2025 France – Japan ITER Data Transfer Tests at 100 Gbps

Denis Stepanov, ITER Organization
With contributions from ITER, QST, NIFS, NII and GEANT teams

GNA-G Community VCs Q4 2025

Disclaimer: The views and opinions expressed herein do not necessarily reflect those of the ITER Organization

iter
china eu india japan korea russia usa

ITER AT A GLANCE

- A next-generation tokamak-type **fusion machine** under assembly in Cadarache, France
- Construction and operation is managed by the **ITER International Organization** established in 2007, supported by seven Members (China, EU, India, Japan, Russia, South Korea, United States) and 33 nations
- The machine is procured "**in-kind**", meaning each Member provides a piece and benefits from 100% of the knowledge obtained in the project
- Licensed as a **nuclear installation** under French law, with safety requirements comparable to those of the nuclear power plants
- Machine **assembly** started in 2020, **integrated commissioning** foreseen from 2033 and **start of research operation** in 2034

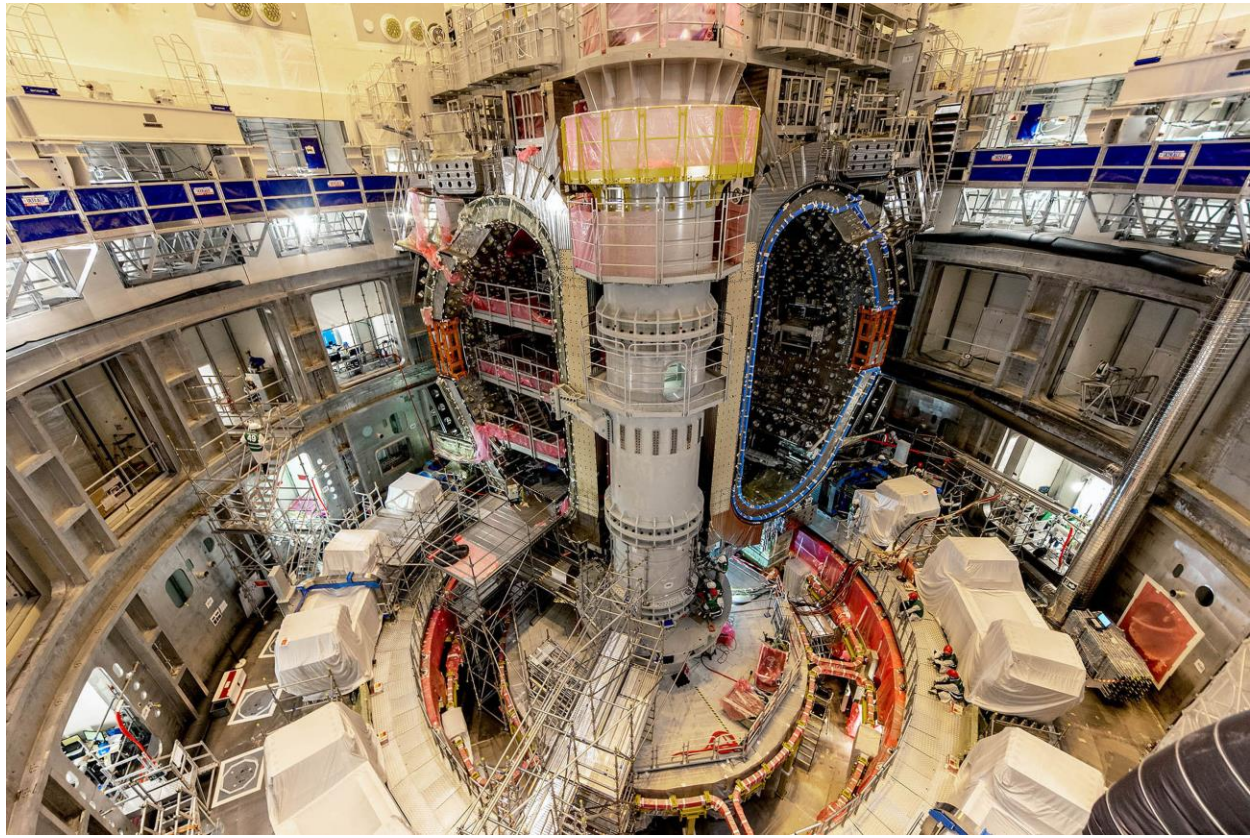

ITER REMOTE PARTICIPATION ARCHITECTURE

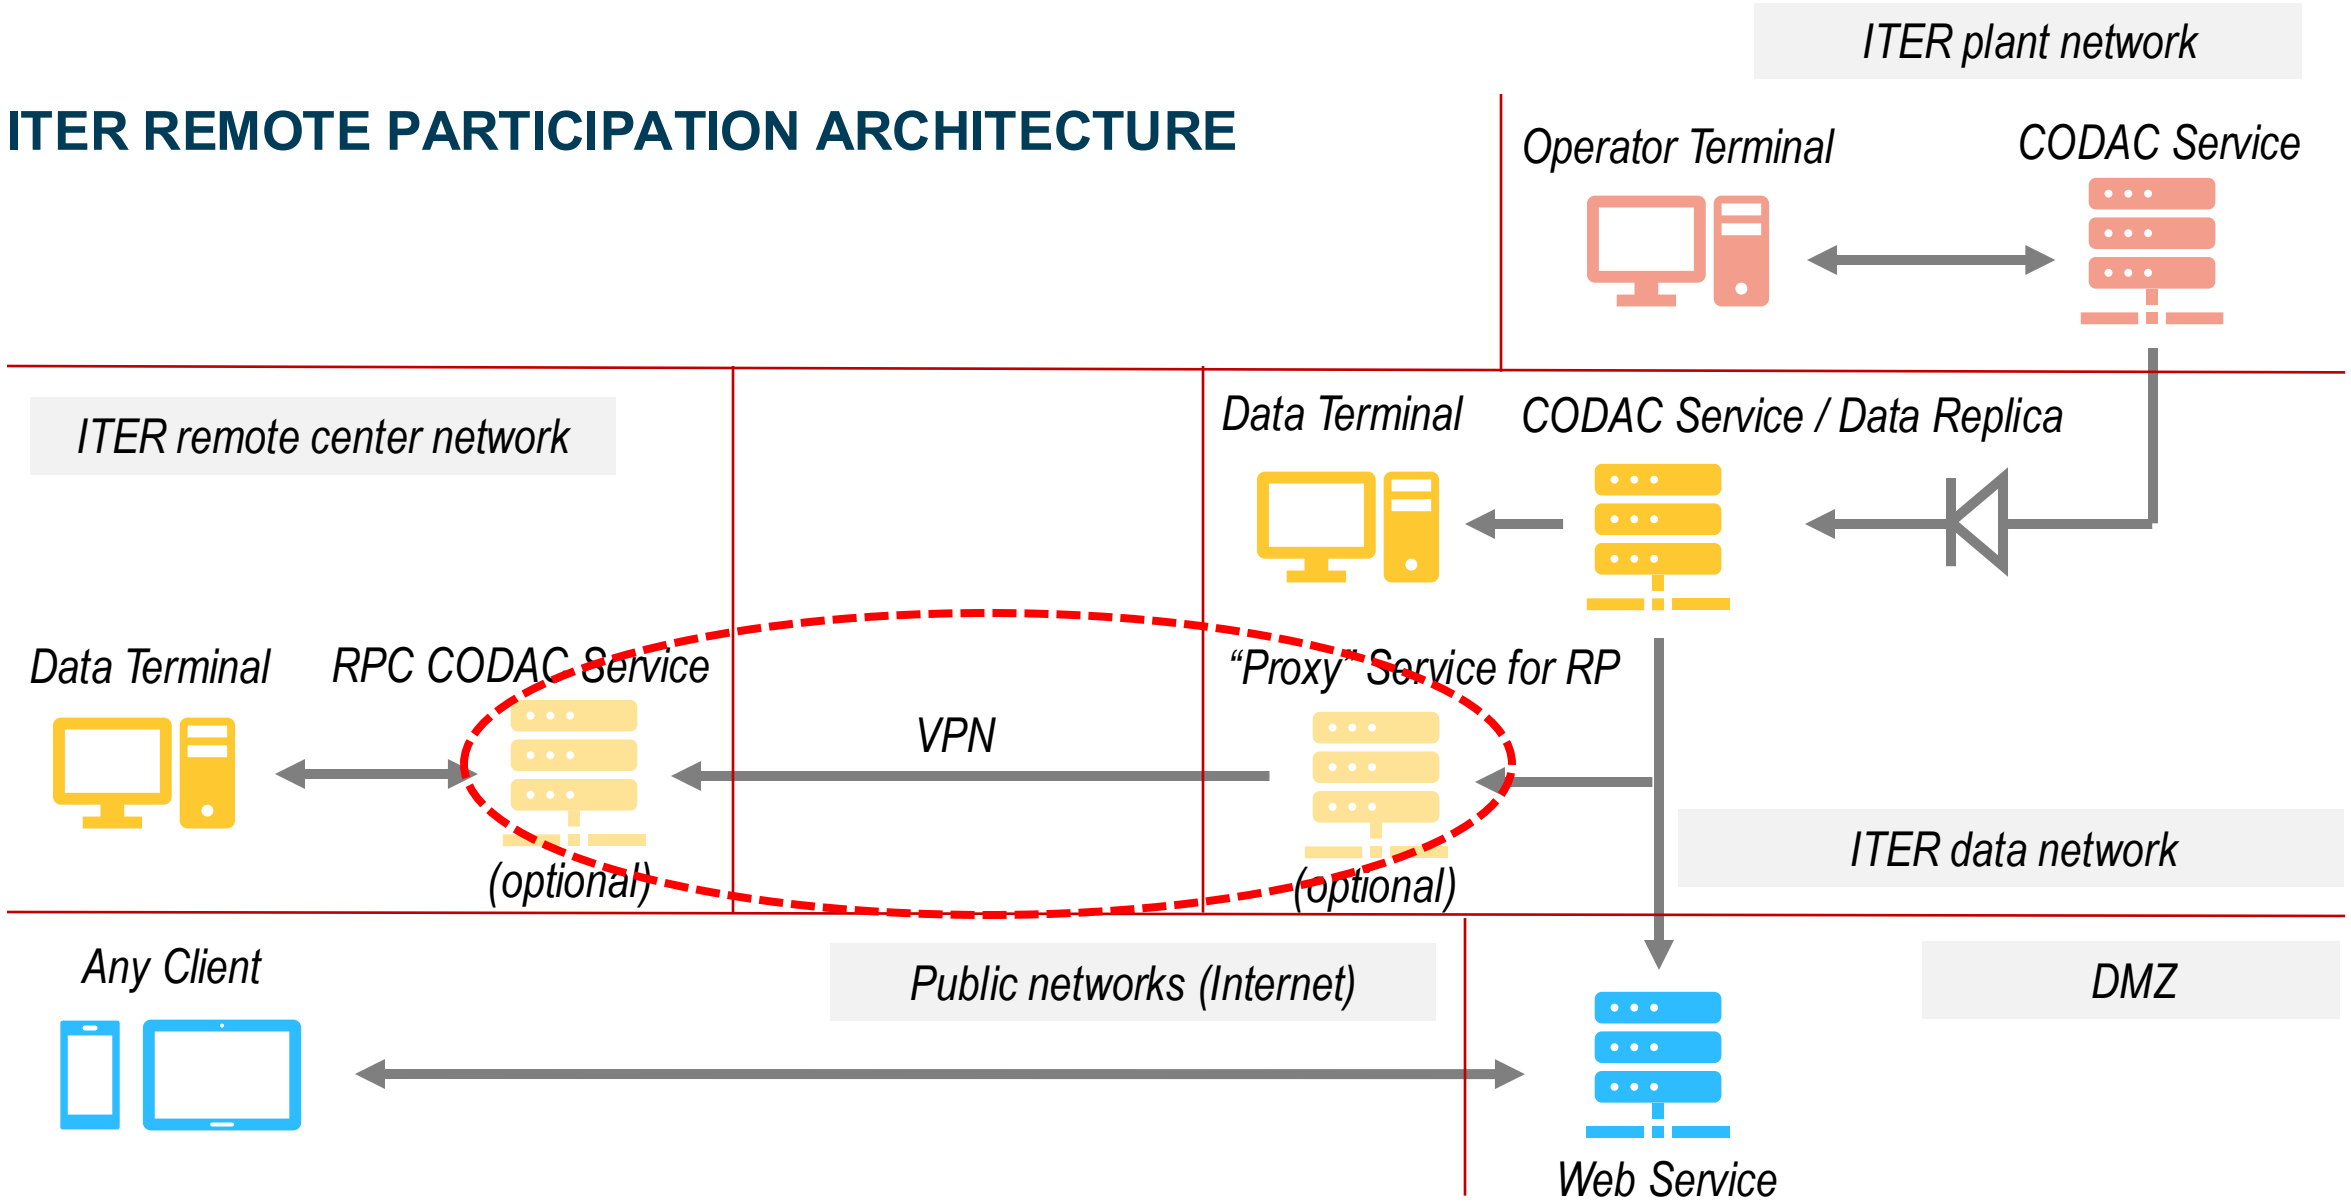

RP = Remote Participation
RPC = Remote Participation Center
CODAC = Conventional Control System
DMZ = De-Militarized Zone

ITER DATA TIERS (DRAFT)

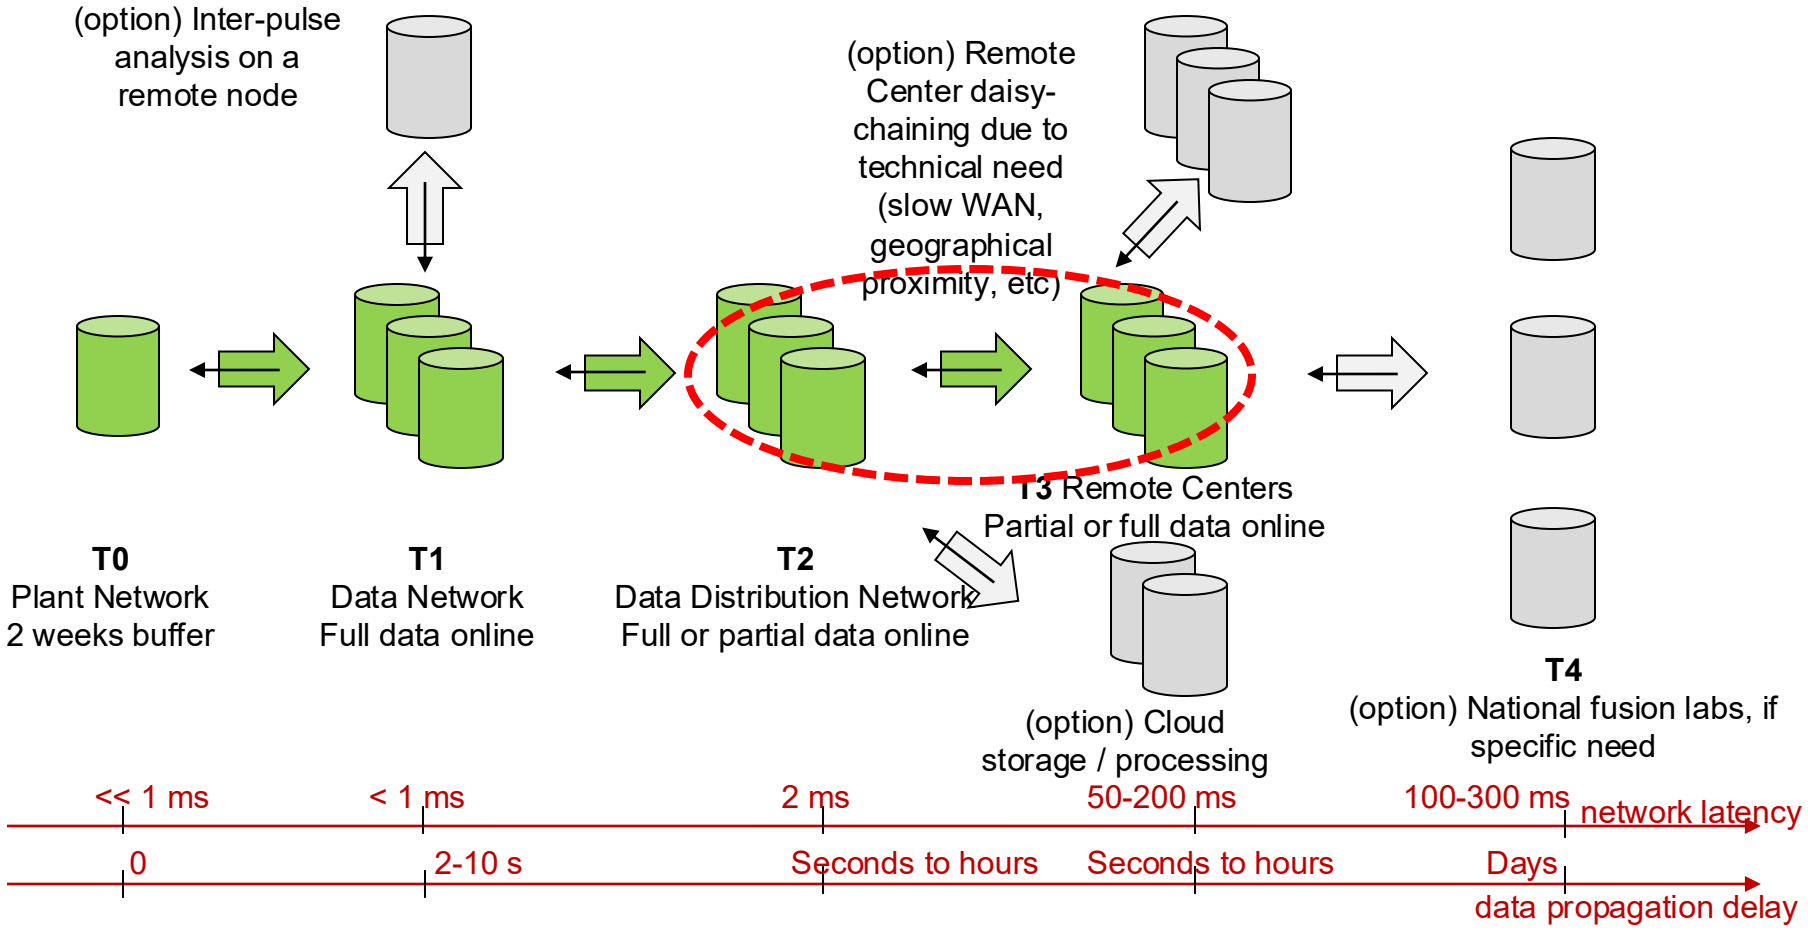

ITER DATA PRODUCTION NEEDS

- ITER is a pulsing device, with pulse duration ranging from one second to one hour
- ITER will be operated in three shifts, eight hours each
- Typical duty cycle would probably be around 25%
- ITER Control System is designed to support data rates of 2 to 50 GiB/s
- Real data production rates can vary significantly depending on the experiment configuration
- For this test we chose a reference Q=10 scenario with a 400 s pulse delivering max data of 50 GiB / s once per hour. This gives 20 TiB / h = 160 TiB (176 TB) per 8h shift
- In the case of full sync these data would need to be transferred at the rate of production (or faster) to avoid clogging

EXTERNAL CONNECTIVITY REQUIREMENTS

(2020 estimates)

Parameter	Recommended Value
Bandwidth IO → RPC, basic functions	100 Mbps
Bandwidth IO → RPC, select post-pulse data mirroring	1 Gbps
Bandwidth IO → RPC, full data mirroring (first plasma scope)	10 Gbps
Bandwidth RPC → IO, basic functions	100 Mbps
Bandwidth RPC → IO, submission of processed data	1 Gbps
Bandwidth IO → occasional RP user, observation mode	10 Mbps
Bandwidth occasional RP user → IO, collaboration mode	10 Mbps
Maximal network roundtrip latency (ping)	200 ms
Maximal jitter	30 ms
Maximal packet loss	0.0001%

"First plasma" is a beginning of operations. Full power would likely need one order of magnitude more ➔ **100 Gbps?**

IO = ITER Organization
RP = Remote Participation
RPC = Remote Participation Center

2016 TESTS

ITER to Japan at breakneck speed

https://www.iter.org/node/20687/iter-japan-breakneck-speed

10 OCT 2016 - R.A.

Print Read the latest published articles

When ITER starts running plasma shots, operators in the control room on site won't be the only ones processing the huge amount of data collected by the tokamak's diagnostics systems. Ten thousand kilometres from ITER, in Japan's Remote Experimentation Centre (REC), other operators will be crunching the same numbers. Problem: how to transfer data as fast as possible to a site half a world away?

" … From 30 August to 5 September, 50 terabytes/day were transferred from ITER to Rokkasho, Japan at an average speed of 7.9 gigabits per second—some 1,600 times faster than the average global broadband connexion. The operation, the largest ever inter-continental high speed data transfer, marked a major advance in state-of-the-art information science and technology"

MARSEILLE CONNECTION

ITER is located in a rural area with conventional business network availability (tens of Gbps).

Works have been done in 2022-2023 to arrange dedicated redundant fiber optics to Marseille area (80 km from ITER). As a result, ITER has got 2 x 400 Gbps connectivity readily available.

Marseille is an international network hub with research networks and cloud providers already present with hundreds of Gbps capacity.

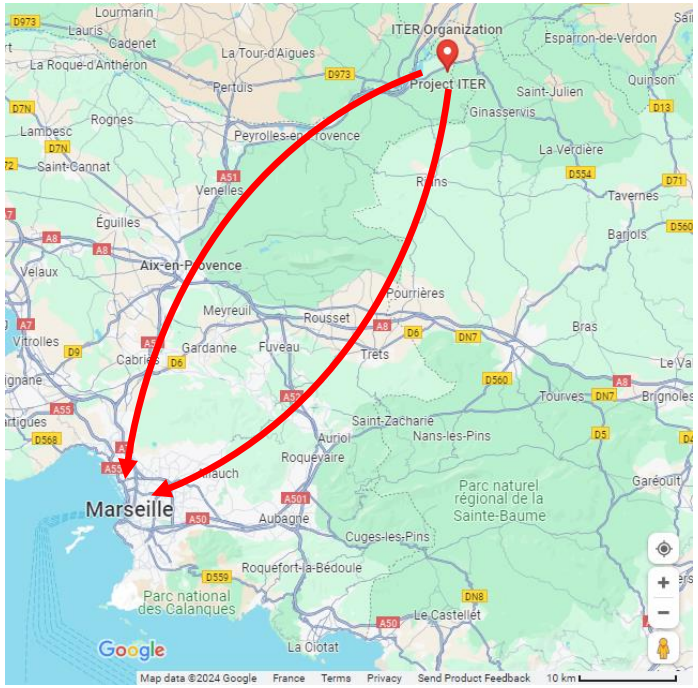

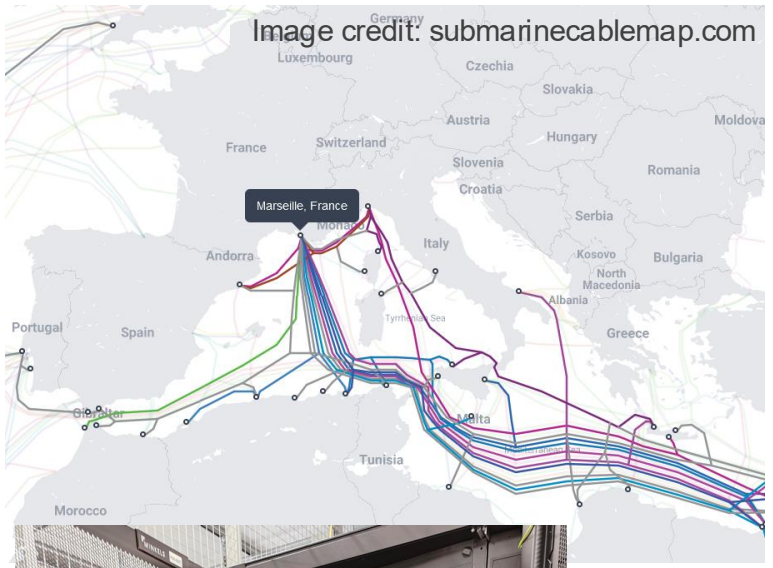

Digital Realty MRS3 data center in Marseille

SINET PROGRESS

- SINET – the Japanese NREN, operated by the National Institute of Informatics (NII)
- Since the 2016 tests, the international connectivity was upgraded to 100 Gbps
- “SINET6” since 2022 (largely 400 Gbps domestic)
- In 2024 the North America link Tokyo – Amsterdam upgraded to 400 Gbps
- In 2025 GEANT – SINET connection in Amsterdam upgraded to 400 Gbps

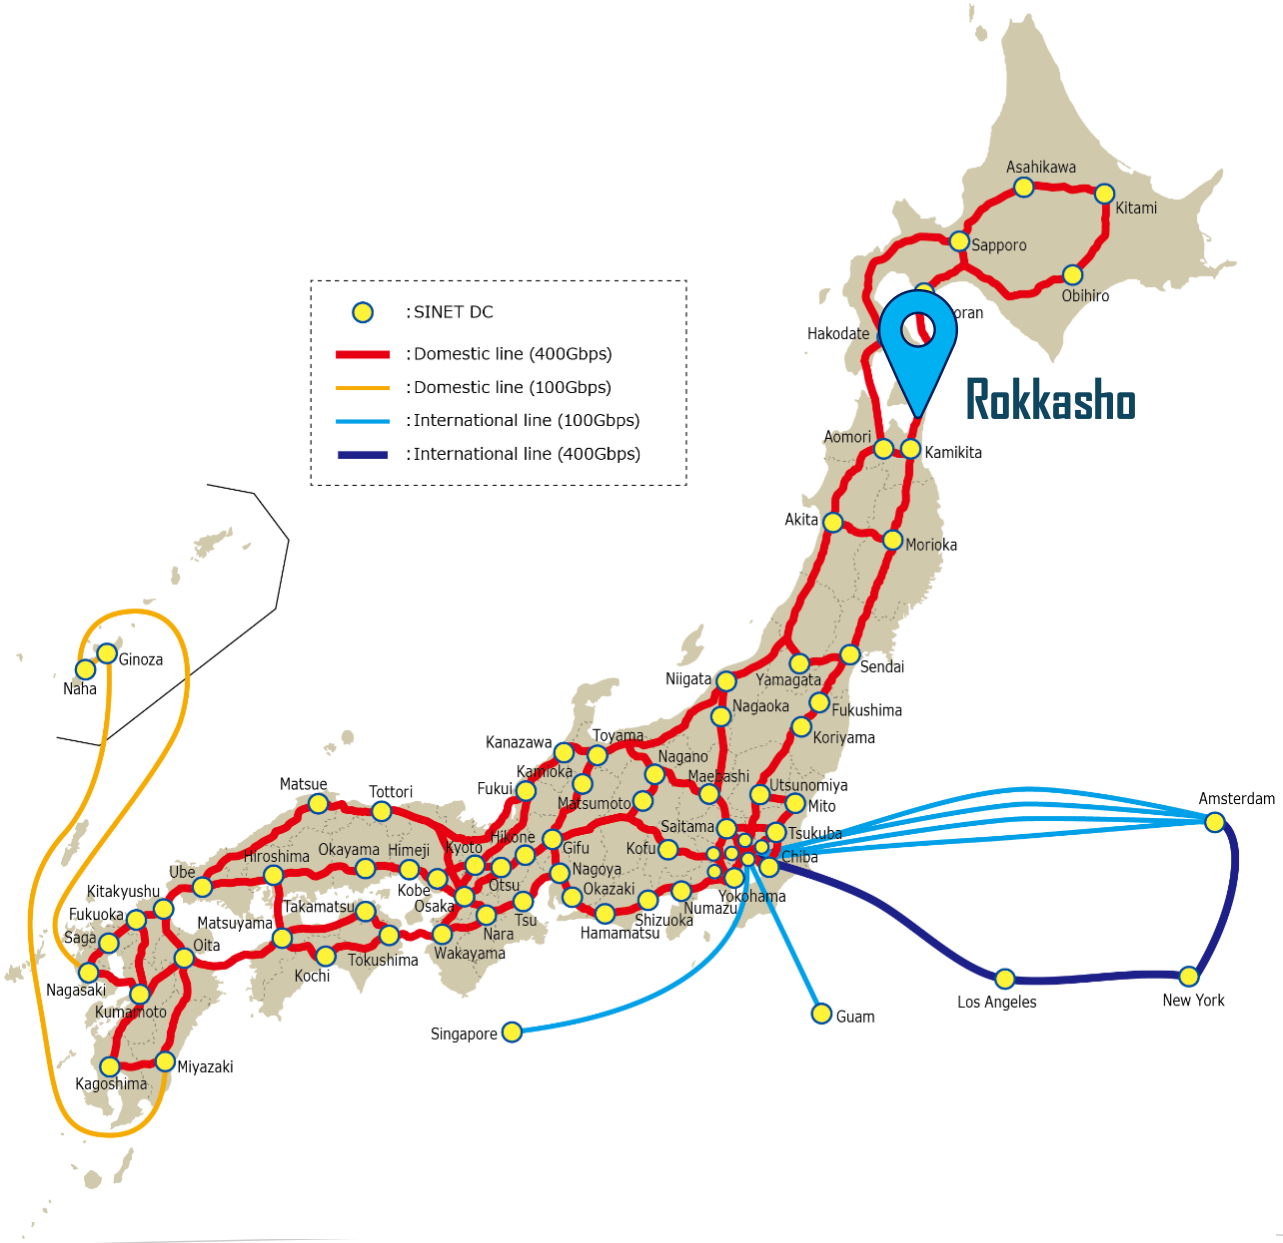

JAPAN'S REMOTE EXPERIMENT CENTER CONNECTIVITY TO ITER

ITER Cadarache ↔ RENATER ↔ GEANT ↔ SINET ↔ QST Rokkasho

- 2021: connected to ITER 24x7 using L2VPN via the shortest "Trans-Siberian" path
- 2024: switched to a longer "North America" path
- 2024: reconnected via ITER Marseille, therefore removing the last bottleneck of 10 Gbps in Cadarache and becoming 100 Gbps end-to-end
- 2025: GEANT configured for us their second 100 Gbps link IC-1 to Asia via Singapore (SingAREN)

FINAL CONFIGURATION FOR THE TEST

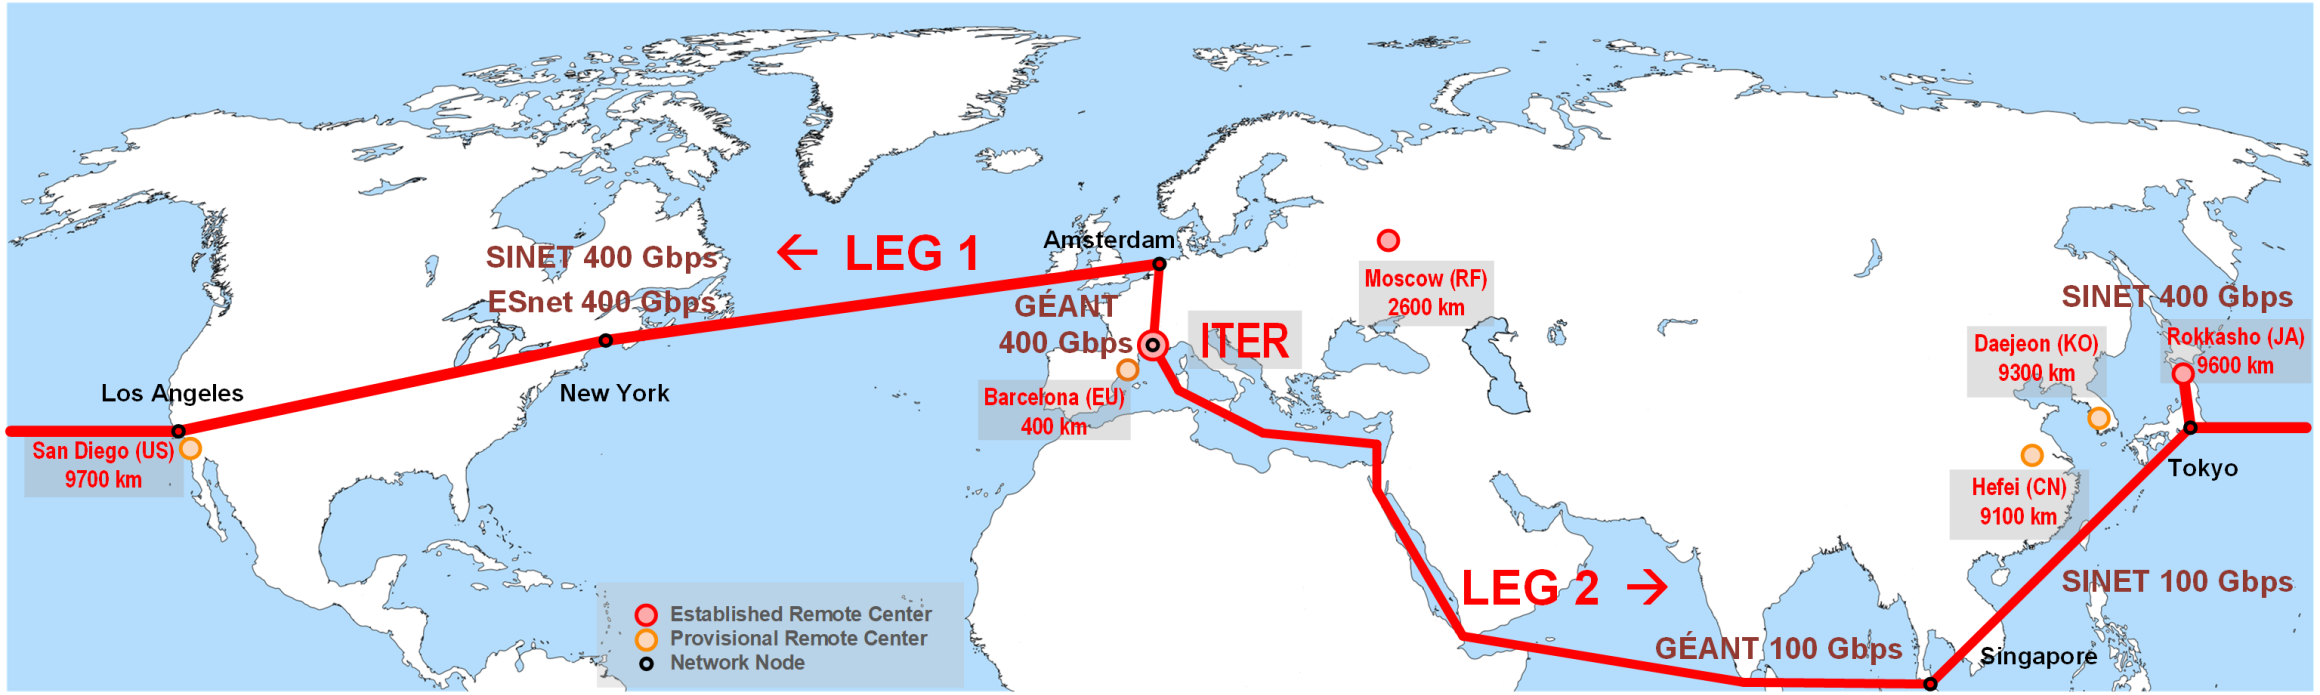

ECMP option:

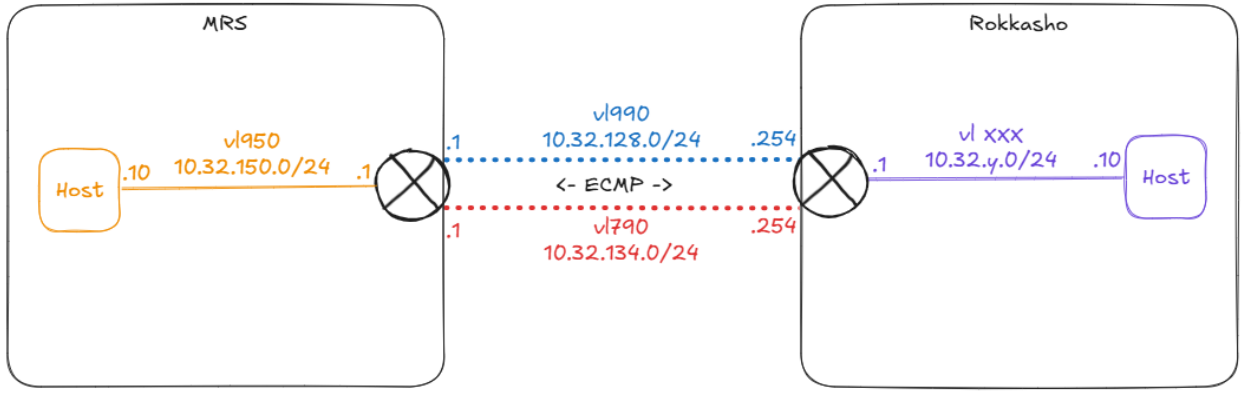

TEST HARDWARE SETUP

Marseille

Multiple servers:

- Storage IBM ESS 3200:
  - SSD 160 TB NVME
  - 8x100 Gbps InfiniBand
  - FS: Spectrum Scale / GPFS

- DTN HPE DL360 Gen10+ x 8:
  - CPU 2 X 40 core Intel Xeon Gold 6138 @ 2 GHz
  - RAM 256 GB
  - 100 Gbps Infiniband
  - 100 Gbps Ethernet
  - OS: RHEL 9.2

Rokkasho

Single DTN:

- CPU Intel(R) Xeon(R) w5-3425 (12core/3.2GHz) Ice Lake
- RAM 512 GB
- Samsung SSD 990 PRO 2x20 RAID 40 TB
- NIC Mellanox ConnectX-4 SP (MCX415A-CCAT)
- OS: AlmaLinux 9.5
- FS: ?

On both sides the kernel network parameters and system limits need to be tuned for long distance TCP transfers

TEST CONDITIONS

Differences with regard to the 2016 test:

- x20 network capacity (10 Gbps → 2 x 100 Gbps)
- ITER connection via the network node in Marseille (simulating a data distribution tier after the main data center permanent storage in Cadarache)
- Simultaneous transfer on two physical inter-continental channels. Simulation of sub-marine cable breakdown
- Two different tools testing (MMCFTP and ITER.sync)
- Use of ITER real plant data + fusion data converted to the ITER format
- ECMP load balancing option

Constraints:

- One direction only (IO → JA)
- No physical link could be loaded to its full capacity for a long time
- Disk space shortage on both ends
- Data sets prepared in advance; no live data from the plant participating
- No data pre- or post-processing (in a scientific sense)
- Client: basic plotting tools directly from HDF5; no data service involved on either end

ORGANIZATION / TIMELINE

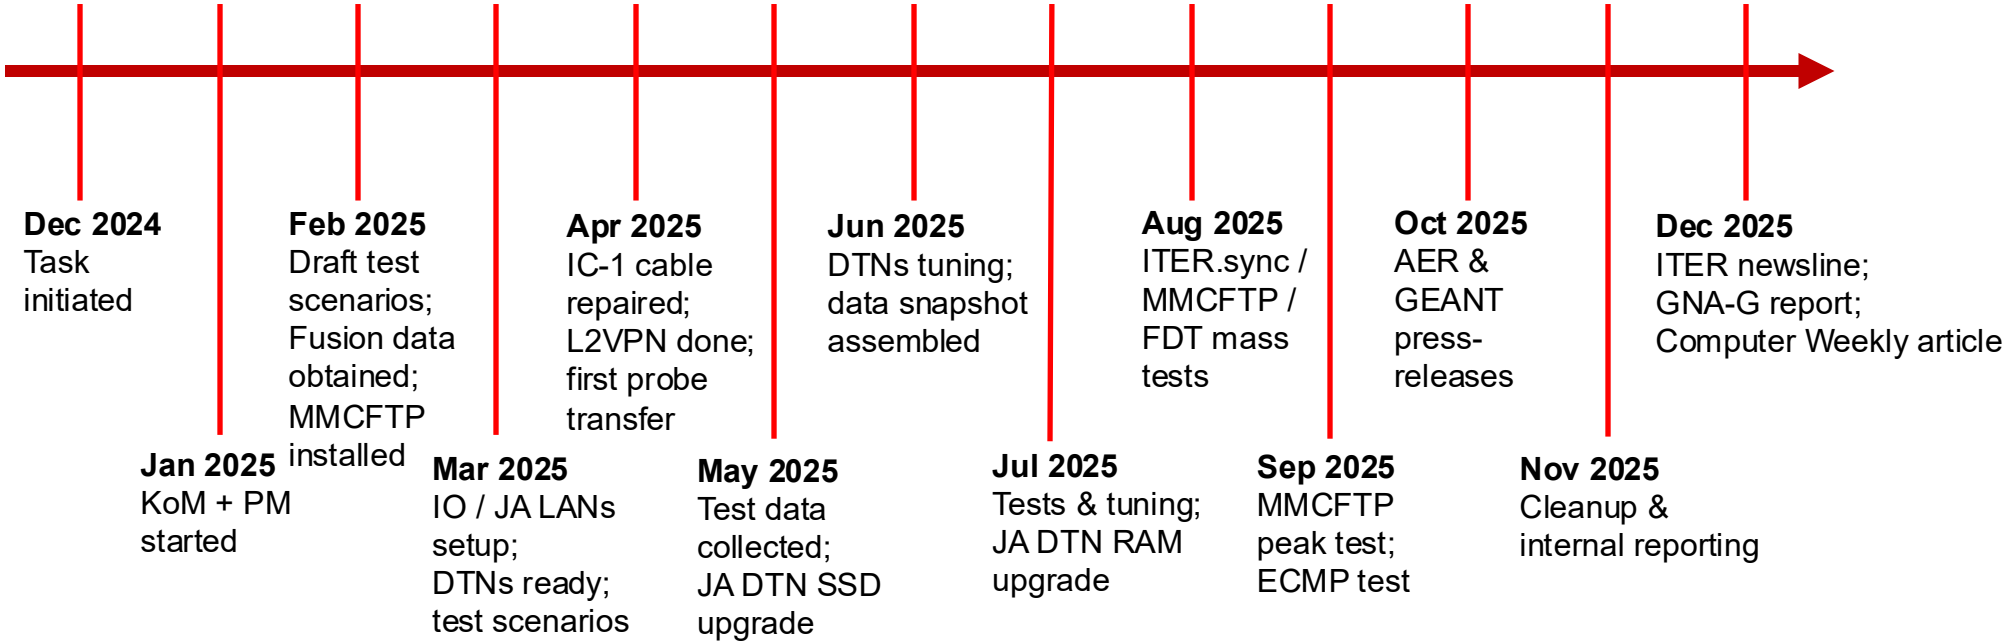

- Bi-weekly progress meetings IO CODAC / IO IT / QST / NIFS / GEANT
- 2nd L2VPN setup with RENATER / GEANT / SingAREN / SINET
- Permission from GEANT to routinely use 2 x 50 Gbps

TEST CAMPAIGN AUGUST – SEPTEMBER 2025

ITER.sync tests in first half of August followed by MMCFTP and FDT tests in second half of August and September. The 07/09 peak pictured is the 100 Gbps network saturation test.

ITER.SYNC

A management layer above ‘rsync’, splitting payload in large chunks and transferring them using many ‘rsync’ processes on many servers.

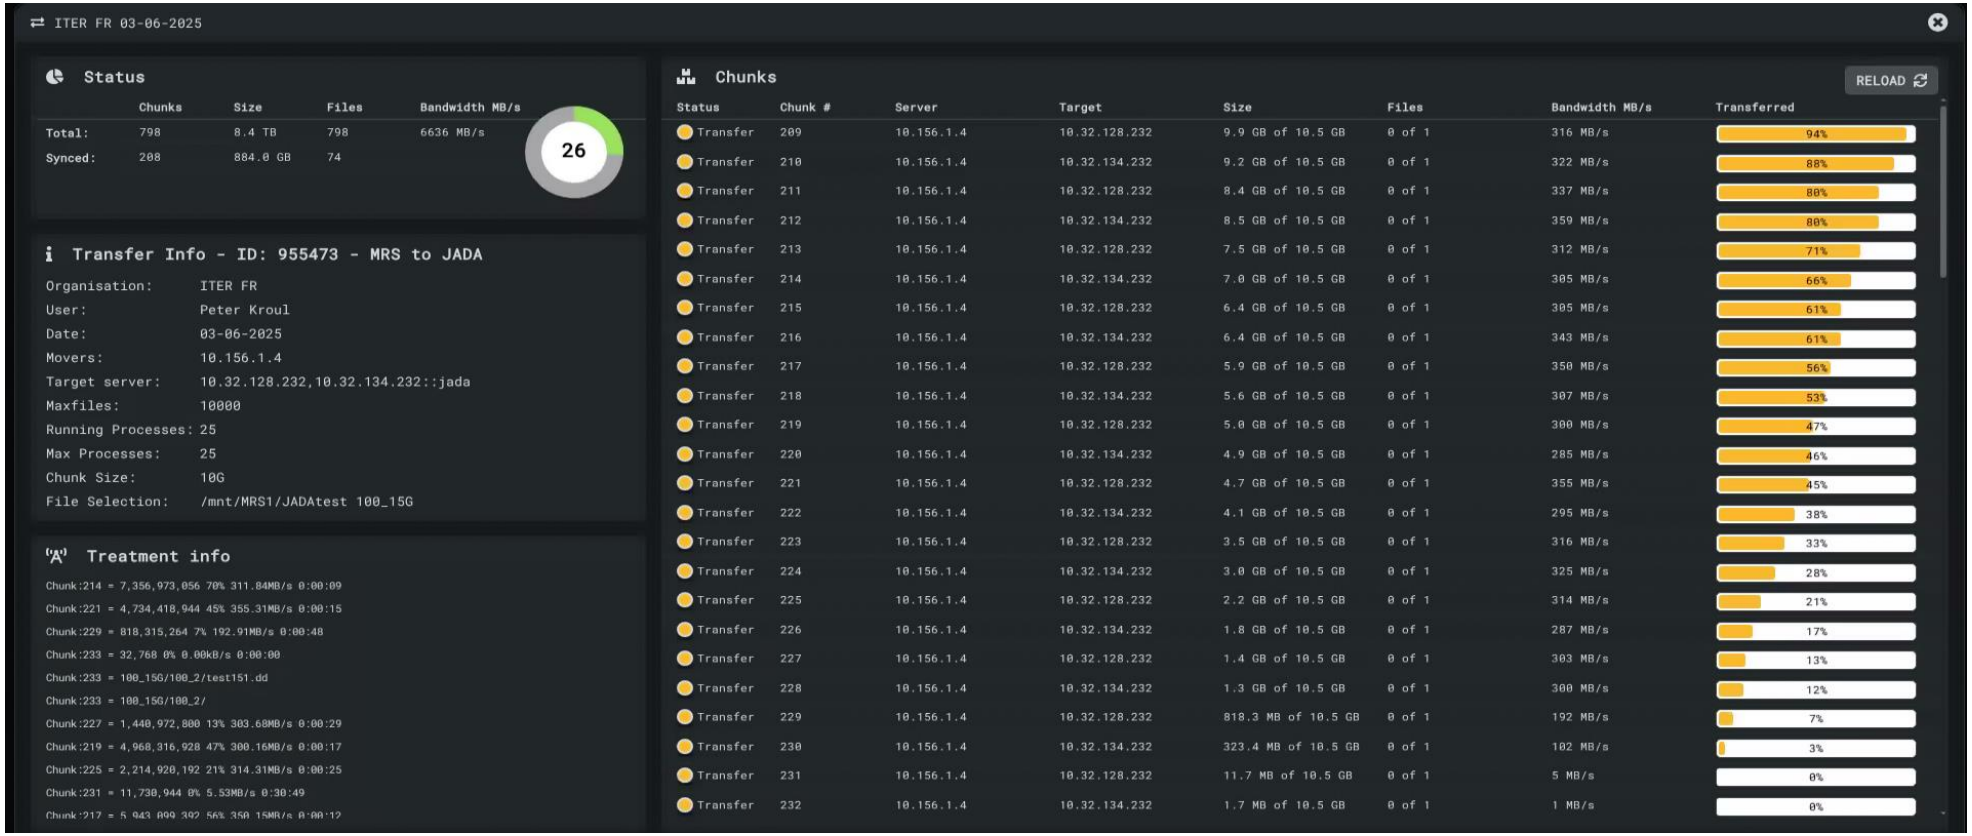

ITER.sync transferring a batch of 8.4 TB in 798 files in chunks of 10 GB from ITER to Japan using 25 parallel rsync processes.

24 HOURS TRANSFER

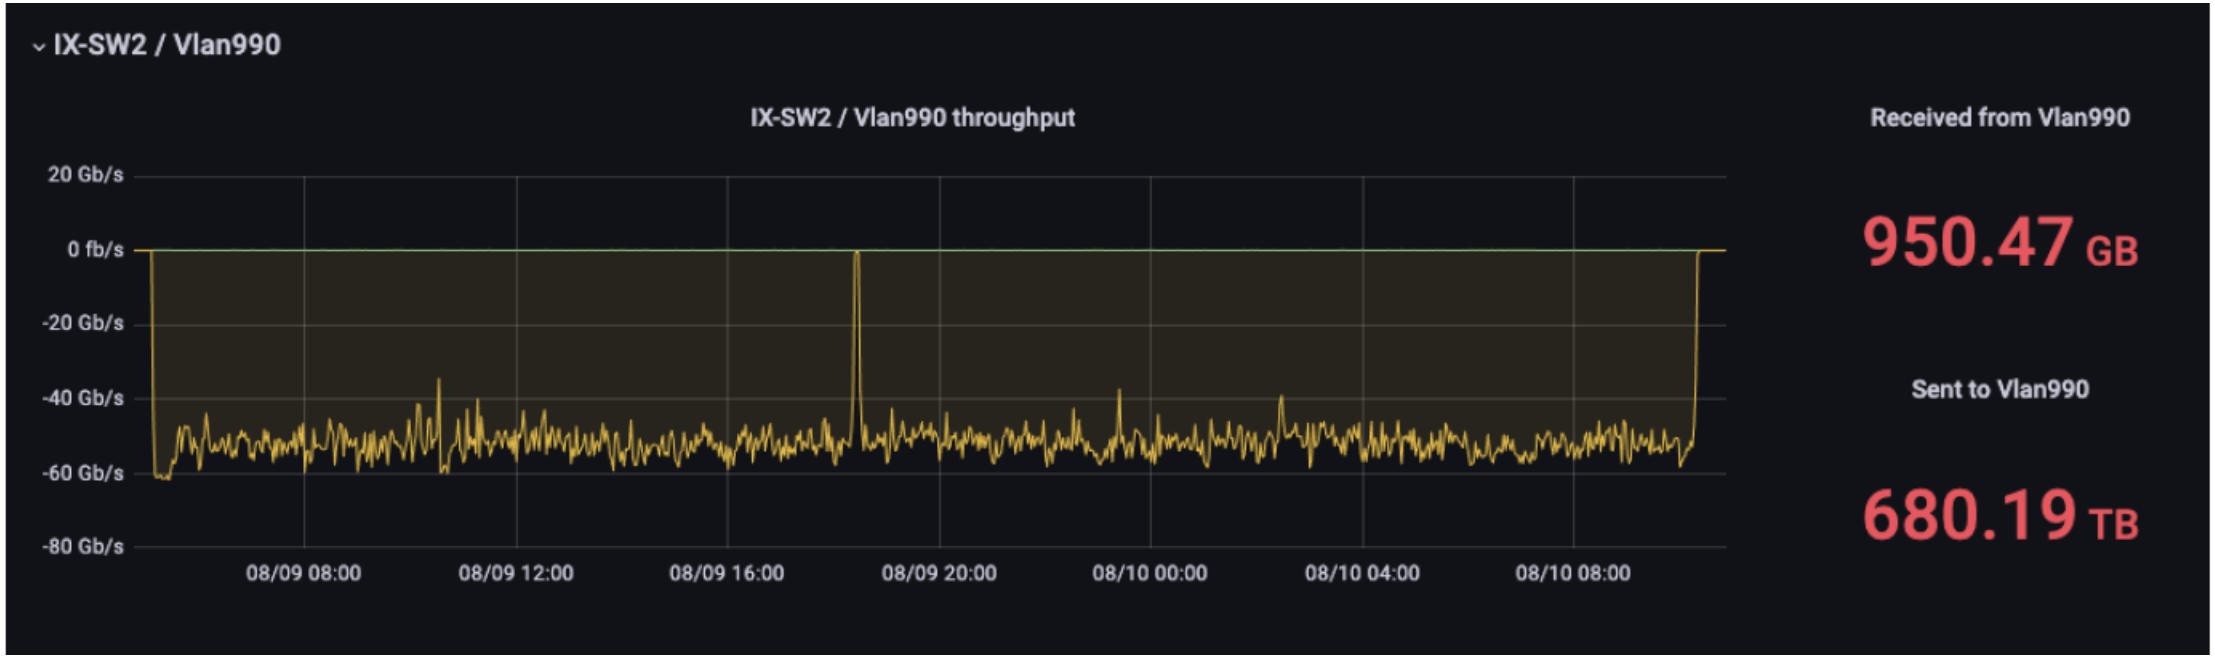

ITER.sync transferring 3 shifts of 176 TB worth of data within 24 hours (pictured duration is ~30h). The speed cap is coming from the limitations in server setup.

MMCFTP MAX NETWORK UTILIZATION TEST

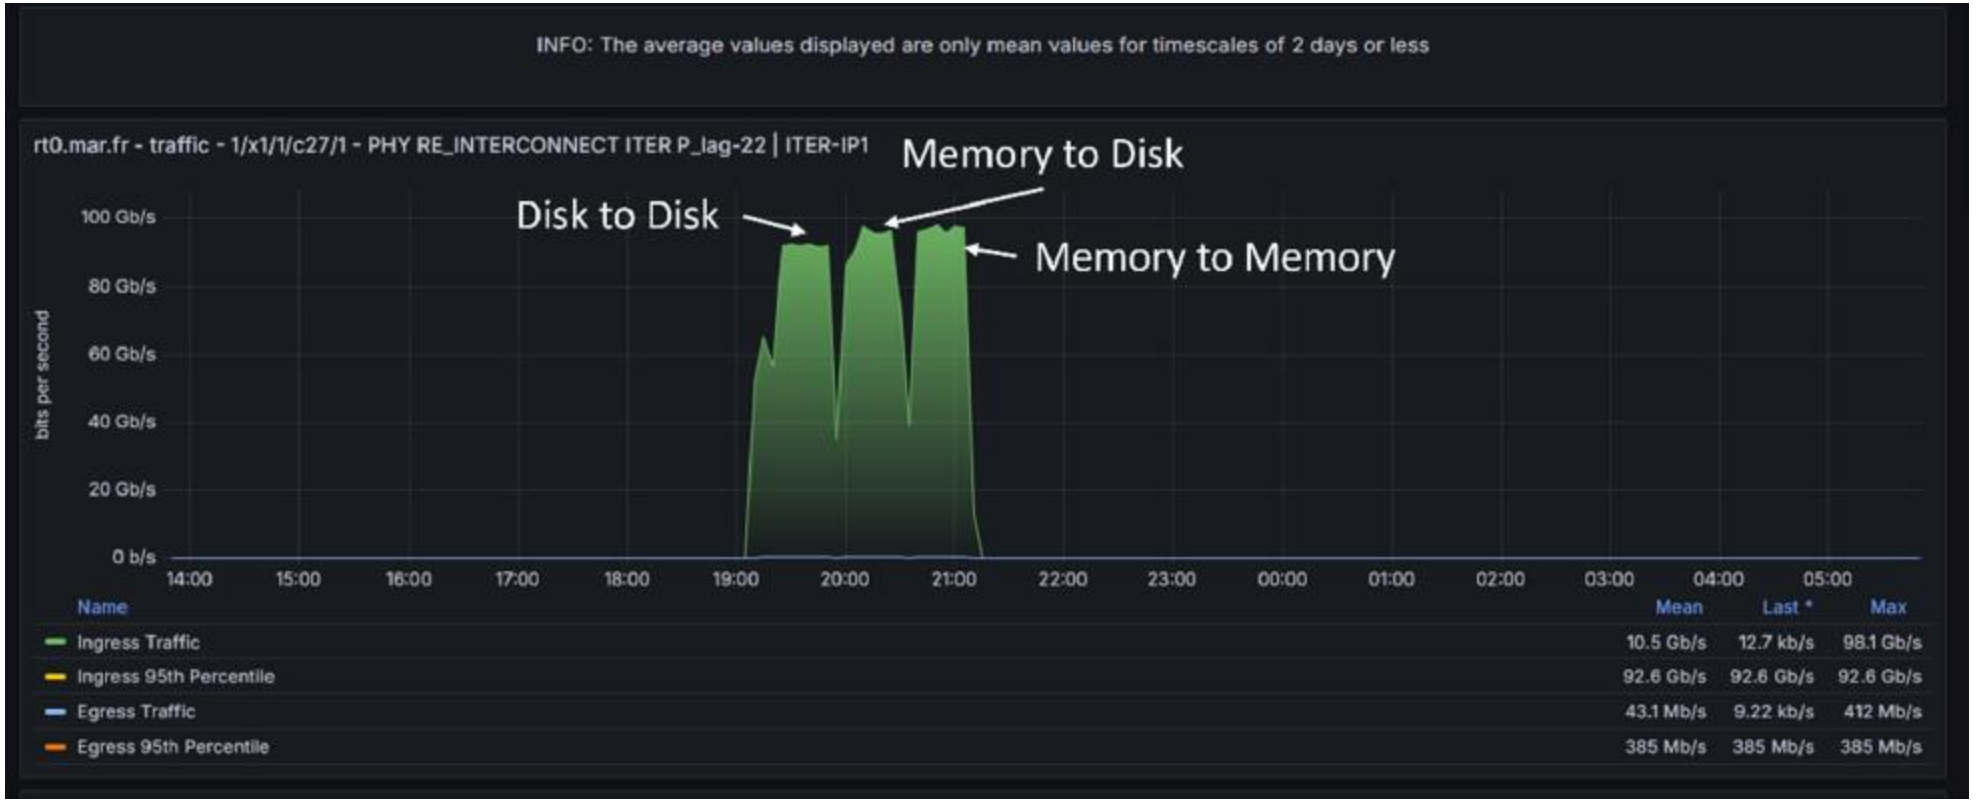

MMCFTP max speed test results using different data transfer modes: disk to disk (achieving 91 Gbps), memory to disk (94 Gbps) and memory to memory (96 Gbps).

MMCFTP USING TWO NETWORK PATHS SIMULTANEOUSLY

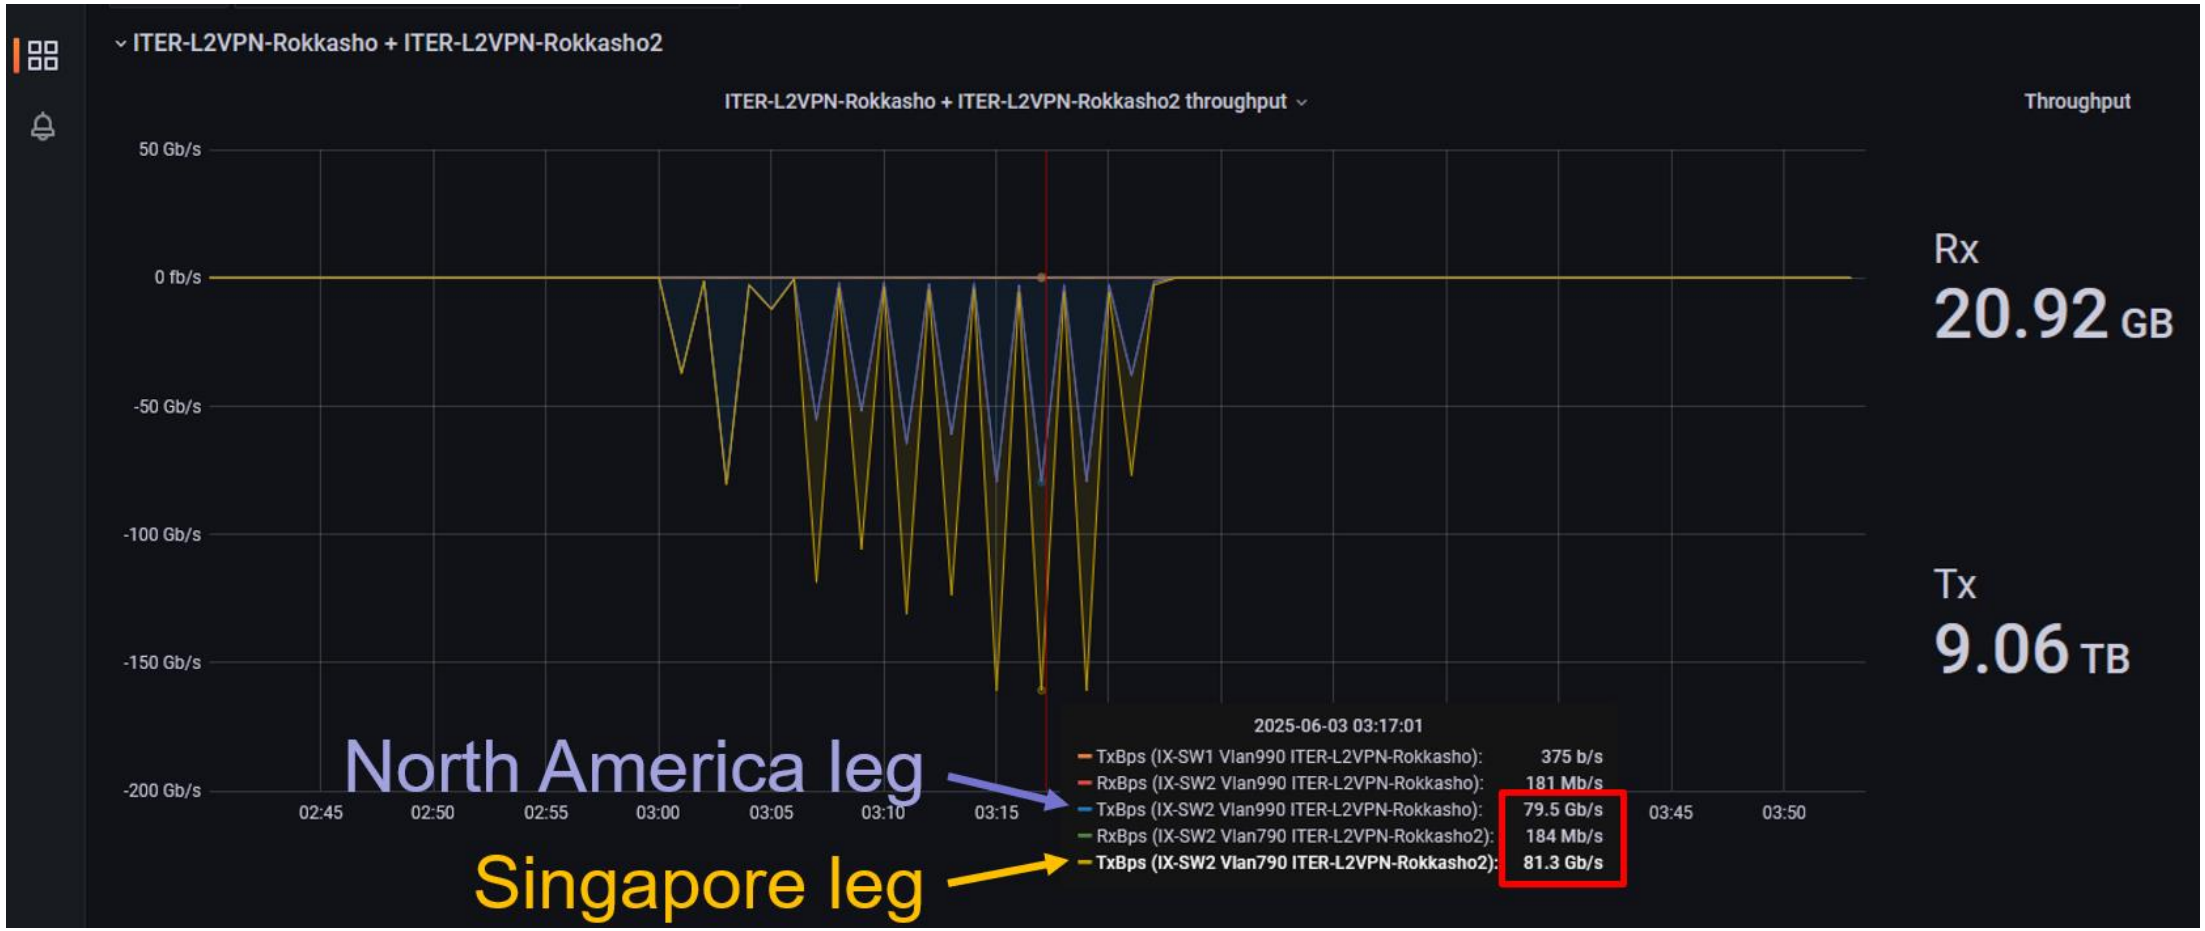

MMCFTP managing both inter-continental 100 Gbps physical paths while transferring 9 TB of data. 200 Gbps are not attained in this example because of bandwidth ceiling set in the tool.

FINDINGS AND TAKEAWAYS

- Network performance was good and reliable
- Good working relationships with NRENs
- Server setup is crucial and it should be chosen in accordance with data transfer scenarios and tools used. File system choice is also important
- Saturation of a 100 Gbps link was demonstrated (MMCFTP). ITER.sync could reach link saturation with two or more DTNs on the receiving end – scales well with the number of servers, in expense of higher server resources needed
- Network transfer is only a part of the workflow. There are extra steps to be done between producing data on local end and accessing data on remote end, which might easily take as much time as the transfer itself
- Scenarios with a remote experimentalist in the loop will certainly require a reserve network channel (hot standby or dynamically managed)

CONCLUSIONS

- Intercontinental 100 Gbps mass transfer successfully demonstrated on data sets reaching 500 TB / day
- Parallel use of two 100 Gbps channels for the same transfer demonstrated
- The results will be used to improve software and steer the final server setup
- For the currently projected data rates and project schedule ITER is quite confident that networks would not represent a bottleneck

CREDITS

Special thanks:

- Peter Kroul (ITER)
- Thierry Reboul (ITER)
- Kenjiro Yamanaka (NII)
- Shinsuke Tokunaga (QST)
- Hideya Nakanishi (NIFS)
- Mario Reale (GÉANT)

In the news:

- AER: https://aer-network.net/?p=269
- GÉANT: https://connect.geant.org/2025/10/27/fast-redundant-resilient-iter-japan-proves-the-power-of-aer-at-100-gbps
- ITER: https://www.iter.org/node/20687/iter-demonstrates-fast-data-transfer-japan-and-us

Thank you!

iter

china eu india japan korea russia usa