amsix
amsterdam internet exchange

Optical Switching at AMS-IX

A great tool in platform migration

Romeo Zwart

Amsterdam Internet Exchange

NANOG 32 - October 2004

Scope

Remember "IX's lessons learned" ?

(Steve Feldman, NANOG #30, Feb 2004)

- Myth #1: "10Megabit is enough"

It worked great, until it got full

This talk today:

– How we are dealing with the *"it gets full"* part

– Our upgrade path so far

– The value of Optical Switching

AMS-IX Traffic volume

- Deal with exponential volume growth

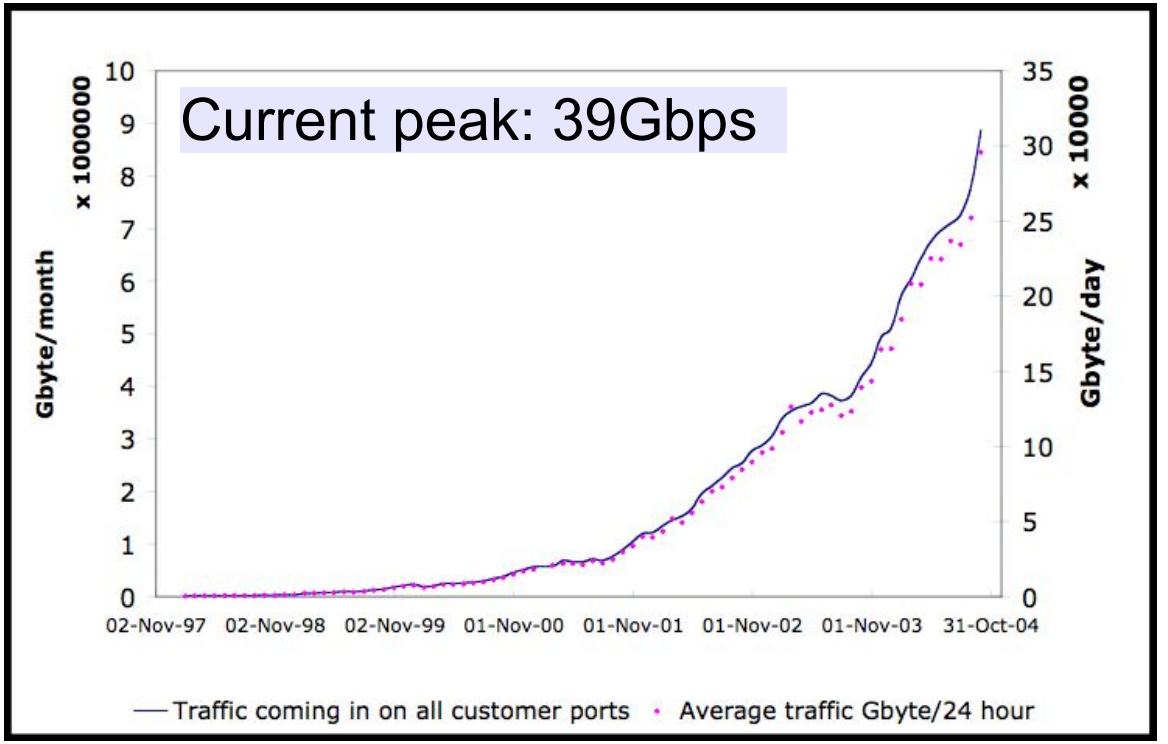

Architectural changes

- Ring vs Spoke-Hub Architecture
  - Inter-node bandwidth issues with a Ring
    - Need to scale up the entire ring to cover backup load in case of link failures
  - Concentrate high-volume customers at one location with a Spoke-Hub topology
    - Bringing down inter-node traffic where possible
  - Expansion can be handled more easily in a Spoke-Hub environment

So ....

how do we get this all done

without a 4 week outage?

Enter... the Photonic Switch

The Photonic Switch

What is it....

Think of it as an ultra-fast, automated, patch panel

- Glimmerglass System 300 Fiber Connection Server
- Built up around Micro Electro Mechanical Systems (MEMS)
  - – Miniaturized, moving components
  - – Built with IC manufacturing technology
    - » Cost effective

The Photonic Switch

How does it work...

- 2D-array of tiny mirrors
- Each mirror reflects light from one fiber into one other fiber.
- Each mirror is individually adjustable

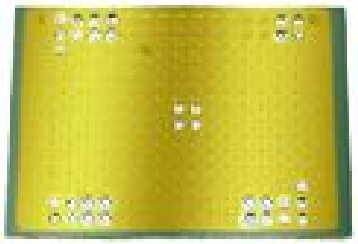

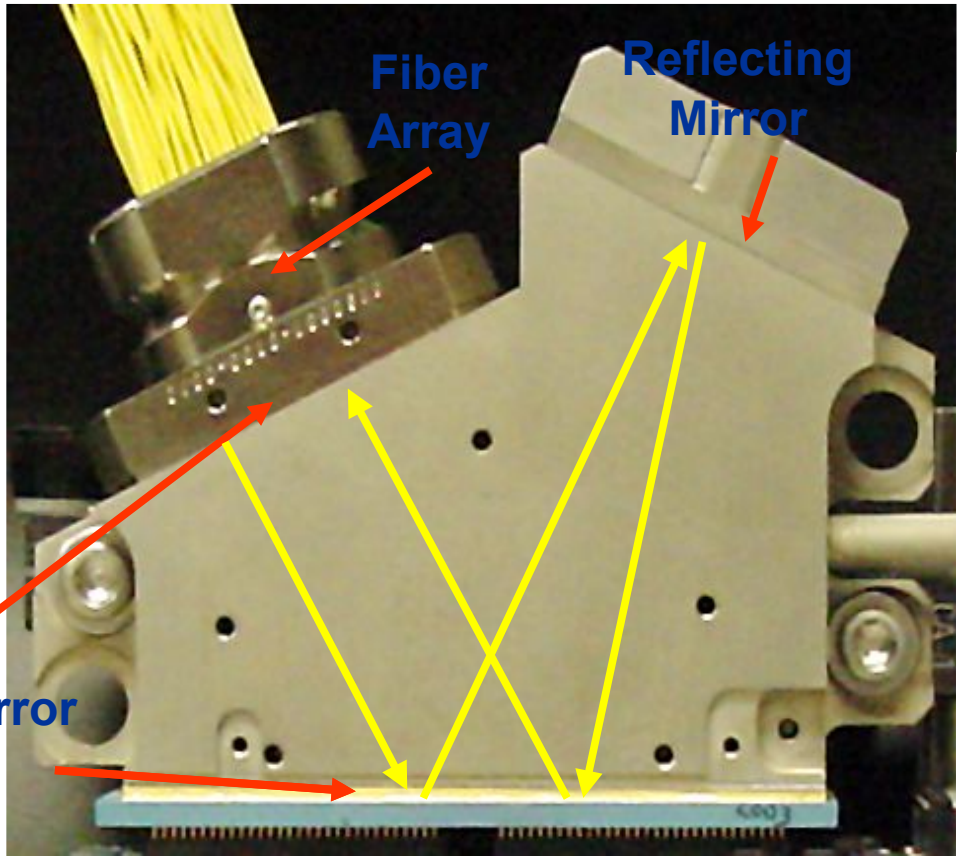

The Photonic Switch

Now for the bad news...

- Induces some signal loss (~2dB)
- There have been some reliability issues with the first revision
  - Two of the photonic switches have been replaced bcs of hardware issues

But

- No issues with the MEMS
- Switches have been fitted with new electronics under warranty
- **No service interruptions** caused by these issues

Optical Switching at AMS-IX

- How do we use the Photonic Switches?
  - How are they controlled?
  - Some use cases and examples

Optical Switching at AMS-IX

Keeping Control

- Photonic switches controlled by TL-1 interface
  - Web interface (port dis/connection)
  - Shell access - VXworks
    - Manual TL-1 commands
- Home-grown remote control tools
  - Remote TL-1 command tool
  - PSCD (Photonic Switch Control Daemon)
    - VSRP state change triggers core switch-over

Optical Switching at AMS-IX

Some use cases and examples

Photonic switches provide flexibility in

- Migration support
  - Ring to Star transition
  - Adding/swapping core node(s)

- Supporting 10G ports resilience
  - Switching the active core node (supporting 10G customer resilience)

Optical Switching at AMS-IX Use Case I – Transition to Core

- Introducing the first photonic switch

NOTE:

Not the CORE switch (yet)

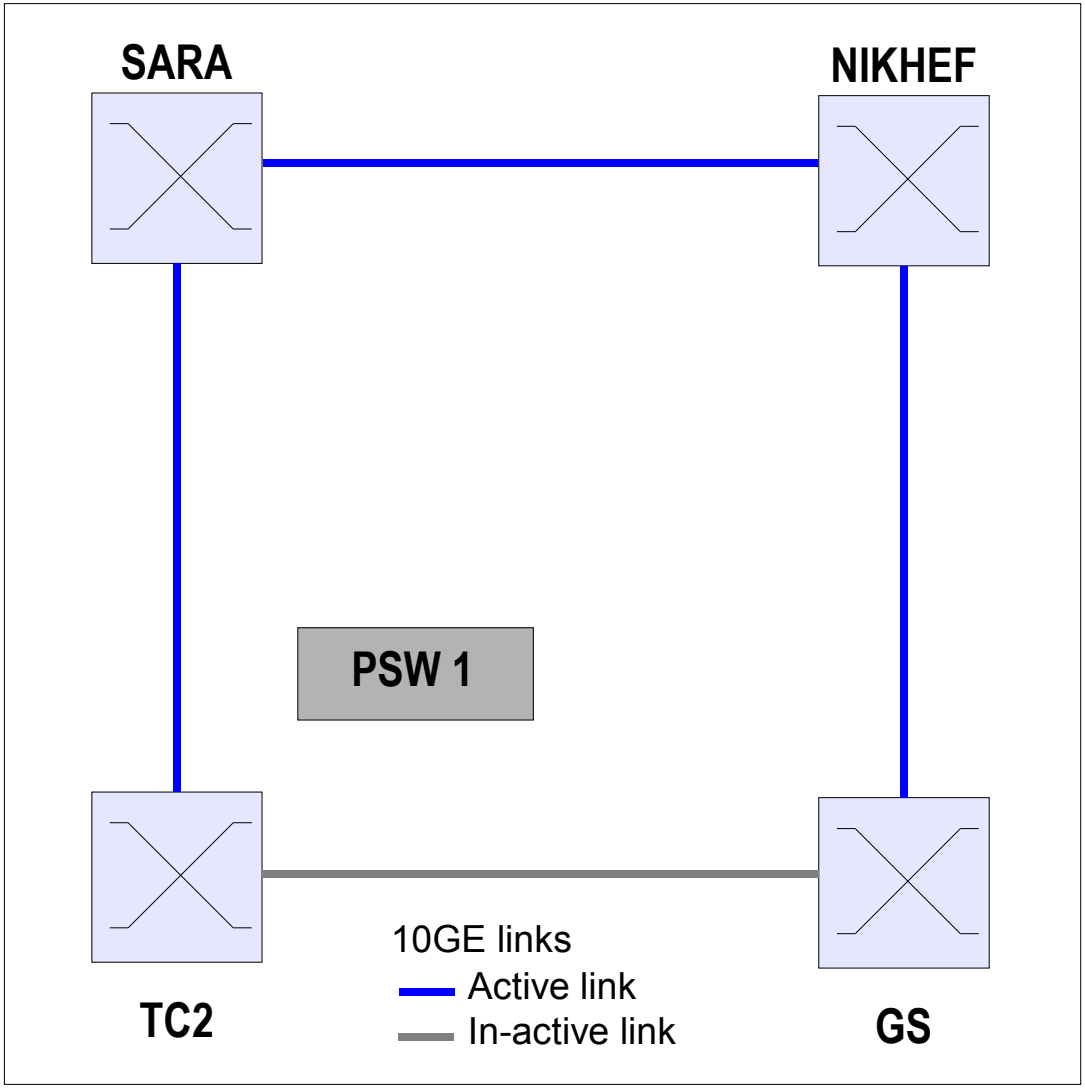

Optical Switching at AMS-IX Use Case I – Transition to Core

- Introducing the first photonic switch
- Adding some links
- Adding 1st Core switch

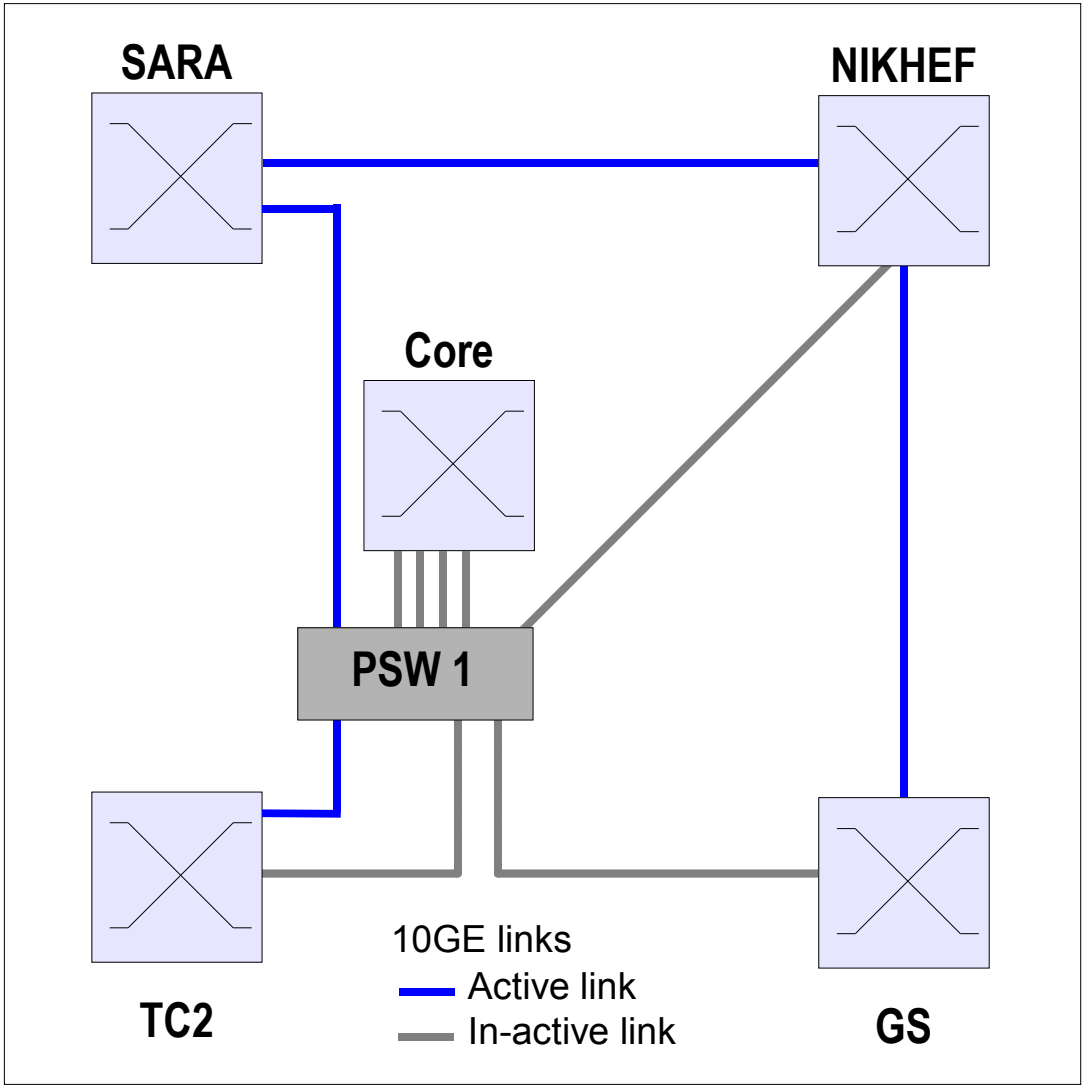

Optical Switching at AMS-IX
Use Case I – Transition to Core

- Introducing the first photonic switch
- Adding some links
- Adding 1st Core switch
- Switching over.... (you missed that!)

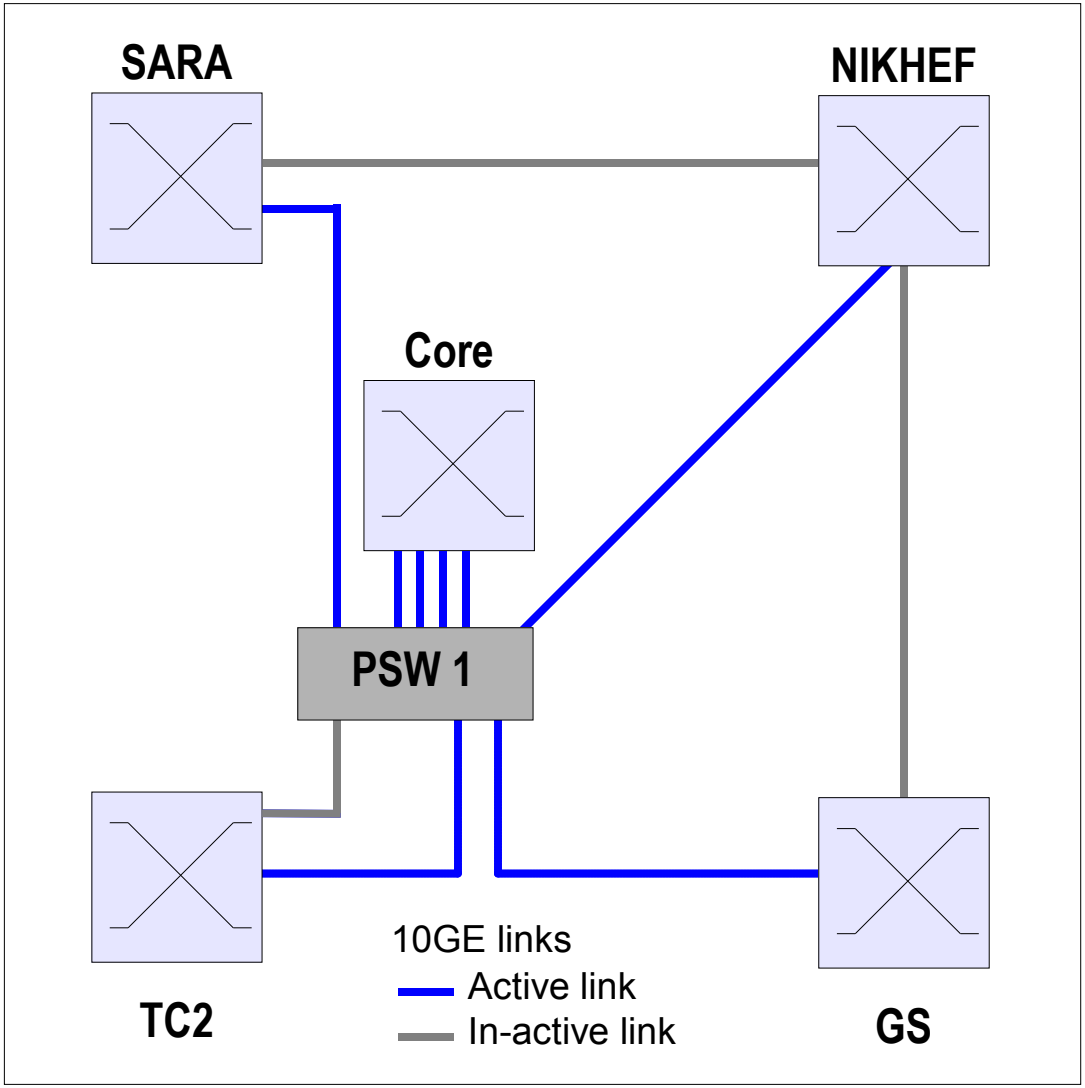

Optical Switching at AMS-IX Use Case II - Swapping Core

- Taking another angle ....
  - Same topology – different view

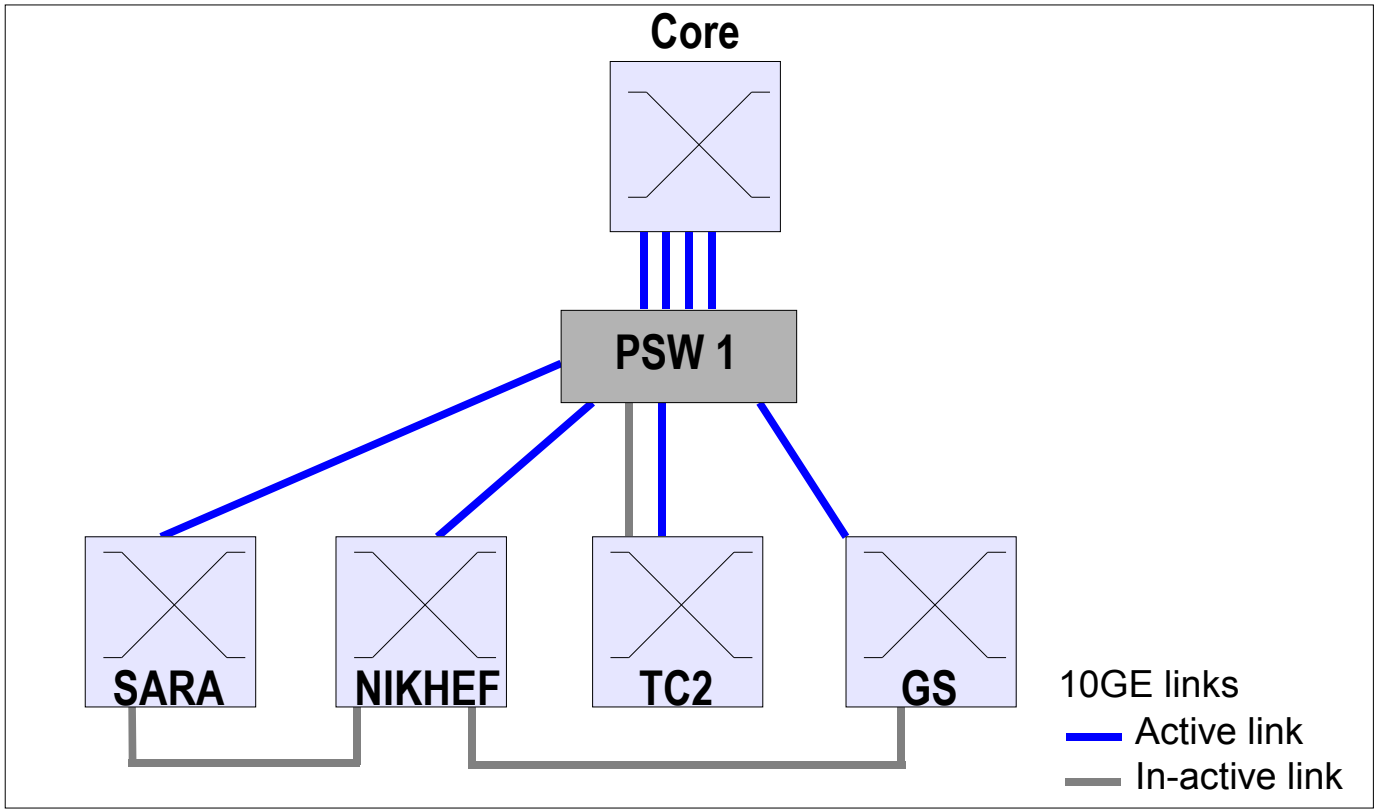

Optical Switching at AMS-IX Use Case II - Swapping Core

- Adding the backup ... and replacing it

- After the second core system came the “Mucho Grande” (MG8)
- Swapping between MG8 and BigIron (Core-2)

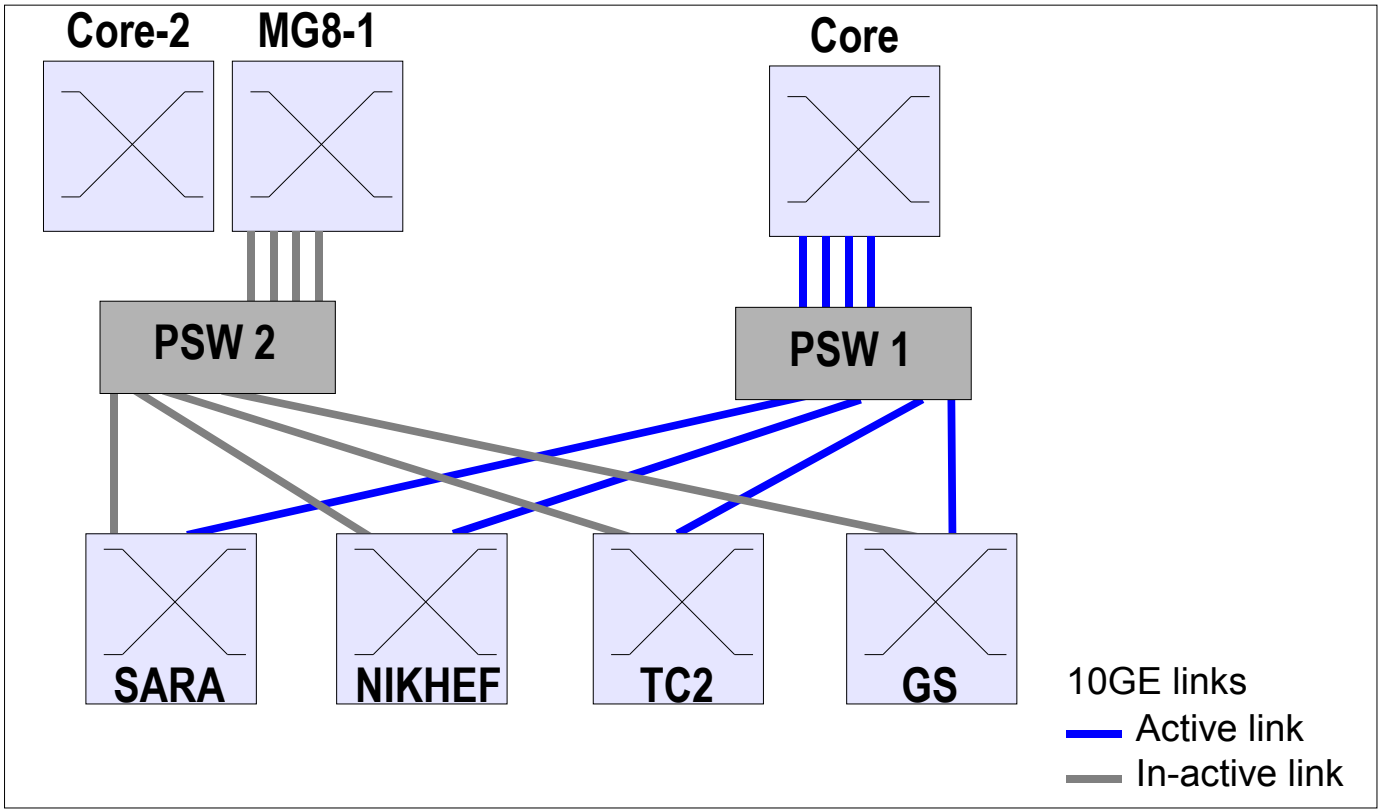

Optical Switching at AMS-IX Use Case III – 10G customers

- Current state – Two Mucho Grandes as Core
  - Only one 'active' at any time – VSRP controlled

- 10G customer connects to one Photonic switch
- Photonic switch connects it to active core switch
- VSRP state-change triggers swap-over

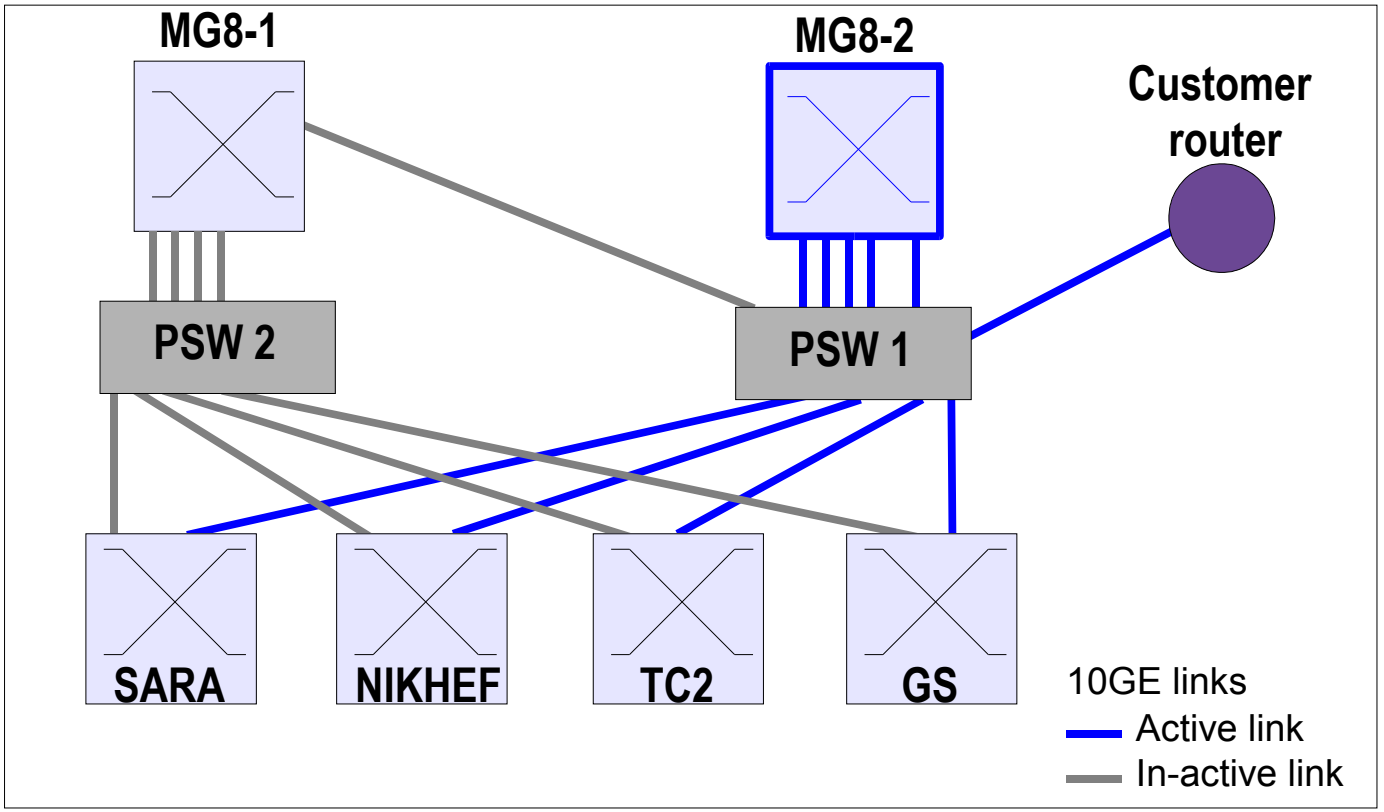

Optical Switching at AMS-IX Use Case III – 10G customers

- Current state – Two Mucho Grandes as Core

- 10G customer connects to one Photonic switch
- Photonic switch connects it to active core switch
- VSRP state-change triggers swap-over

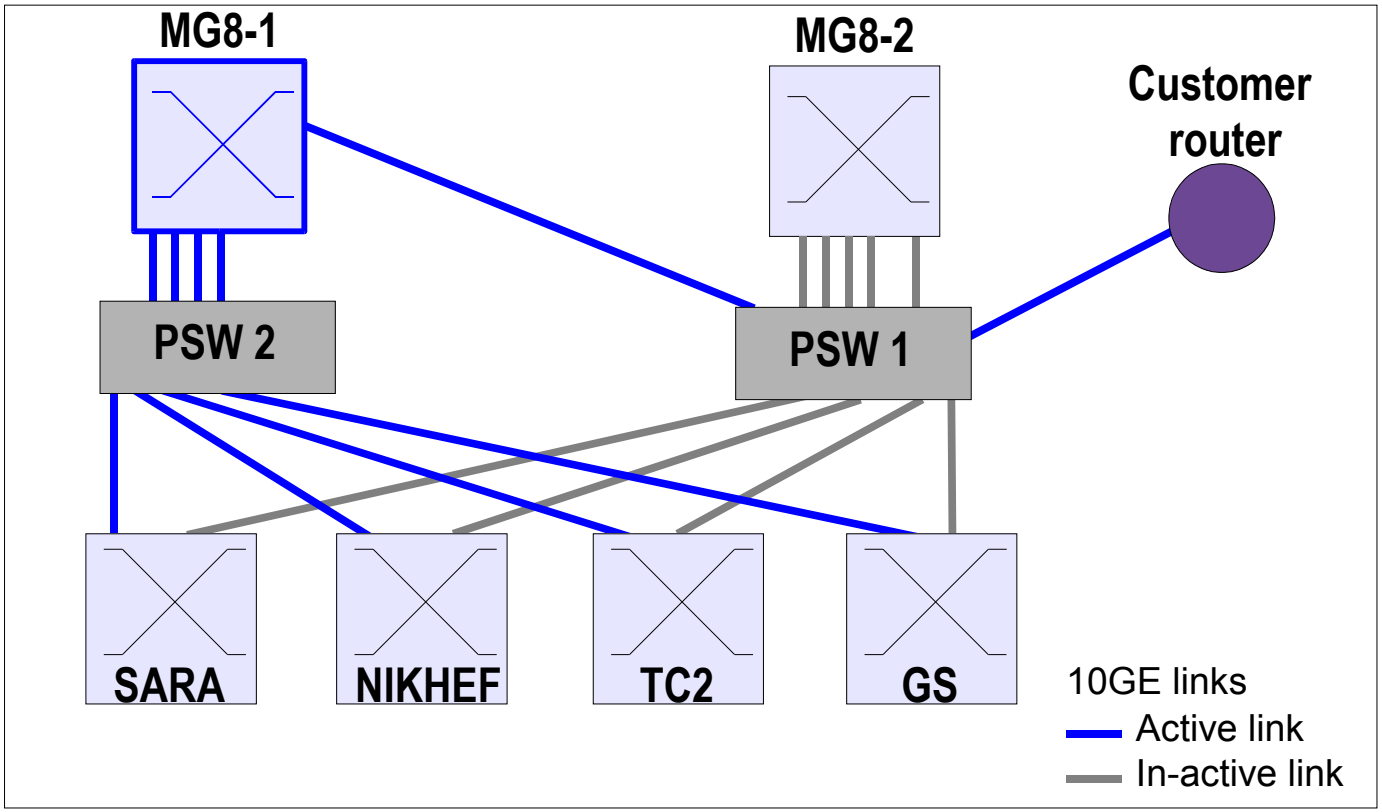

Summarizing - Are we done?

Well, will we ever be done?

- Remember Steve's Myth #6 ?
  - "100Mbps is enough"
- Today we exceed 39 Gbps peak traffic
  - Growth predictions: Q1 2006 100 Gbps peak?
- Start planning ....
  - Mucho Grande to the edge?
  - Move to DWDM and bring everything to core?
- the Photonic switches will be part of it

amsix
amsterdam internet exchange

Optical Switching at AMS-IX

Any questions?