



# GEANT2: The European Network for R&E

Roberto Sabatino, DANTE  
UCC,  
20 March 2007



Connect. Communicate. Collaborate

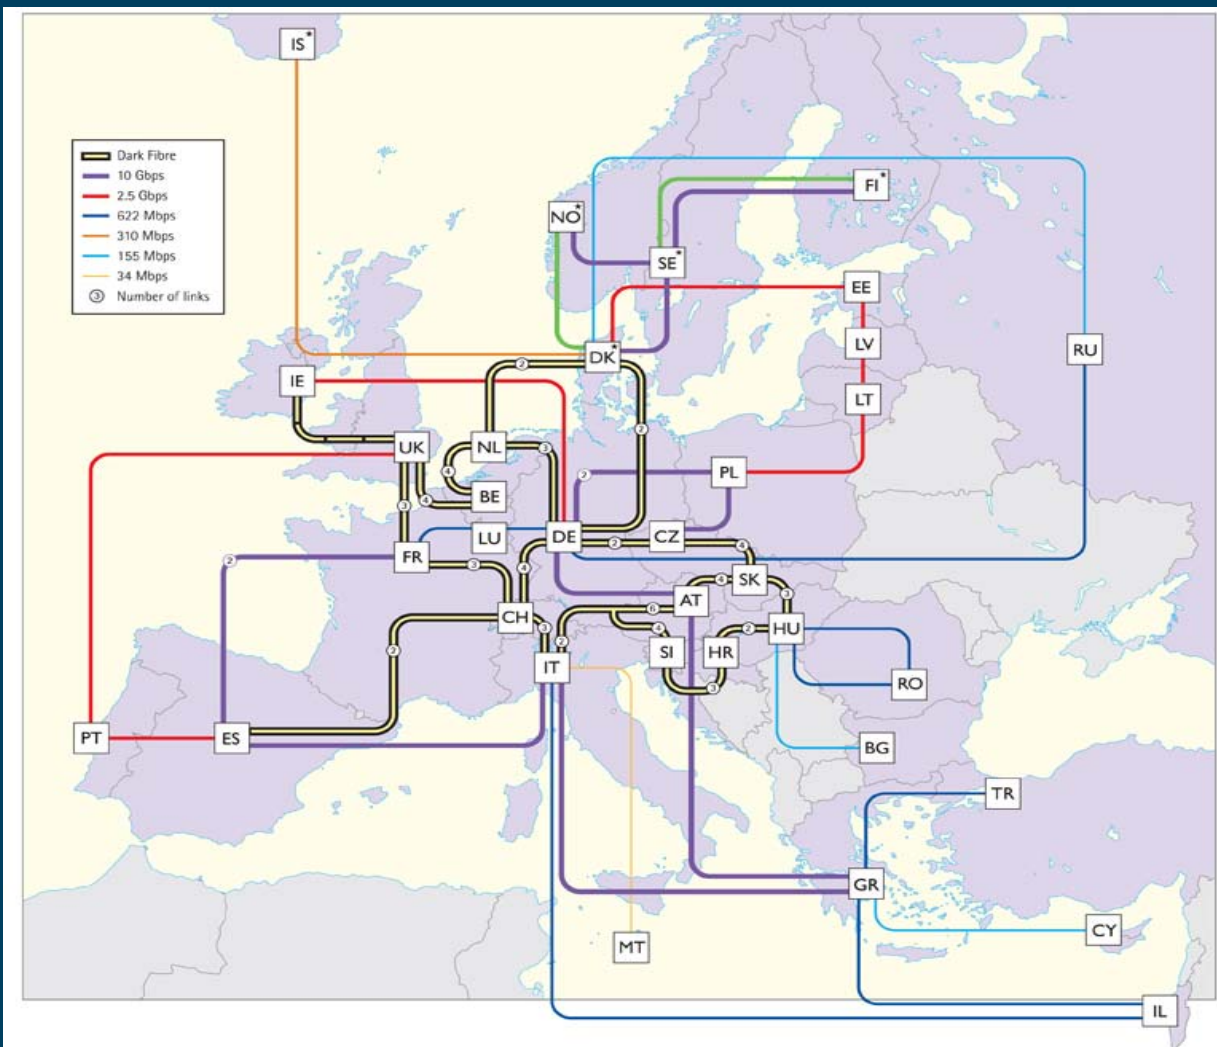
# Introduction

- The GÉANT2 network interconnects the National Research and Education Networks in Europe
  - HEANET in Ireland
- It is implemented, managed and developed by DANTE in Cambridge
  - 35 staff, 19 engineers
- Governed by a Policy Committee
- Co-funded by EC (200M Eur budget over 4 years)





Connect. Communicate. Collaborate





Connect. Communicate. Collaborate

# Evolution

- 1996-1998 TEN-34: a 34Mbps backbone based on IP over ATM technology
- 1998-2001 TEN-155: a 155Mbps (later 622Mbps) backbone, offering an IP and an ATM service
- 2001-2005: GÉANT. A major step forward. 10Gbps IP network
- 2005- GÉANT2. Another massive step forward: Nx10Gbps backbone exploiting fibre and optical technology. IP and Ethernet Virtual Line services





Connect. Communicate. Collaborate

# Agenda

- High level overview of the network
- IP Routing
- Security + Netflow
- Monitoring tools development
- Demo of some of the tools
- A brief introduction to advanced services: multicast, ipv6, MPLS, E2E GE
- Q&A

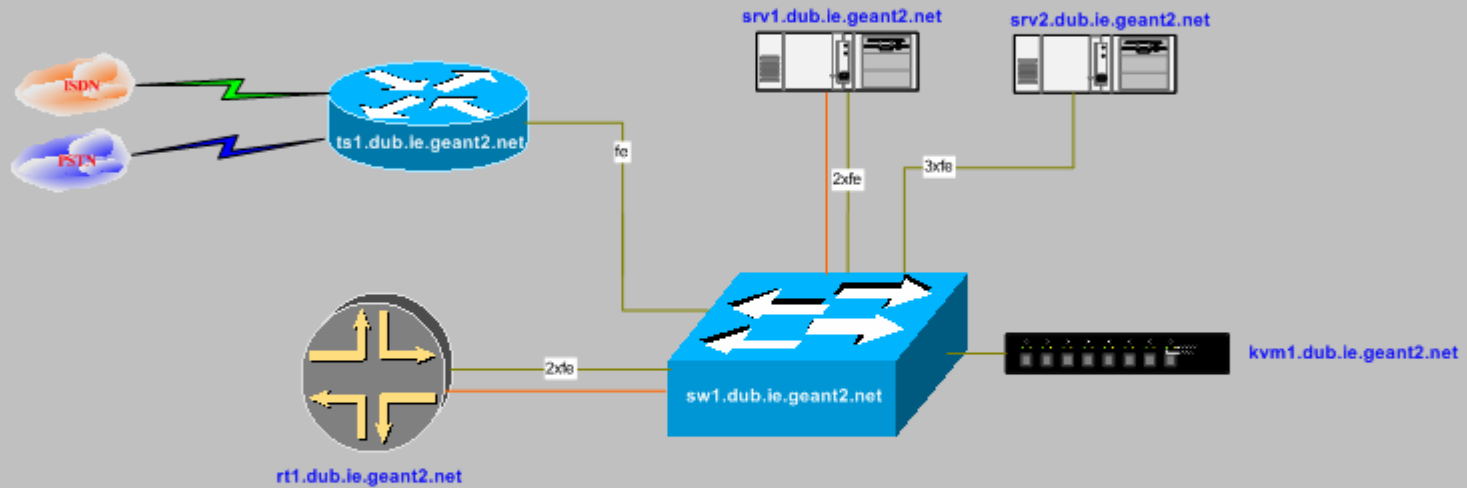


Connect. Communicate. Collaborate

# POP setup

- 1 big router per PoP
  - Juniper M40 (8 slots at 2.5Gbps each)
  - Juniper M160 (8 slots at 10Gbps each)
  - Juniper T640 (8 slots at 40Gbps each)
- Connections to NREN (typically at 10Gbps)
- Backbone connections (2 or more at 10Gbps)
  
- Local LAN
  - Terminal server, Ethernet switch (Cisco)
  - workstations, servers (SUN)

# GEANT2 IE PoP LAYOUT-ETHERNET CIRCUITS



LEGEND	
	Gigabit Ethernet
	Fast Ethernet
	Multimode
	Fibre





Juniper  
NETWORKS

M40

REAR PANEL CONTROLS

REAR	REAR	REAR	REAR	REAR	REAR	REAR	REAR	REAR	REAR
------	------	------	------	------	------	------	------	------	------

POWER

STATUS

REAR PANEL CONTROLS

M40



Connect. Communicate. Collaborate

# Juniper routers

- Operate at line rate
  - Each line card has a forwarding table, updated by central routing engine
- Based on ASICs (Application Specific Integrated Circuits)
  - Pro: very fast forwarding and treatment of packets (forwarding, ACLs)
  - Con: slow development: typically 6 months to 1 year
- Underlying OS (JUNOS) based on Unix, allows all-in-one SW release
- Multi-chassis options for T-640



Connect. Communicate. Collaborate

# IP setup

- Internal protocol IS-IS
  - Was OSPF
  - Moved to IS-IS in 2002 to enable ipv6
- External BGP-4
  - Communities
  - MEDs
- AS number 20965
- Address space 62.40.96.0/20



# IP addressing – IE POP

Connect. Communicate. Collaborate

Router	Interface	Interface Address	Network Address ( /30 )	Usage	Destination address
ie1.ie	pos-7/0/0	62.40.96.166	62.40.96.164	Circuit to DE1	62.40.96.165
	pos-6/0/0	62.40.96.137	62.40.96.136	Circuit to UK	62.40.96.138

*Access and peering circuits (NRENs, INFONET, ... ) :*

Router	Interface	Interface Address	Network Address ( /30 )	Usage	Destination address
ie1.ie	pos-4/0/0	62.40.103.229	62.40.103.228	Circuit to HEANET	62.40.103.230
ie1.ie	pos-5/0/0	62.40.103.241	62.40.103.240	Backup circuit to HEAnet	62.40.103.242



Connect. Communicate. Collaborate

# Network monitoring

- Monitoring is the enabler for
  - Planning
  - Security
  - Performance analysis
- Passive : SNMP, netflow
- Active: inject measurement traffic into the network (one way delay, end to end throughput)



Connect. Communicate. Collaborate

# What is netflow ?

- Netflow is a standard way of exporting information on traffic flows (*srcadr, dstadr, srcport, dstport*)
  - Byte, packet count
  - Protocol
  - Next hop
  - Src, dst AS (peer or origin)
- Information on a flow is exported every 30 mins
- You need to collect and analyse the data



Connect. Communicate. Collaborate

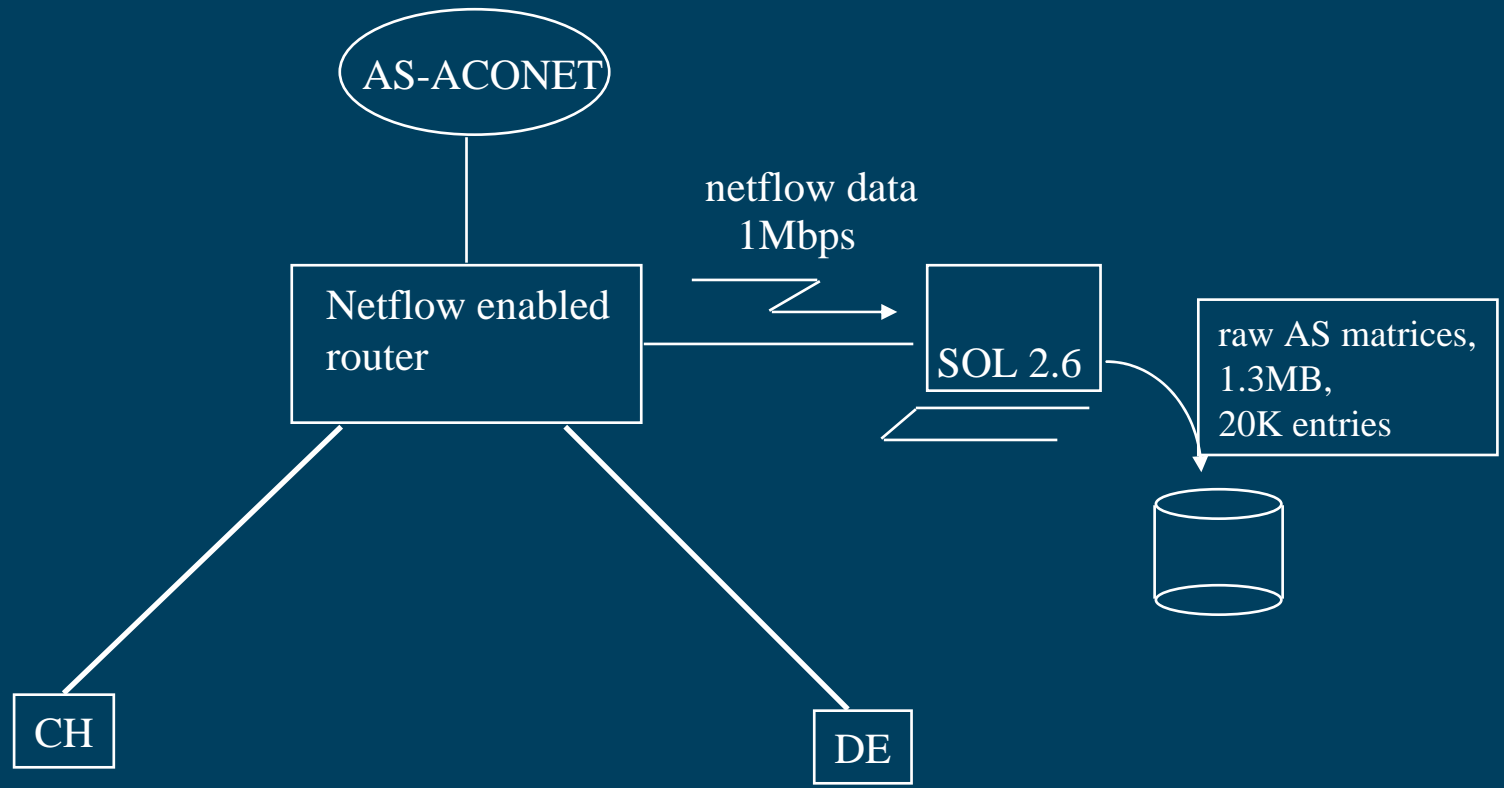
# Why is netflow useful ?

- It helps get an insight into the traffic patterns on your network
  - Planning & dimensioning
- It helps you detect anomalies
  - Security

# NETFLOW DATA COLLECTION



Connect. Communicate. Collaborate

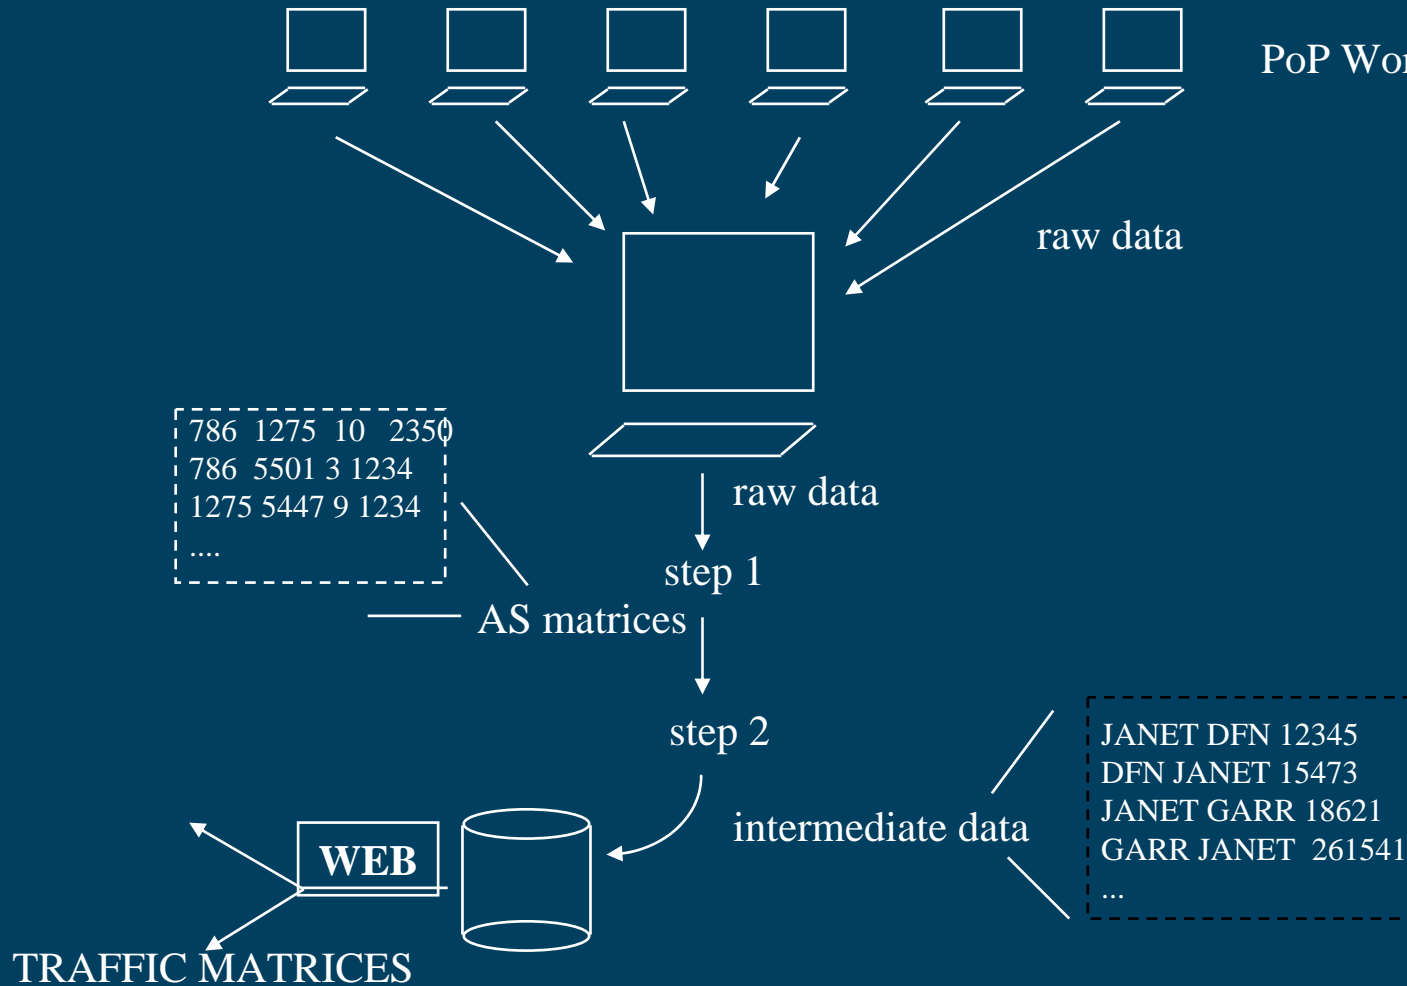


# NETFLOW DATA ANALYSIS



Connect. Communicate. Collaborate

PoP Workstations





Connect. Communicate. Collaborate

# Using netflow for security

- A typical use case is DoS attack detection and routing anomalies
- DoS attacks:
  - checks flows with similar target destination IP
  - raises alarms when a (configurable) limit is reached
  - logs the attack
- Routing anomalies: checks flows with sources AS = 0 (and destination AS = 0)



Connect. Communicate. Collaborate

# Some stats


- Alarms raised / real DoS attacks: 97%
- Between 10 and 40 DoS attacks detected per day (average: 25 attacks/day, up to 5-6 concurrent).
- Only 7% of the attacks last more than 30 minutes
- 95% of the attacks: “cclass” attacks (spoofed source address within the correct cclass of the source)
- 75% of attacks: tcp, 40 bytes packets



# A typical case (i)

Connect. Communicate. Collaborate

Duration of attacks






# A typical case (ii)

Connect. Communicate. Collaborate


Amount of packets





# A typical case (iii)

Amount of traffic



- <0.3 Mbps
- 0.3 - 1 Mbps
- >1 Mbps

# How to deal with a DoS attack ?



Connect. Communicate. Collaborate

- Classify the attack and consider the severity.
- If deemed to be causing a significant enough effect on services then,
  - Filter all DoS packets at the victim NREN access point.
  - Track the DoS attack back through the Network to the Ingress points and filter there subsequently removing initial filters on NREN access router
- Contact those parties involved and proceed with CERT investigation where possible.



Connect. Communicate. Collaborate

# TOOLS DEMO





Connect. Communicate. Collaborate

# Advanced networking

- Multicast
- Differentiated services (Premium, Scavenger)
- Ipv6, ipv6 multicast
- MPLS
- Gigabit ethernet transport over optical networks



Connect. Communicate. Collaborate

# TDM

- Time Division Multiplexing
- Traditionally telco territory to multiplex voice and data traffic: SDH(EU), SONET(US) standards
- Units of data: VC-3, VC-4
- Exp. An STM-1 (155Mbps) corresponds to 1 VC-4 (Virtual container-4). An STM-64 (10Gbps) is made up of 64 VC-4s
- Each VC-4 is allocated a very specific time slot on the transmission link: timing and synchronisation!
- Next Generation TDM switches: mapping and transport of Gigabit Ethernet



Connect. Communicate. Collaborate

# WDM

- Wave Division Multiplexing
- Fibre optics cables capable of transporting 96 wavelengths (colours), of 10Gbps capacity each
- DWDM (long haul), CDWM (metro area)
- Amplification : every ~70km
- Regeneration: 500-1000km
- Factors: quality of fibre, engineering of fibre routes: need to understand physics to cope with this !



Connect. Communicate. Collaborate

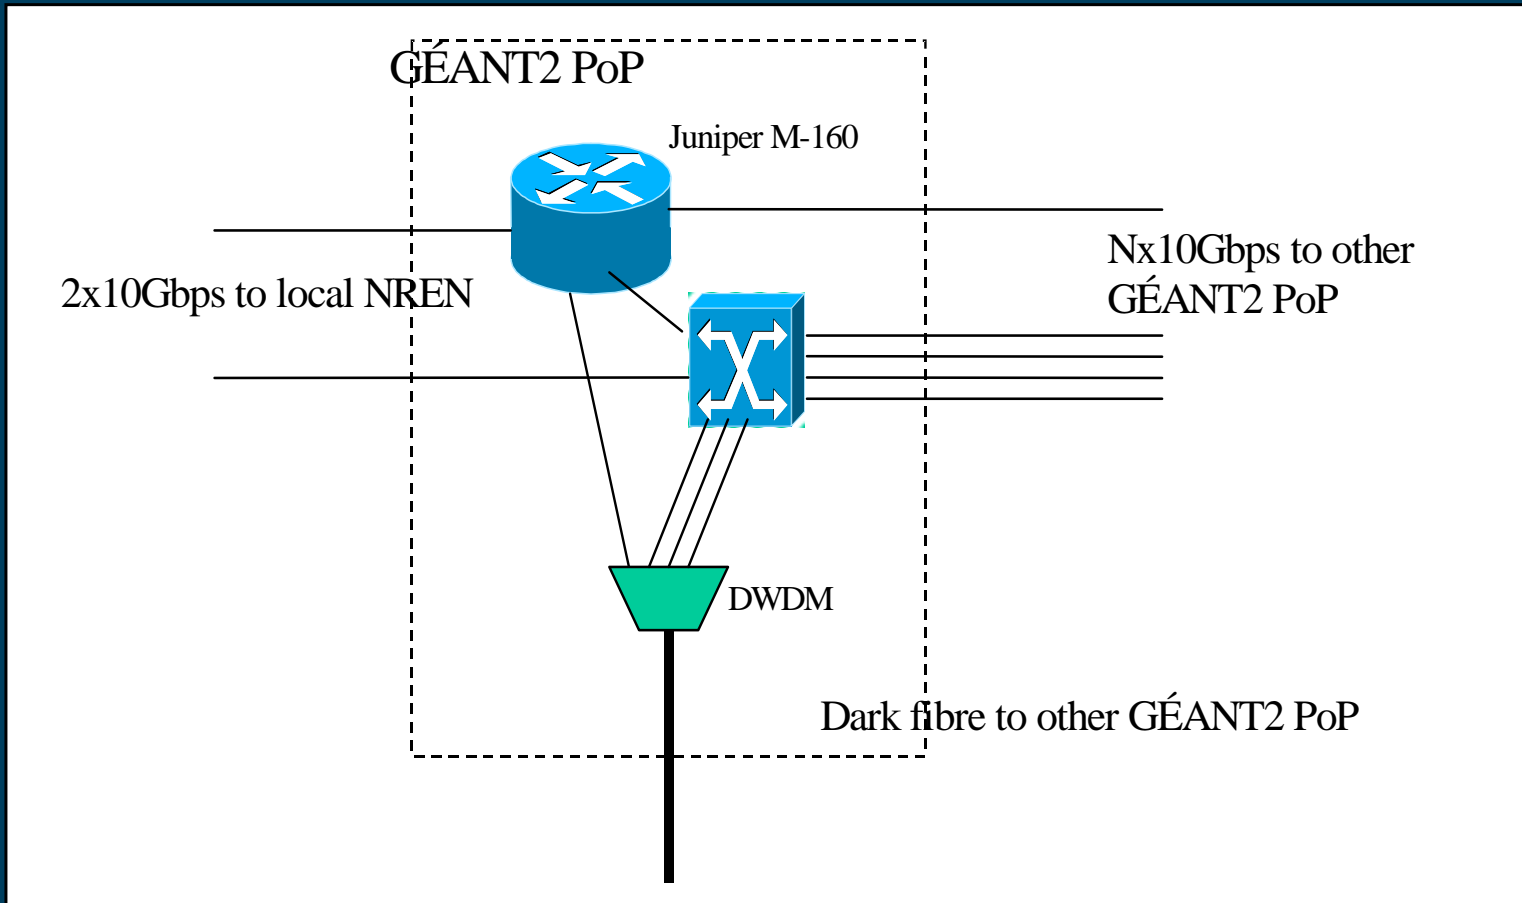
# TDM and WDM on GEANT2

- Combination of both technologies
- WDM to offer multiple wavelenths, at 10Gbps capacity, on each fibre route. Engineered to support 40. Planned 14 by 2008
- TDM to offer GE services. 1 wavelenths can carry 9 GEs
  - Each GE occupies 7 VC-4s. Total 63 VC4-s out of 64

# GÉANT2 PoP Design



Connect. Communicate. Collaborate



# End 2 End Gigabit services



Connect. Communicate. Collaborate

