't* we do

**US link:** we did not managed to get the US link stable for long enough to run real experiments in that direction

**Large scale experiments:** we did not managed to run large scale high resolution experiments on O(50K) cores and 1~2 km resolution

**Heterogeneous runs:** we did not managed to run experiments with heterogeneous combinations of machines (CPU/GPU/XeonPhi)

**3 site runs:** we did not managed to run experiments on more that 2 sites



# Lessons learned

## It works!

We can run an climate model on a distributed system with acceptable performance!

But tuning the wide-area TCP links is still an issue ?

Balancing the number of cores on the sites is black magic!

The **hardest part** (by several orders on magnitude) is the **last 10 feet!**



# Lessons learned

Even if you have a 9000 km lightpath up and running with ends up in the datacenter, actually getting it connected to the supercomputer is **hard!**

Weird usecase

Initially goes against the gut feeling of the sysadmins and network admins

No standard way to do it

Need to bridge networks (requires some gateway machine) and tuning

System and network admins were **very helpful**, but support time still became an issue when quick fixes were not enough!



# Possible solutions

Providing infrastructure (network & core hours) is not enough. Need to add **person hours** for setup and support

Provide standardized way to **connect running simulations**, which is different from the existing storage-to-storage data transfers

**Co-scheduling** would be great, although we managed amazingly well without



# Many thanks to:

Henk Dijkstra  
Henri Bal  
Frank Seinstra  
Michael Kliphuis  
Sandra Brunnabend  
Ben van Werkhoven  
Maarten van Meersbergen  
Sander Boele  
Alexander van den Hil

Jules Wolfrat  
David Wallom  
Nick Hill  
Dieter Kranzlmuller  
Christian Grimm  
Shantenu Jha  
Frank Bryan  
Benjamin Kirtman  
Edward Moynihan  
Dinesh Ganpathi

Mark Santcroos  
Yaakoub El Khamra  
Sylvia Kuijpers  
Migiel de Vos  
Patrick Storm  
Huub Stoffers  
Nathaniel Mendoza  
Reinhold Bader  
Gerben van Malenstein



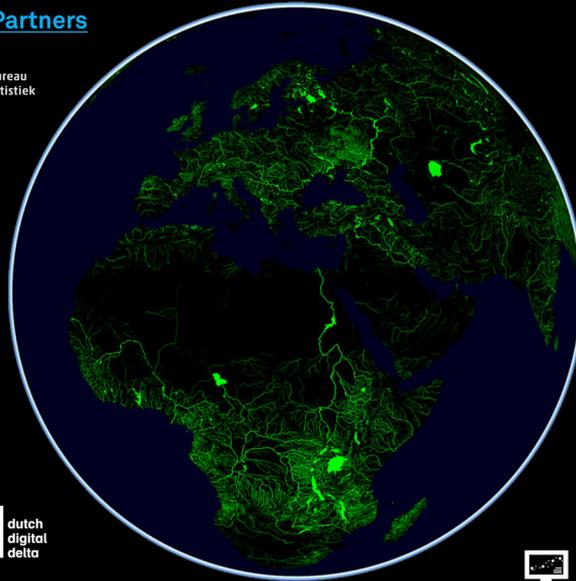
4<sup>th</sup> National eScience Symposium

# Science in a Digital World

## Partners



Leiden Centre of Data Science



## Sessions



Social Data



Astronomy



Smart Energy



Sports & eHealth



Data Science