



## D4.2 “Field Trial of OPST ring demonstrating PC virtualization services”

Status and Version:		
Date of issue:	09.01.2012	
Distribution:	Project Internal	
Author(s):	Name	Partner
Kassinis	George	PRIMETEL
Georgiades	Michael	PRIMETEL
Fullaondo	Jaime	UAM
Aracil	Javier	UAM
M O'Connor	Peter	INTUNE
Fletcher	Nigel	INTUNE
Basham	Mark	INTUNE



## Table of Contents

<b>Table of Contents</b>	<b>2</b>
<b>1 Executive Summary</b>	<b>6</b>
<b>2 Introduction</b>	<b>8</b>
2.1 Motivation and Requirements for using two Testbeds	9
<b>3 Field Trial &amp; Setup</b>	<b>12</b>
3.1 Description	12
3.2 Topology	13
3.3 Node Specifications	15
3.4 Configuration Description	16
<b>4 Deployment and Operational Aspects</b>	<b>17</b>
4.1 OPST Design & Deployment	17
4.2 Service Integration	17
4.3 Network Interfaces	17
4.4 Configuration considerations of the OPST ring network	19
4.5 Deployment Steps	20
4.6 Troubleshooting	26
<b>5 Performance Evaluation and Quality of Experience</b>	<b>27</b>
5.1 Preliminary Performance Testing	27
5.2 Application Video Captures	30
5.3 Use Case Scenarios	31
5.4 Scalability Considerations (Increase number of users and demand).	34
5.4.1 Stress Test	34
5.5 Average Workload Test	36
5.5.1 Scalability/Dimensioning Conclusions	37
<b>6 Conclusions</b>	<b>38</b>
<b>7 References</b>	<b>39</b>



## Table of Figures

Figure 1: MAINS Metro Ring/Mesh Hybrid Architecture (In Blue: Testbed Scenario 1; In Red: Testbed Scenario 2) [1].....	9
Figure 2: Field Trial Scenario at PrimeTel.....	12
Figure 3: Geographical view of testbed scenario in Limassol.....	14
Figure 4: Geographical view of testbed scenario .....	14
Figure 5: Network configuration .....	16
Figure 6: INX Beta Front Panel (Source: Intune Networks).....	18
Figure 7: Out of Band Management & Control Network .....	19
Figure 8: Photonic Manager Screen Capture.....	20
Figure 9: OPST nodes installed at PrimeTel HQ Data Centre.....	21
Figure 10: 10Gbit XFPs connecting Optical Switch to OPST nodes.....	21
Figure 11: Switches used for management purposes over LAN access, connecting all OPST nodes with management desktop PC.....	22
Figure 12: Research servers 2 & 3 .....	22
Figure 13: Test Traffic Generator.....	23
Figure 14: Field Trial equipment configuration with Traffic Generator .....	23
Figure 15: Field Trial scenario .....	24
Figure 16: Print-screen from OPST Photonic Manager .....	24
Figure 17: Creation of Services and allocation of bandwidth.....	25
Figure 18: Node-to-Node direction.....	25
Figure 19: Mean Throughput performance for large MTU's for our protocol ARQ implementations as opposed to TCP on a 10GE P2P topology .....	28
Figure 20: Achieved throughput for 3GB RAM-to-RAM transfer.....	28
Figure 21: Packet Reception progression throughout the transfer .....	29
Figure 23: Server selection. Orange smiley icon represents us, blue server is the selected server, and red servers are available servers .....	30
Figure 22: Login at MAINS Virtual PC homepage.....	30
Figure 24: Test Case scenario 1 .....	31
Figure 25: Test Case scenario 2.....	32
Figure 26: Test Case scenario 3.....	33
Figure 28: Dimensioning - CPU Load .....	35
Figure 29: Typical Workload CPU load .....	37



## Acronyms

AN	Access Node
ADSL	Asymmetric Digital Subscriber Line
API	Application Program Interface
BPRS	Bi Directional Path Switched Ring
CAPEX	Capital Expenditures
CO	Central Office
CPE	Customer Premises Equipment
COS	Class of Service
CN	Concentration Node
CS	Concentration Switch
DSLAM	Digital Subscriber Line Access Multiplexer
DWDM	Dense Wavelength Division Multiplexing
E2E	End to End
EMS	Element Management System
EN	Edge Node
E-NNI	External Network-to-Network Interface
FTTH	Fibre To The Home
GbE	Gigabit Ethernet
GMPLS	Generalized Multi-Protocol Label Switching
L2VPN	Level 2 Virtual Private Network
NCS	Network Centric Services
NMS	Network Management System,
OAM	Organization Administration and Management
ODP	OPST Design Principal
OLT	Optical Line Terminal
OPEX	Operational Expenditures
OPST	Optical Packet Switching Technology
PCE	Path Computation Element
POP	Point of Presence
QoE	Quality of Experience
QoR	Quality of Resilience
QoS	Quality of Service
REST	Representational State Transfer



SC	Sub-wavelength Concentrator
SOAP	Simple Object Access Protocol
TE	Traffic Engineering
TMF	Telemanagement Forum
TNA	Transport Network Address
VM	Virtual Machine
VoD	Video on Demand
VONO	Virtual Optical Network Operator
VoQ	Virtual Output Queue
WADL	Web Application Description Language
WSON	Wavelength Switched Optical Networks
XML	eXtensible Markup Language
XSLT	eXtensible Style sheet Language Transformations

## **1 Executive Summary**

The ICT MAINS project proposes new network architecture with sub-wavelength optical switching technologies supporting the dynamics required for the data transport plane as well as enhancements in the GMPLS architecture control plane to support these.

This deliverable provides a report on the field trials of the OPST rings and the performance evaluation of PC virtualization services running over the proposed transport solution. It will report on the Quality of Experience to real users running pc virtualization services in a Telecommunication Operators infrastructure". More specifically the field trial consisted of Primetel's Optical Fiber Network test-bed, Intune's OPST switches and UAM's Virtual PC Server application, which was used to test part of the overall MAIN's architecture and functionality.

Intune Networks has invented a new networking platform called Optical Packet Switch Transport. OPST is the complete merging of Optical Transport with a Distributed Switch and both functions operate under a single control system. The OPST nodes contain a fast tunable laser capable of tuning between wavelengths at nano-second speeds. Each node also contains a fixed wavelength filter which is the wavelength-routed address of the ports of the system. Basic operation involves reading incoming packet addresses, translating them into wavelengths and queue them in virtual output queues.

The OPST system has significant advantages over existing systems because it grants simple and rapid access to direct sub-wavelength flows of packets. This new capability of accessing sub-wavelengths in a highly responsive manner can enable Cloud Computing and Software Service applications such as the PC virtualization and the OPST beta nodes will demonstrate this capability.

Section 2 gives an introduction of MAINS Metro Ring/Mesh Hybrid Architecture. The motivation and benefits of using two separate testbeds for performance testing and validation of the proposed work is also explained. In addition to this, a discussion is included on how this is more effective in order to cover all testing logistical and technical aspects under consideration by using a laboratory testbed and a real network. A Telecom Operator can provide a real operator's network, with real customers, and obtain an idea of the Quality of Service from real users. On the other hand a university provides a lot more flexibility in terms of testing and will be utilized for testing the proposed control plane between two separate sub-wavelength technologies.

Section 3 provides a description of the field trial scenario, including a description of PrimeTel's Optical Fiber Network testbed, Intune's OPST Beta version switches and UAM's Virtual PC Server application were used in order the remaining test cases to be evaluated. Moreover this section includes a description of the topology, the detailed specification and the specs related to the interconnection configuration.

Section 4 describes the deployment and operational aspects of the field trial. The deployment is subdivided into three more detailed configuration steps undertaken at Primetel's network. Step 1 was for configuring all components used at Primetel's racks both at the HQ as well as Ayia Paraskevi premises (e.g. Servers, Switches), setup an internal LAN for the Ring Configuration, connect a PC and Laptop to the LAN. Step 2 was for positioning Intune's OPST Switches in place, connecting these to the Optical Ring network and step 3 was for the installation of the Virtual PC Application and creation of Virtual Machines for testing.

Section 5 discusses the work undertaken once all configuration and deployment in section 4 was achieved. The first part discusses the enhanced TCP-like protocol proposed by UAM



to improve on the maximum throughput performance of File and in this case Virtual Machine transfer. This is done because a high bandwidth network is set up for transferring large Virtual Machines from one server to another and hence to get maximum utilization of the offered bandwidth, the transport layer protocol e.g. TCP should not become the bottleneck in the performance achieved. Results show that the proposed TCP-like protocol provides an improvement on packet throughput as compared TCP and other protocols.

The second part of section 5 discusses the application used for the creation, modification, shut down and transfer of Virtual Machines. The application allowed in a very user friendly fashion to select a desired server for creating or transferring a Virtual Machine both for real usage as well as testing. The last phase of the work included testing of a number of use case scenarios to validate the operability of the network setup and secondly to test usability and Quality of Experience of users under different access situations.

The final Chapter describes the conclusion of the work and some migration issues that should be noted for the second field trial at the University of Essex which targets to validate interworking of two subwavelength technologies with the support of a common GMPLS plane.

## 2 Introduction

The EU ICT MAINS project proposes a new approach regarding sub-wavelength support for metro architectures that makes easier the support of dynamic and high traffic volume internet services such as high definition video streaming, social networking, 3D telepresence, Peer-to-Peer (P2P) gaming and cloud computing. [1] At present, metro network solutions which are based on electronic packet switching find it very difficult to scale at high traffic volumes. In addition, metro networks in order to support high bandwidth service depend on real-time independent wavelength allocation without exploring individual wavelength sharing.

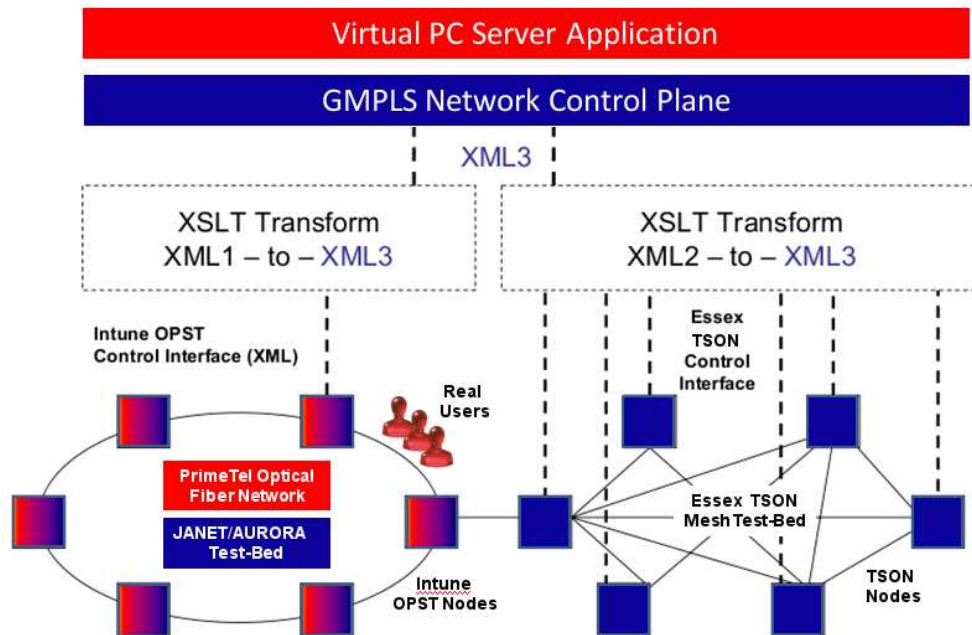
Despite MAINS project, there are also other ICT EU projects which are investigating dynamic and high traffic volume service issues for future networks that are based on either optical fiber or copper wire transport technologies.

Projects like ICT ETICS [2] raise similar issues in that although more and more optical networks are integrated in current networks and are able to provide QoS and protection guarantees, they do not sustain on-demand provisioning or dynamic changes across multiple carrier networks. The ICT STRONGEST project [3] on scalable, tunable and resilient optical networks guaranteeing focuses on extremely-high speed transport and also addresses issues with the need of ultra-high bandwidth networks, scalability, flexibility and assurance of end-to-end quality of service. GEYSERS [4] on generalized dynamic services for future networks aims at studying architectures for service-aware optical networks with a generalized network-centric approach to the end-to-end transport service delivery. FP7 NoE BONE was proposed for building the Future Optical Network in Europe: The BONE project [5] aiming to strengthen European research on optical networking through integration and collaboration amongst academic centers targeting on improving related issues.

The main goal of MAINS project is to propose a new architectural approach that will make use of sub-wavelengths. This new architecture will be based on dynamic time-shared use of single wavelengths that would be achieved with the use of ultra-fast tunable components as well as with time-shared optical network (TSON) systems. By this operators would be allowed to offer different services over independent sub-wavelengths flows. In order to make it more clear, a "full" wavelength is able to provide up to 10G bandwidth, whereas in MAINS proposed architecture, 3 operators could be given access to 2G, 3G and 5G sub-wavelength flows. For achieving this, a new approach which is going to extend GMPLS control plane is proposed. As a result, the architectural approach proposed by MAINS would allow efficient deployment of distributed application servers inside a metro network. This will offer support for dynamic bandwidth demands for services. As a final result, both CAPEX and OPEX will be reduced, whereas improved capabilities for service operation are achieved.

The new network architecture with sub-wavelength optical switching technologies [1], which ICT MAINS project is proposing, supports the dynamics that are required for the data plane along with improvements in the GMPLS architecture control plane to support these. Below, Figure 1 illustrates the network architecture overview that is proposed by MAINS. The network is consisted by two optical domains, which includes the OPST Metro Ring network and the TSON Metro Mesh network.





**Figure 1: MAINS Metro Ring/Mesh Hybrid Architecture (In Blue: Testbed Scenario 1; In Red: Testbed Scenario 2) [1]**

One of MAINS project objectives is to demonstrate transportation and switching of sub-wavelength flows across the two networks that are being administered by a GMPLS control plane, which is common in both networks. A service request received by the GMPLS control plane will result in network control commands being sent to the OPST and TSON networks in order to configure the required sub-wavelength path across the network. In order to allow the 3 systems to interwork in configuring the required sub-wavelength path across the network, the MAINS project proposes the implementation of an XML-based interface. The development of the GMPLS extension for sub-wavelength granularity, together with the XML interface for flexible, cross-technology, transport resource configurations, provide a uniform way to describe TSON and OPST services to GMPLS CP and thus provide the platform to create seamless end-to-end services among OPST metro rings interconnected with TSON metro mesh [1].

For verification and testing purposes, two separated testbed environments are going to be used such as Ring and Mesh Metro networks. Figure 1 separates in blue and red color the different functions/components. In blue color illustrates components that will be used during the first testbed using the University of Essex Lab Testbed along with JANET/Aurora research Network. On the other hand, in red color, it's the second testbed scenario, where PrimeTel's optical commercial network is going to be used. These two testbed environments are also used to verify parts of the proposed MAINS architecture in isolation and under different service loads to prove particular concepts related to traffic demand and adaptation of the architectural components to dynamic network changes [1].

## 2.1 Motivation and Requirements for using two Testbeds

This section provides a description and reasons why two separate testbeds are being selected for performance testing and validation of the proposed work. In addition to this, methodology is provided on how this will be effective in order to cover all testing aspects under consideration.

Network Research can be validated in several ways such as, analytical modeling, simulation modeling, emulation or implementation testing. For the purposes of this project, in order to assess the architecture proposed by MAINS, implementations of several testbed are being used. Implementations are being divided into field fiber testing and real consumer network implementation/integration testing. Table 1 below represents some of the main key aspects comparing research and consumer network.

	Research Network	Operator Network
Flexibility for Repeatability	✓	
Reproducibility	✓	✓
Evaluation of Quality of Service (QoS)	✓	✓
Evaluation of Quality of Experience (QoE)		✓
Real users		✓
Evaluation costs	Low	High
Connection to GEANT	✓	
Further research	✓	✓
Support Commercial acceptance		✓

**Table 1: Research vs. Operator Network Environments [1]**

The evaluation of carrying out the testbed aims to examine the effectiveness and performance improvements of the MAINS architecture both qualitatively and quantitatively. Quantitatively factor will be achieved by performing tests that have to check the performance of the proposed functionalities, control plane support, network performance under high traffic conditions, as well as Quality of Service measures. On the other hand, qualitatively factor will be achieved by user questionnaires, Quality of Experience feedbacks from real users and evaluation of performance compared with typical network service provisioning performance criteria.

Furthermore, there are two more important testbed attributes that are considered for selection: repeatability and reproducibility. Repeatability aims the ability to repeat test runs in a consistent manner. Also known as test-retest reliability it is the ability to perform the same test under the same conditions. When it comes for reproducibility, this aims the ability for researchers to reproduce results derived from previous experience, typically but not necessarily, performed by other researchers. The University of Essex/Janet-Aurora is selected for progressive tests regarding the architectural functionality.

For choosing two different testbed networks some considerations have been made and these had to do with technical and logistical aspects. In the case of a Telecommunication Operator's network the following considerations hold:

- A Telecom Operator can provide the operator's network with real customers and commercial OPST nodes.
- The ability to test Virtual PC services in a commercial environment.
- The possibility of conducting surveys and performance tests with real users to extract results from the clients with regards to quality of service/quality of experience (QoE) using questionnaires or other means you might already use.

- More costly in terms of trial execution. Requires dedicated resources for test-setup from the network, network of which is also used for consumer services.
- Support Qualitative Research
- A ring topology is already established at the Operator's premises.

On the other hand using the UEssex-JANET/AURORA test-bed:

- Has been preferred for testing horizontal integration of OPST-TSON and vertical integration of both transport solutions with extended GMPLS CP to provide end-to-end services.
- To be used for demonstration of data plane OPST-TSON interoperability.
- Can be used for verification of the control plane extensions proposed within MAINS framework.
- Multi-technology interoperable GMPLS CP.
- Can be used as a platform to test a variety of applications before selecting for real customer validation.
- A dedicated test-bed for optical network research is less costly in contacting project's experiments and verifying implementation on a stage-by-stage basis.
- Support mainly Quantitative Research

Based on Table 1 and the above discussion the following tests have been illustrated (see Table 2) and have been selected to take place in each testbed environment accordingly.

	Univ. of Essex / Intune JANET/Aurora (Research Network)	PrimeTel Network / Intune (Operator Network)
Control Plane Verification	✓	
Virtual PC Servers Testing		✓
Optical Ring Network		✓
Optical Mesh Network	✓	
Mesh/Ring Interoperability	✓	
User Experience Evaluation		✓
Commercial Use Evaluation		✓


**Table 2: Research vs. Operator Network Tests**

In this Deliverable we discuss the First MAINS's field trial which was undertaken from September – November 2011 at PrimeTel's premises. Partners involved for this primary field trial included PrimeTel PLC (Cyprus), Intune Networks (Ireland) and University of Autonoma of Madrid (UAM). For this setup we hence focused on examining the functional, operational logistical aspects at a real operator's network. The following section discusses the Field Trial setup at PrimeTel which was necessary in order to achieve the first set of field trial testing.

### 3 Field Trial & Setup

#### 3.1 Description

This section aims to provide with a description of the field trial scenario, where PrimeTel's Optical Fiber Network testbed, Intune's OPST Beta version switches and UAM's Virtual PC Server application were used in order the remaining test cases to be evaluated. The figure below illustrates the field trial scenario at PrimeTel's testbed.



**Figure 2: Field Trial Scenario at PrimeTel**

PrimeTel's existing network infrastructure in Nicosia was used for the needs of the field trial. This included facilities such as optical fiber, Data Center and PoP where the iNX Beta nodes and Virtual PC servers were being installed.

For the purpose of the trial, three OPST nodes had being used; two of them were installed at PrimeTel's HQ Data Center and one at Aya Paraskevi PoP. The distance between the two locations was an important factor for choosing those two locations, since there is a distance limitation for the Beta version of the nodes. Besides the two OPST nodes, Data Center also hosted the two Virtual PC Servers, which were plugged through a switch to the OPST nodes.

Using PrimeTel's testbed the MAINS architecture was tested in a real environment, thus it was easy to evaluate ring's performance and have a feedback regarding the quality of experience.

For the purposes of MAINS project, Intune Networks developed a new networking platform called OPST switch, however for field trial purposes they provided a Beta version of the OPST switch called iNX Beta instead of their commercial iNX 8000. OPST stands for Optical Packet Switch Transport and consists of a tunable laser that it's capable of tuning between wavelengths at nano-second speeds.

The operation of the OPST switch is to read incoming packet addresses, then translate them into wavelengths and finally put them in queue in virtual output queues. Afterwards, from the packet queues and with the aid of a scheduler, bursts are created. The scheduler is also responsible to modulate a burst onto a tunable laser transponder, where its

wavelength is rapidly tuned to the destination wavelength. Finally, the packet is sent through the ring.

During the field trial, three OPST iNX Beta Version switches were used and installed in two different locations in order to create a ring topology.

University of Autonoma of Madrid was responsible to design and develop a Virtual PC server application that aimed to achieve a new TCP protocol with better throughput yet avoiding introducing small packets into the network. In other words, the objective of this new protocol is to maximize the bandwidth utilization on a time-shared optical fiber. [1]

For this project, two servers are going to be used as virtualization servers and according to UAM requirements both servers had to be upgraded in order to fulfilled project's goals. Below, there are specifications for the two servers as used during the project.

- 1 Sun Fire X2270 Server
  - 4-core Xeon E5504@2.0GHz
  - 2-port 1 GE NIC (igb)
  - 18GB RAM
  - 1TB RAID1 Array
- 1 Sun Fire X 4150
  - 8-core Xeon E5450@3.0GHz
  - 4-port Gigabit NIC (e1000e)
  - 16GB RAM
  - 1TB RAID5 Array

### **3.2 Topology**

The first scenario for the field trial was based around PrimeTel's existing network infrastructure in Limassol, like in figure below. However, an important factor made us to change the location and scenario, thus the field trial was relocated at our premises in Nicosia. The factor had to do with distance limitation between the two locations, where the OPST nodes were being installed. According to the Beta version specifications sent by Intune, there was a distance limitation factor between the points, where the OPST nodes were going to be installed.

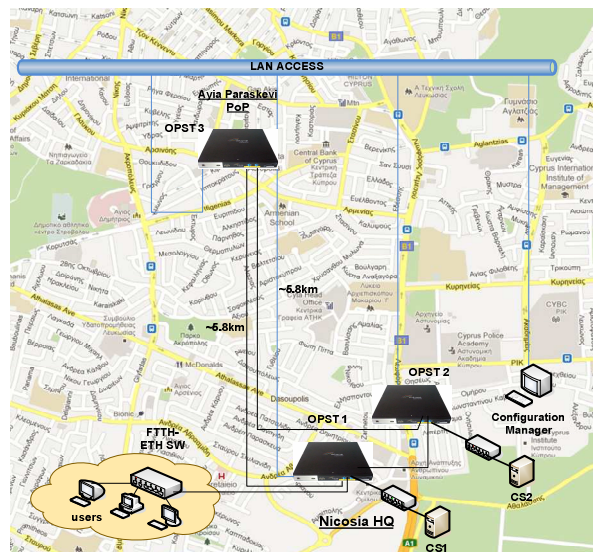
If the field trial scenario was in Limassol, the fiber route from PrimeTel's HQ Data Center to Atlantis PoP was around 25km and according to the specifications, the maximum fiber span that can be supported by the nodes between two points, must be 15km. As a result, we had to find an alternative scenario in order to establish a ring topology, where the fiber route must be less than 15km. In addition to this, there was one more limitation for the Beta version of OPST nodes and that had to do with the length of the ring. According to Intune, maximum ring length which is supported by the nodes is 60km, which also was one of the factors of relocating the scenario.



**Figure 3: Geographical view of testbed scenario in Limassol**

In order to face the distance limitation between two points, we came up with a new solution. However, with the new solution we had to relocate the scenario at Nicosia. Therefore, PrimeTel’s Nicosia HQ Data Center and Agia Paraskevi PoP was going to be used as the two locations, where the equipment for field trial would be installed. The fiber route connecting those two points is approximately 5.8km, so it was inside the distance limits supported by the Beta version OPST nodes.

Figure 3 illustrates a geographical view of the new scenario. According to this, two OPST nodes were going to be installed at HQ Data Center and one OPST node at Agia Paraskevi PoP in order to establish a ring topology using existing optical network infrastructure. In addition, along with the two OPST nodes at the HQ Data Center, two servers were also installed, which used as virtualization servers for running the Virtual PC application that UAM developed.



**Figure 4: Geographical view of testbed scenario**

In order to run the Virtual PC application and create traffic through the ring, servers must be connected with the OPST nodes. For this purpose, we needed a switch that would be able to convert the 10GE interface of the nodes to 1GE, since both servers only support 1GE network interface. Thus, to overcome this problem, Intune provided a Netgear switch with two 10GE SFP uplink ports able of making this conversion. For management purposes of the ring, all nodes had to be connected through a LAN access and being controlled by desktop PC management software that comes with them. As for the LAN access, we created a separate VLAN on our network, where all devices could be added and managed by the Configuration Manager.

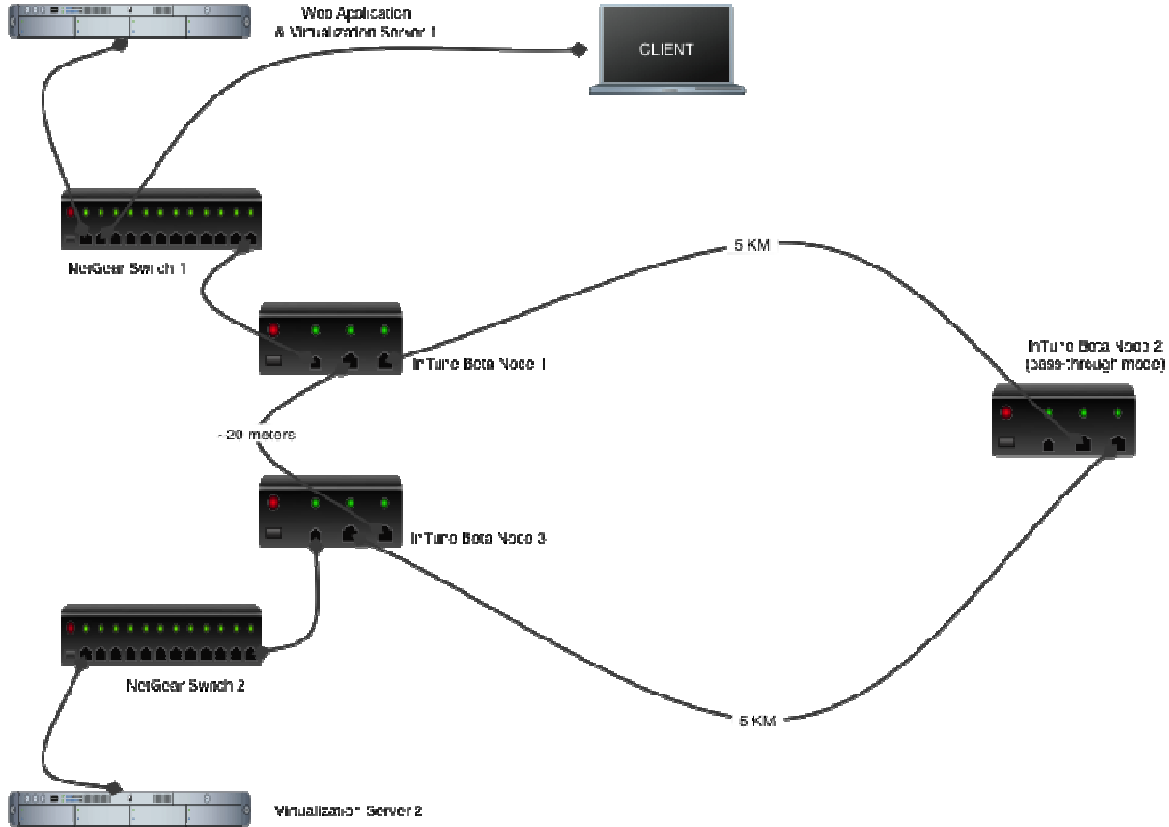
### **3.3 Node Specifications**

The node specifications are as follows:

- 2 Netgear Switches: 48 1Gbps ports - 2 10GE SFP uplink ports.
- 1 Sun Fire X2270 Server
  - 4-core Xeon E5504@2.0GHz
  - 2-port 1 GE NIC (igb)
  - 6 GB RAM
  - 1TB RAID1 Array
- 1 Sun Fire X 4150
  - 8-core Xeon E5450@3.0GHz
  - 4-port Gigabit NIC (e1000e)
  - 8 GB RAM
  - 1TB RAID5 Array
- 3 InTune Beta NXT Nodes
  - Node 1 1x10G Client Port Drop Channel 16, MAC Address 00:50:C2:80:10:2C
  - Node 2 1x10G Client Port Drop Channel 10, MAC Address 00:50:C2:80:10:2F
  - Node3 1x10G Client Port Drop Channel 19, MAC Address 00:50:C2:80:10:2B
  - Line Physical Interface – Fast Tuning Laser, 2 xLC/APC
  - Client Physical Interface – 10GBase-R (IEEE 802.3ae) 1 xXFP
- 1 Intune Networks Beta Node Photonic Manager
- 1 Intune Networks Beta Node Connection Manager

### 3.4 Configuration Description

The field trial configuration is depicted in the following picture



**Figure 5: Network configuration**

Two virtualization servers are connected to Netgear switches, which in turn connect to Intune Beta nodes 1 and 3 in the symmetrical ring of 5 Km. span per segment. The ring is formed by three Intune Beta nodes and one of them serves as a pass-through (node 2 in the picture).

The Netgear switches provide interface between the Server 1G ports and the 10G client interfaces on the Intune Beta nodes.





## 4 Deployment and Operational Aspects

### 4.1 OPST Design & Deployment

An Intune Networks Systems Engineer worked with the time prior to the network deployment to ensure, the network topology, service configuration and test scenarios were viable and achievable in the field.

Intune Networks dispatched an Engineering Support Manager and Field Installation & Commissioning Engineer to the field trial location at Primetel. Once on site the 2 Intune Networks personnel liaised with the MAINS field trial team to commence deployment of the Beta nodes. Once the installation and turn up of the nodes and the network and service configuration was complete test traffic was applied to the network. With the trial network installed and operational the Engineering Support Manager left the field trial site and returned to Intune Networks leaving the Field Installation & Commissioning Engineer on site to provide ongoing support during the life of the field trial.

### 4.2 Service Integration

The service integration and setup was fairly straight forward. Actually, for the setup, the only difference with respect to a regular LAN setup is the need to set up the ARP tables on the servers (and the clients, actually). This is due to the fact that ARP requests are sent to broadcast MAC addresses (obviously) and InTune's forwarding tables in place didn't contemplate this. You therefore need to set up the ARP tables for the remote machines you wish to reach (nodes not connected to the same Ethernet switch). Example, server A with IP's (192.168.0.30, 10.0.30) needs to reach server B with IP's (192.168.0.20 and 10.0.0.20 with respective MAC 00:DE:AD:BE:EF:00 and FF:DE:AD:BE:EF:FF addresses):

```
root@192.168.0.30# arp -s 192.168.0.20 00:DE:AD:BE:EF:00
```

```
root@10.0.0.30# arp -s 10.0.0.20 FF:DE:AD:BE:EF:FF
```

### 4.3 Network Interfaces

As it was mentioned in section 3.2, two OPST nodes were being installed at Nicosia HQ Data Center and one OPST node at Agia Paraskevi PoP. For connecting the nodes with the optical network, we used fiber patch-cords SC/SC type, where one end was connected to the nodes and the other end was connected to our Optical Distribution Frame (ODF), which is connected with optical fiber linking the two locations. Two ports on ODF were used one for transmitting purposes and one for receiving. The link loss for the two links at 1550nm was 2.12dB and 2.15dB respectively.

Moreover for providing LAN access to the nodes, we used CAT5 patch cables with RJ45 connectors in order to connect the nodes with the network through our backbone switch.

The iNX Beta platform contains all of the key elements of OPST technology, packaged in a compact prototype equipment format, for deployment in lab test facilities.

All of this innovative OPST technology is re-used for the commercial grade platform, iNX8000, with alternative packaging to provide the levels of resilience and operational readiness that a network operator would require for commercial deployment.

The implementation of the iNX Beta platform contains and demonstrates that the fundamental technical capabilities of OPST can be tested and proven.

¡Error! No se encuentra el origen de la referencia. below details the front panel of the iNX Beta platform and indicates the location of the various external interfaces provided, and described below.

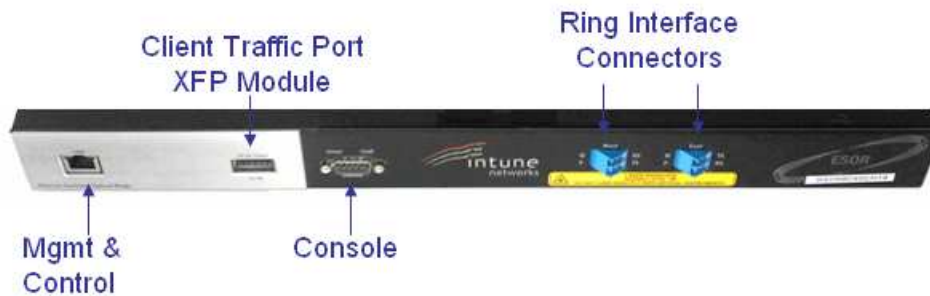


Figure 6: iNX Beta Front Panel (Source: Intune Networks)

### Power

The iNX Beta chassis is equipped with two load-sharing power supply units that operate from nominal -48V DC inputs, using two rear-mounted TYCO 6609075-4 connector.

### Client Traffic Interface

An XFP module presents the client-facing 10Gbps Ethernet interface as per IEEE 10GE-BaseLAN definitions. The following vendor variants of XFP module has been tested as part of the iNX Beta development program.

Finisar - FTLX1411M3, 1310nm SMF, 10GBase-LR (10km) with LC Duplex connector

Please note: Whilst the client interface software of the iNX Beta platform is configured to permit other XFP vendors and variants to be equipped, their use is currently **not** supported as this is an untested configuration.

### OPST Ring Interface

The iNX Beta chassis has two LC-Type optical fibre terminations to connect the OPST transmitter and receiver into a physical ring topology that forms the shared fabric of interconnectivity with the other Beta nodes in the ring.

The fibre specification required to link the iNX Beta Nodes together is ITU G.652 [as per normal industry practice] for single mode fibre (SMF) network installations.


The OPST ring interface carries both the Ethernet traffic on the ring and a proprietary control channel for the signalling of real-time traffic demands between the Nodes.

The beta system will support at 60km circumference ring network with 5 nodes deployed, and this approximates to 12km per span. The PrimeTel network consisted of 2 colocated nodes with a third at a distant site 5km away, so total circumference was 10km approx., There was no measurements made on the actual loss in the spans in Cyprus, however we would typically expect to see 0.25 dB/km for SMF. The loss per fibre to fibre connection of 0.15dB per connection would be typical.

### Management & Control

An RJ45, 10/100BaseT Ethernet interface provides out-of-band connectivity for commissioning, control and management information to each of the iNX Beta Nodes via

the PC GUI 'Photonic Manager' and 'Connection Manager' applications. This is illustrated in **¡Error! No se encuentra el origen de la referencia.**Figure 4 below.



**Figure 7: Out of Band Management & Control Network**

### Remote Access

The PC hosting the iNX Beta applications can be remotely accessed for support purposes by utilising a second, separate network connection on the host platform (not shown in the figure above). Further details on this can also be found in an Appendix to Reference 2 with a description of the various firewall restrictions and security provisions that are appropriate.

### Console Interface

A 9-pin RS232 serial console port is provided for the purposes of engineer level command-line operations and certain aspects of Node configuration. These operational processes are detailed in Reference 2. Remote access to this interface can also be provided.

#### 4.4 Configuration considerations of the OPST ring network

The configuration considerations were fibre lengths between nodes and patch cords. The Optical manager automatically optimizes the calibration required bases on optical loss measurements between nodes.

Screenshot of Photonics Manager shown below in **¡Error! No se encuentra el origen de la referencia.** below:

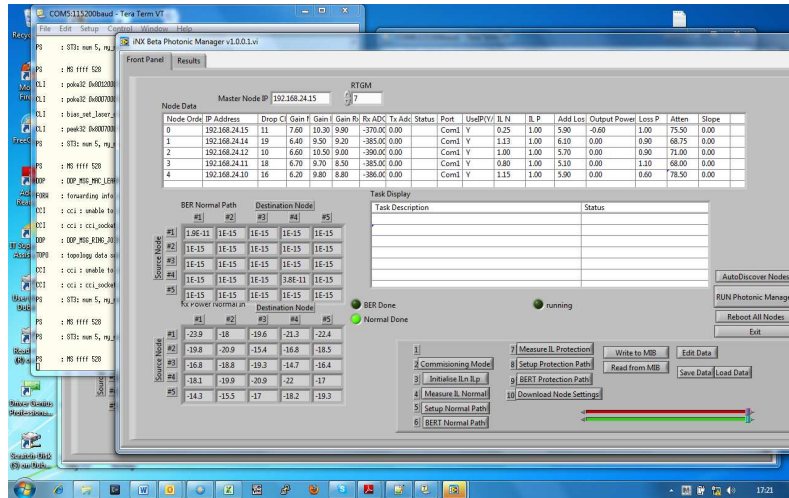


Figure 8: Photonic Manager Screen Capture

The iNX Beta connection manager required the MAC Addresses of the Nodes, Servers and Clients for configuration of the forwarding tables

#### 4.5 Deployment Steps

We divided the deployment into 3 major steps:

- Step 1 was for configuring all components used at Primetel's racks both at the HQ as well as Ayia Paraskevi premises (e.g. Servers, Switches), setup an internal LAN for the Ring Configuration, and connect a PC and Laptop to the LAN.
- Step 2 was for positioning Intune's OPST Switches in place, connecting these to the Optical Ring network.
- Setup 3 installation of the Virtual PC Application and creation of Virtual Machines for testing.

The steps are described below in more detail.

##### Step 1: Testing of servers

Figure 9 shows two Intune nodes installed at the Nicosia HQ, Shown in Green are the ring connections going out on East on the Top Node (Node 1) and returning on West at the Bottom Node (Node 3). These two nodes are interconnected back-to-back. The East connection goes out towards Node 2 which is installed at the second site of Primetel (Ayia Paraskevi), 5.5Km away, and returns back to Node 3 on the West optical port.




**Figure 9: OPST nodes installed at PrimeTel HQ Data Centre**

Figure 10 shows 2 10Gb XFPs which are used to connect from the Optical Switch to the Intune Node and provide client connectivity, both for transmitting and receiving.




**Figure 10: 10Gbit XFPs connecting Optical Switch to OPST nodes**

Figure 10 also shows the Ethernet ports used for management and configuration purposes, initially for the optical calibration, to setup the services for each of the clients, and to monitor any LAN traffic services etc.



**Figure 11: Switches used for management purposes over LAN access, connecting all OPST nodes with management desktop PC**

Figure 11 shows the two NETGEAR switches connected with a 10Gbit output from SFP+ to the client port towards the Intune Node. For each client, i.e. for each MAC address there is 1Gb Ethernet connectivity provided.



**Figure 12: Research servers 2 & 3**

Figure 12 shows the two servers provided by Primetel namely Research 2 and Research 3, with each of the Ethernet ports configured with a corresponding MAC addresses, and which are in turn connected to the NETGEAR switches shown in Figure 12.

**Step2: Testing of OPST nodes and server using Traffic generator (all nodes on 1 site)**

At this step, all three iNX Beta nodes were tested in one location. Spirent Traffic Generator was used to self-test Beta Node Optical ring with 20Mbps traffic load with fixed frame length of 128 Bytes. Full mesh traffic was tested to/from each node using the Traffic Generator to simulate servers and clients

Traffic Generator screen shots shown below.

State	Active	Name	Tags	Index	Controlled By	Traffic Pattern	Type	Tx Port	Rx Port	Traffic Group	State	Load	Load Unit	Frame Length Mode	iMIX Distribution	Fixed Frame Length
Ready	<input checked="" type="checkbox"/>	To Server 1	Click to ad...	0	generator	Pair	Port	Port //1/1 (...)	Any		Ready	20	Mbps	Fixed		128
Ready	<input checked="" type="checkbox"/>	To Server 2	Click to ad...	1	generator	Pair	Port	Port //1/1 (...)	Any		Ready	20	Mbps	Fixed		128
Ready	<input type="checkbox"/>	Original to...	Click to ad...	2	generator	Pair	Port	Port //1/1 (...)	Any		Ready	10	Percent (%)	Fixed		128

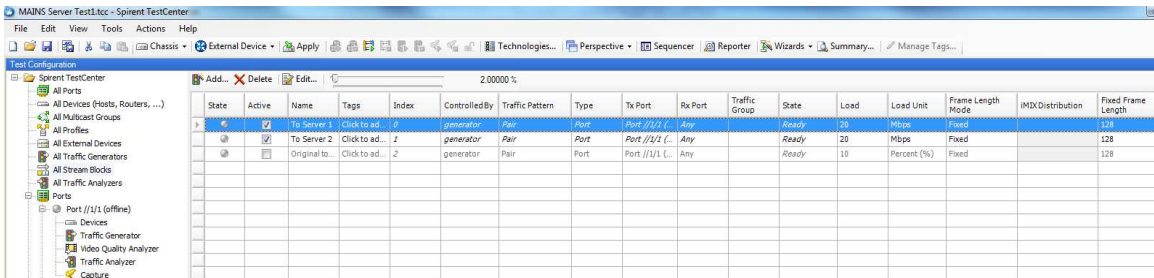
  


Figure 13: Test Traffic Generator

Figure 13 shows a three-node ring formed using three Intune beta node switches configured via a Private LAN setup for management. What we have is a 5km span connected to Primetel’s second site, with the third node been setup on the remote site. Connected as clients on node 1 and node 2 2 switches were also connected acting as a 1Gb LAN to a 10Gb Optical interface. Connected to those switches two servers were connected, which were used for the creation/deletion and modification of Virtual Machines for demonstration. A PC was also connected to the LAN to support configuration and management. Moreover client PCs/Laptops were connected to the NETGEAR switches and where used as users to achieve our use case scenarios (see Figure 14 below).

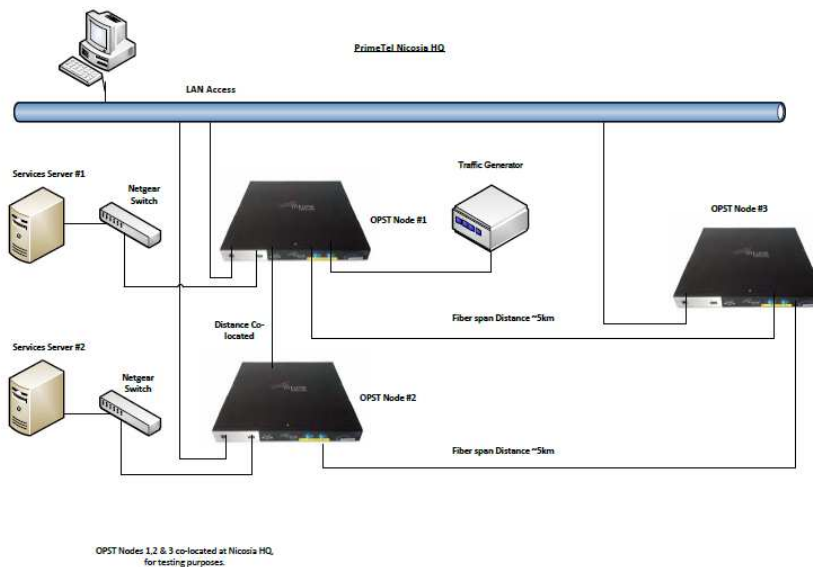
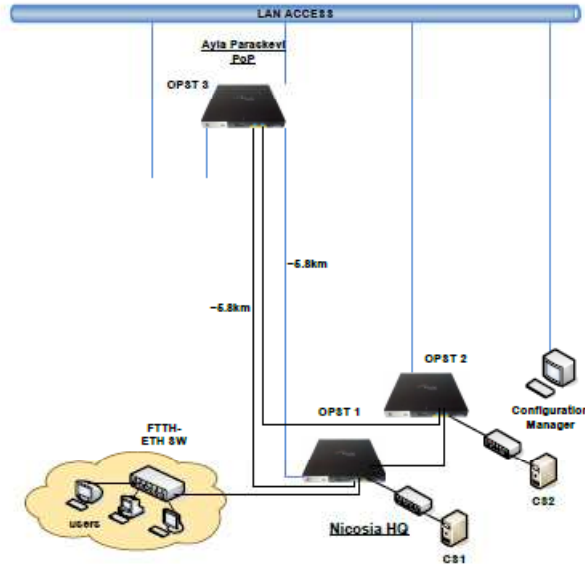


Figure 14: Field Trial equipment configuration with Traffic Generator

Normally when the three-node ring was powered up, the power levels had to be calibrated optically. Figure 2 shows the three nodes with their associated IP addresses for management, the drop channel of each node, 17, 16, 19, respectively, with various optical values for attenuation, insertion losses, and the gain for each optical section connected to the nodes. The configuration manager in Figure 2 also allows for a bit error test, where PRBS test data is generated before action data traffic is allowed on the ring.

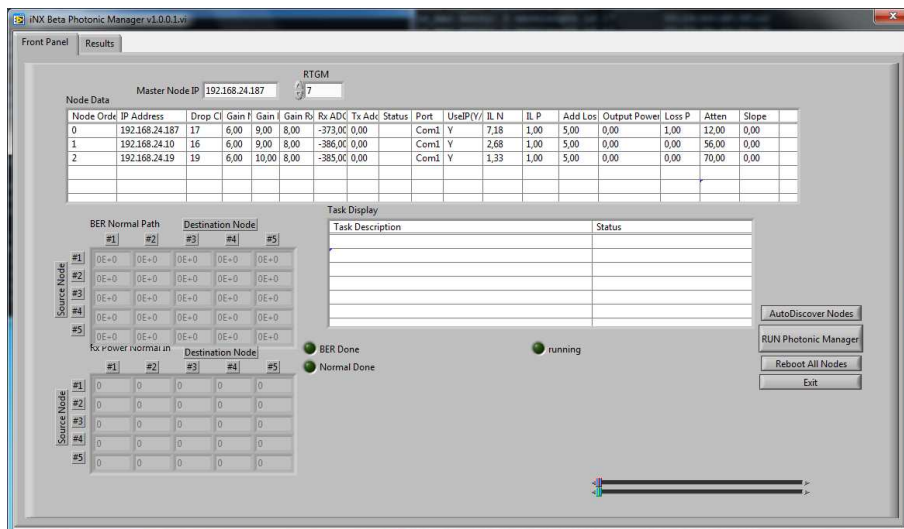
**Step3: Final Testing of Ring Network**



**Figure 15: Field Trial scenario**

In this step the Traffic Generator was removed and node 3 was relocated to the remote location. At this stage Servers and clients were connected to test the trial configuration.

Once the setup was in place all nodes on the rings were rebooted once. When that is done a connection to the master node is made, place each node into service mode, and also configure their MAC ID's (e.g. 00:50:c2:80:10:24), along with a topology number. Each end point of each node has a set of MAC addresses for any client that is connected to them. So for example node 1 has a set of 4 MAC addresses, 2 of them used for the server, and two for different client laptops. Node 2 has no set MAC addresses as it is simply used as a pass through node, and Node 3 has 2 MAC addresses for Server 2.



**Figure 16: Print-screen from OPST Photonic Manager**



Each of this connection in default mode will provide you with best effort traffic. So a number of steps in the configuration setup where undertaken to setup more realistic download connections. So Virtual Connections were setup with 100Mb bandwidth as highest priority, from every source MAC to each destination MAC. To create these services this connections where setup from node 1 to all other nodes and from node 3 to the same connection back to node 1 to complete the ring. Figure 17 show the source allocated bandwidth and the destination allocated bandwidth.




Figure 17: Creation of Services and allocation of bandwidth

Figure 18 shows the how the Node to Node configuration and forwarding was achieved.

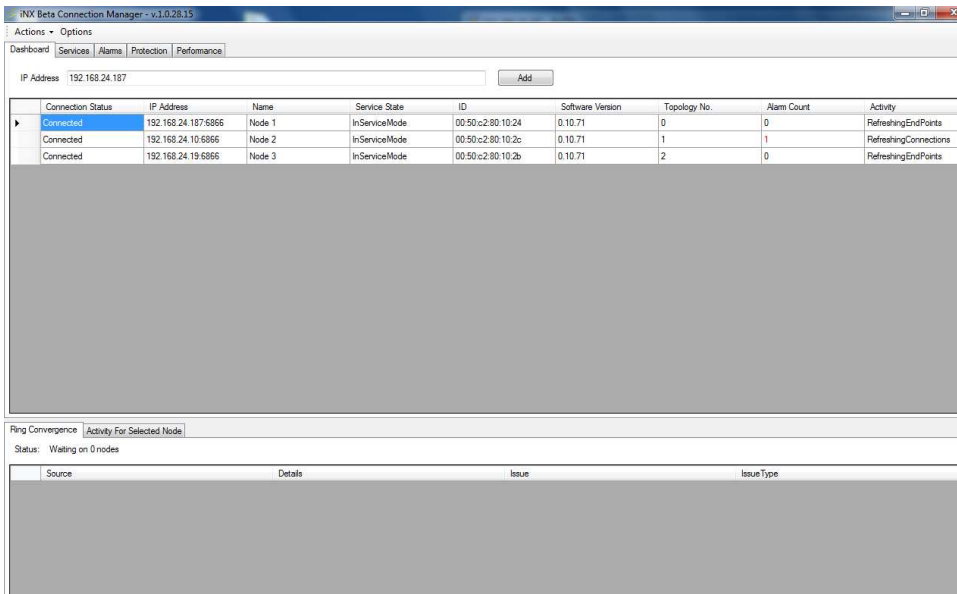


Figure 18: Node-to-Node direction

#### **4.6 Troubleshooting**

As for the VPC application troubleshooting, we found no major difficulties in the field trial setup. However, the fact that the servers had no Internet connectivity implies the password recovery functionality for the Virtual PC service does not work: Email must be sent to reset/recover the password. Additionally, this also limits the Virtual PC service ability to download the installation ISO's if they're unavailable on the virtualization server. Therefore, because we only had the installation ISO for Ubuntu 10.04, that is the only OS effectively available for VM's.

During the commissioning of the Beta nodes the following issues were identified and resolved:

- LCPC to LCAPC fibre patch required for interface from Primetel backbone ring connection. Beta spec not clearly outlined – only shown as “LC Type”. Intune supplied required patch fibres and adapters.
- One Beta node not transmitting – forwarding table reconfigured to other nodes to reduce impact to installation. Later debug showed 10G Ethernet switch with faulty SFP+ - suspect this was cause of failure.
- Insertion loss shown from 2<sup>nd</sup> Primtel site at Ayia Paraskevi between nodes. On investigation damaged LC to SC patch was replaced on ring connection allowing successful Optical calibration and topology of ring.
- Replaced 10G Ethernet switch to resolve bad receive to Server2. Faulty SFP+ suspected to be cause of failure.

## 5 Performance Evaluation and Quality of Experience

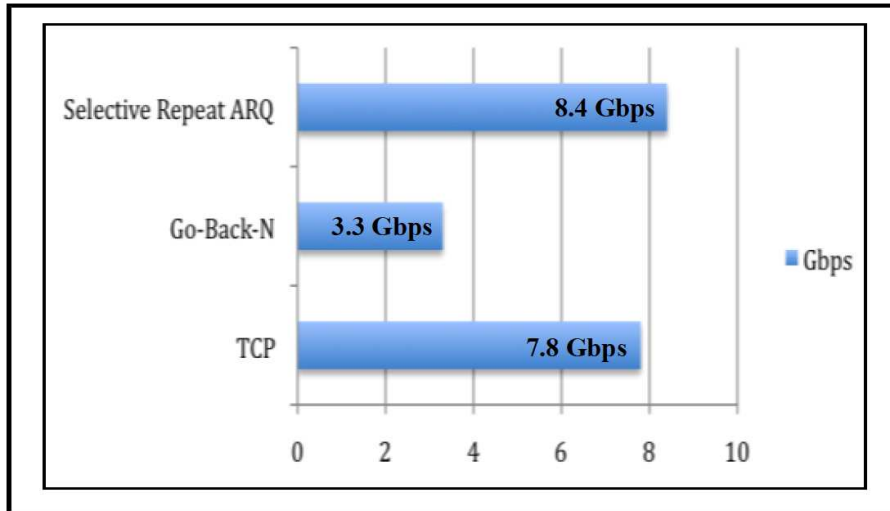
The service is in essence unchanged with respect to the demo shown during July's plenary. Changes have been made in order to integrate the Virtual PC service with the MAINS project website. In addition, efforts were made to improve the interactivity of the website as well as the usability. This was achieved by incorporating AJAX mechanisms, so that the web application will have a more desktop-like feeling.

### 5.1 Preliminary Performance Testing

A Virtual PC service to be delivered to mobile users requires that the network should have certain parameters such as minimum latency, jitter or transmission speeds with acceptable threshold values. Note that when we refer to mobile virtual PC service users we mean users who will demand the service from different locations, and not necessarily users who will be consuming the service while roaming.

A user who intends to use this virtual PC service will only tolerate low delay before growing frustrated, which denotes that the Virtual Machine will execute on the edge node. This node should be the nearest to the user's location in order the service to be commercially feasible and acceptable. The sizes of the VM image in regards to MAINS project will be up to 60GB. Having in mind the conditions mentioned above, the VM's should be moved as quickly as possible from the previous edge node to the current one using the optical data plane on the testbed. The time invested in the transfer will be perceived by the user as part of the "boot-up" time, so we must therefore do our best to keep these transfer times at a minimum.

UAM designed a new transfer protocol, which aims to achieve a TCP-like throughput avoiding introducing small packets into the network, thus it maximizes the bandwidth utilization on a time-shared optical fabric. Demanding a time-slot simply to acknowledge correct transmission seems like overkill if we can guarantee reliable transmission otherwise. Additionally, maximizing one-way communication and packet traffic, which is in essence what a file transfer should be, will not just benefit the data plane by hopefully reducing the number of packets in the optical fabric and ensuring these are as often as possible of MTU size, but could also minimize stress on the control plane by reducing the requests for lightpaths. Using TCP as our measuring stick we have tested two different ARQ techniques for the protocol: go-back-N and selective repeat. At the moment, results achieved in a point-to-point transfer of a 3GB iso over Ethernet (as tested in both copper and optical connections) have yielded the results shown in Figure 19) [1].

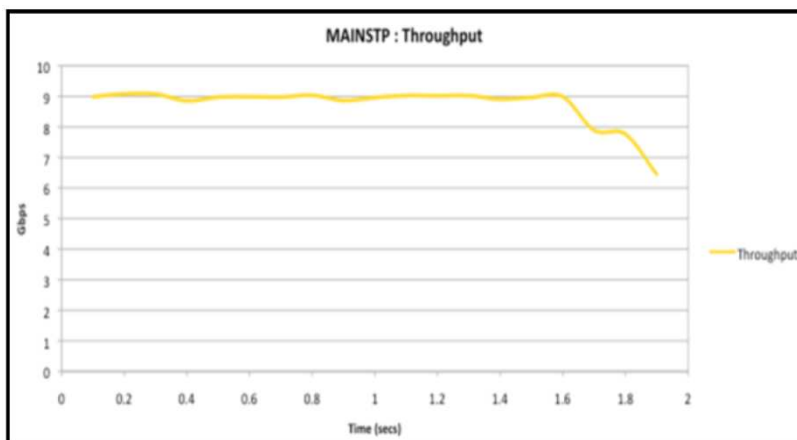


**Figure 19: Mean Throughput performance for large MTU's for our protocol ARQ implementations as opposed to TCP on a 10GE P2P topology**

Figure 20 below, indicates that we were able to maintain a sustained 9Gbps throughput during the transfer; however this can be improved during the handshake phases. Due to time lost during the handshakes the mean throughput drops to ~8.5 Gbps.

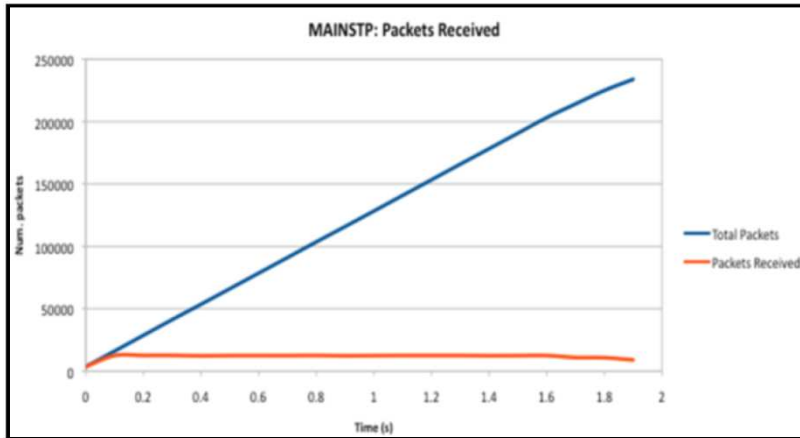
Figure 21 shows packet reception progression throughout the transfer. These results would allow us to transfer a 60GB image in just over 54 seconds, which is an acceptable waiting time for the user. Note that selective repeat ARQ is running in user space whereas the TCP runs in kernel space. Consequently, additional improvement is expected when further elements of the selective repeat ARQ are executed in kernel space as well.

The proposed optical network architecture by MAINS suits a number of services and it is well fitted for the virtual PC as will be used in the field trial using the second scenario, but also video-on-demand services, storage-on-demand. These services require moving great volumes of information (virtual images, HD media content) through the metro network because the user experience benefits greatly when the latency and jitter are minimal.



**Figure 20: Achieved throughput for 3GB RAM-to-RAM transfer**

For instance, in VOD having the media streaming to the user from its nearest edge node is pivotal to reduce factors which hurt user experience such as buffer under runs to a minimum.



**Figure 21: Packet Reception progression throughout the transfer**

The distribution of the content to the edge node as opposed to accessing a more centralized service will be beneficial to the core network by reducing the amount of traffic that is transmitted through it. The new transport protocol developed at UAM is a lightweight protocol and it could be used to maximize the efficiency of transmissions through the metro network core. It is important to note that the proper dimension of the virtualization servers which will lie on the edge nodes is essential. The requirements for desktop user are typically between 512MB and 1GB RAM, whereas a single core should be able to serve 4 virtual PCs. With these estimates, assuming a homogeneous geographical population distribution, we can expect around 70 users per VM node. To actually test the virtualization service on a real world scenario would allow us to gain further, and necessary experience towards providing a consumer-grade virtualization service which delivers a satisfactory user experience.

## 5.2 Application Video Captures

The following pictures show snapshots of the running application in the field trial:

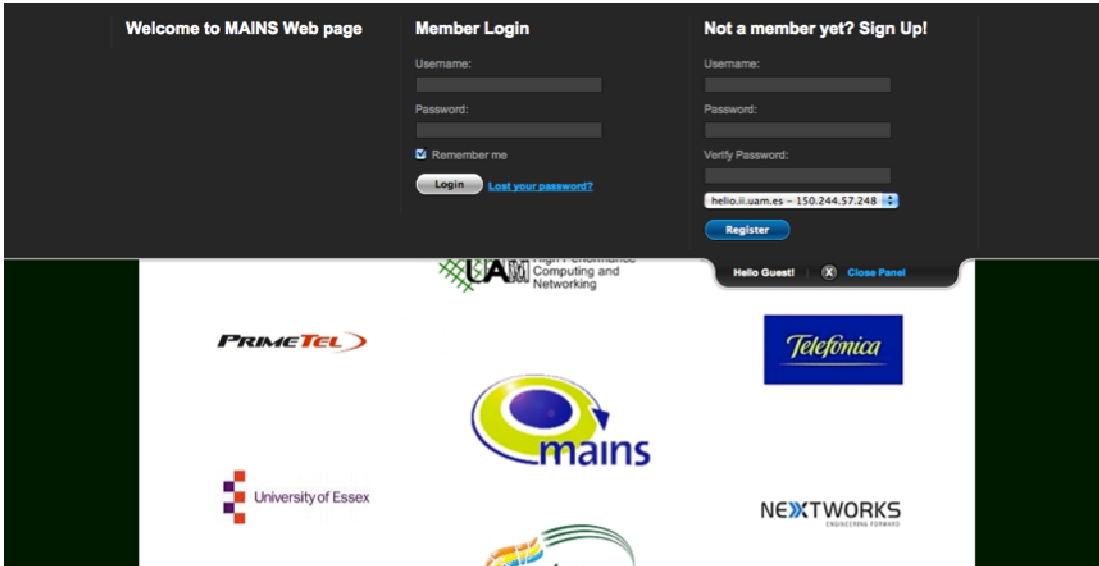


Figure 22: Login at MAINS Virtual PC homepage

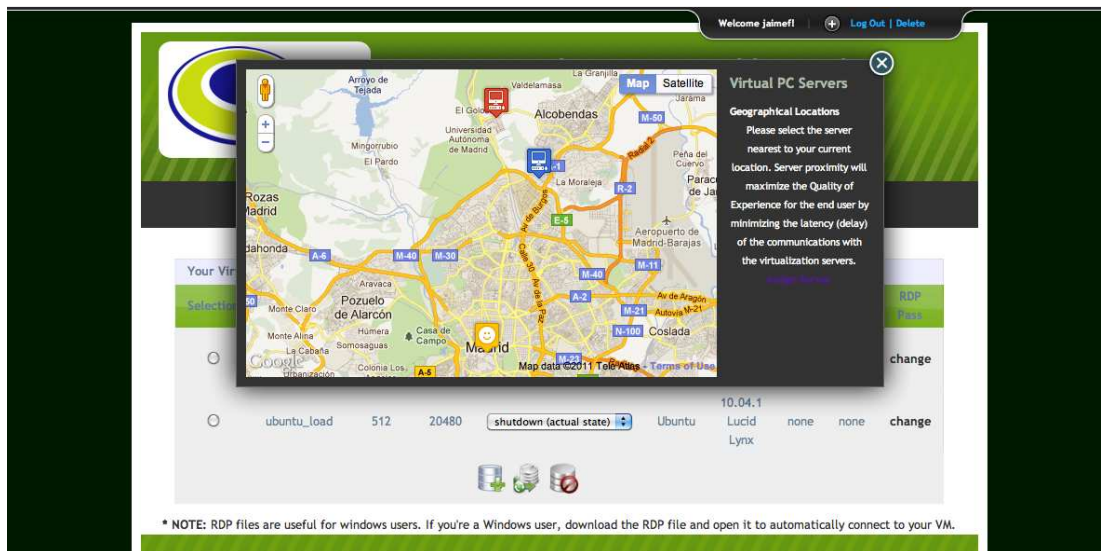


Figure 23: Server selection. Orange smiley icon represents us, blue server is the selected server, and red servers are available servers

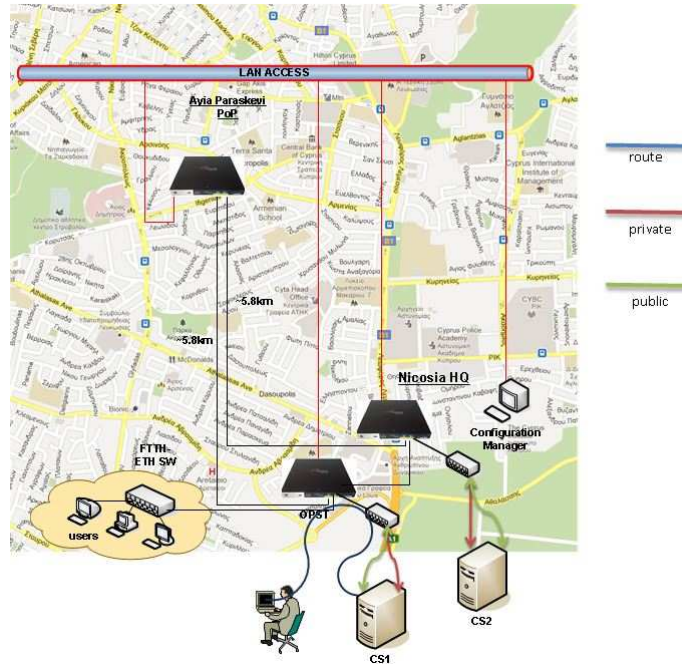
A number of Virtual Machines were created on a single Laptop connected to the prototype network with real user interfaces.

With the Virtual Machine Application we were in position to create a number of Virtual Machines, shut them down, modify them and if necessary transfer them where necessary from one server to another. Figure 10 shows the user friendliness of the application where one is able to indicate on a map the selected destination and hence server location.

### 5.3 Use Case Scenarios

Three use case scenarios were used to test the operation of the proposed Optical Ring prototype.

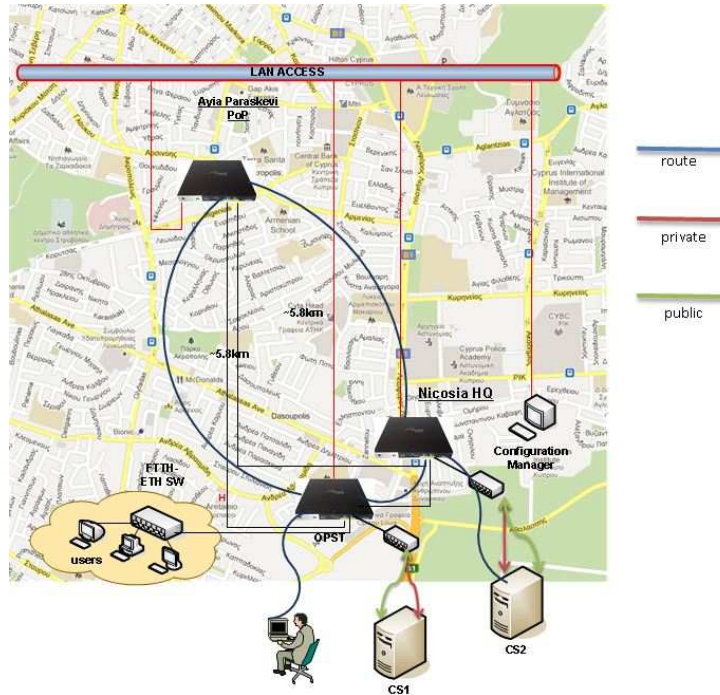
- Scenario 1: User/Multiple Users creating and accessing a Virtual Machine from a Local Server on the ring.



**Figure 24: Test Case scenario 1**

Quality of Experience of Scenario 1: In this use case scenario, the user is accessing his virtual machine from local server on the same network and hence experiences optimum performance.

- Scenario 2: User/Multiple Users creating and accessing a Virtual Machine from a Remote Server on the ring.

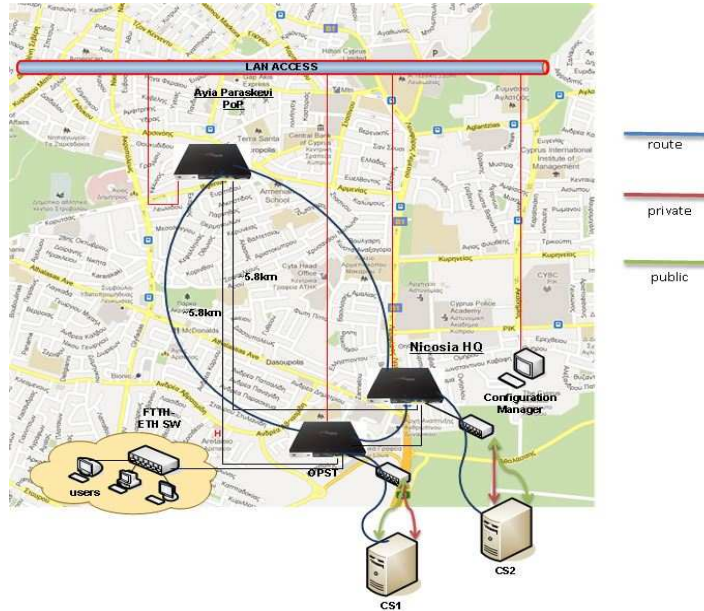


**Figure 25: Test Case scenario 2**

Quality of Experience for Scenario 2: In this scenario the user is accessing his virtual machine from remote server located approximately 11km away. The quality level experienced by the user was not much different to the first scenario with very minor jitter experienced due to the transport layer protocol.



- Scenario 3: Virtual Machine transfer from server to server on the ring.



**Figure 26: Test Case scenario 3**

Quality of Experience for Scenario 3: In this scenario, we tested the transfer of a virtual machine from a remote to a local server, while a mobile user changed its point of access from one location to another. The handover time experienced by the user on the application level was mainly due to the copying time of the virtual machine caused by the transport layer protocol.

## 5.4 Scalability Considerations (Increase number of users and demand).

One of the key tasks to perform prior to the deployment of a service is a dimension and scalability study. In order to offer an acceptable service it proves key to have an estimate, as accurate as possible, of the computational resources necessary to attend a certain demand. In essence it is no different to finding the point of equilibrium in a demand and supply graph.

Given the virtualization software chosen, Oracle Virtualbox 4.1, we had to evaluate the usage of memory, CPU, network and hard drive resources. A good starting point was Sun/Oracle's deployment documentation<sup>1</sup>. According to their estimates:

- 20 virtual desktops per core.
- 0.15Mbps per virtual desktop.
- Virtualization host requires 2GB RAM for host operations
  - And the additional RAM for every VM running on the host.

$$RAM_{host} = 2GB + \sum_{i=1}^n x_i \times RAM_{VM_i}$$

- Enough storage for every VM and host.

$$HD_{total} = HD_{host} + \sum_{i=1}^n x_i \times HD_{VM_i}$$

$$HD_{total} = HD_{host} + HD_{SAN}$$

It is typically split into the actual host hard drive and a SAN where the VM's are usually stored. This however is not the case for our field trial deployment because due to a lack of resources a SAN will not be available for storage.

At first glance, those figures look a bit optimistic. We decided to carry out our own tests to evaluate the performance. About half a dozen volunteers helped conduct these tests, each with up to three virtual machines. The users used the web interface to create their machines and perform the respective tests. Unfortunately, although we wanted to perform the 1-20 VM's/core we were unable to do so due to RAM restrictions on the test servers, more on this later.

The servers used for the tests had the following specs:

- 4 core ('8 cores' with HT) i7 860 @ 2.8 GHz
- 8 Gb RAM.
- 200 Gb HD.
- 5.6 Mbps upload bandwidth.

### 5.4.1 Stress Test

Unfortunately it's rather hard to evaluate virtualization numerically; we have to base our evaluation on a subjective perception of service quality and user experience. We also quickly evaluated memory consumption and cpu load while performing these tests to try and spot a relation between QoE and resource usage. To stress the virtual machines, and

---

<sup>1</sup>Oracle/Sun. Deployment Guide [8].

the host, we wrote a simple factorial calculation program which would load the CPU to 100%.

Since we also needed to also evaluate the QoE, we resorted (concurrently) to a common yet reasonably demanding application: **youtube**. Video comes quite handy when evaluating the quality of the experience because it can quickly reflect a decreasing QoE during playback. We spawned VM's gradually: 1-12 VM's.

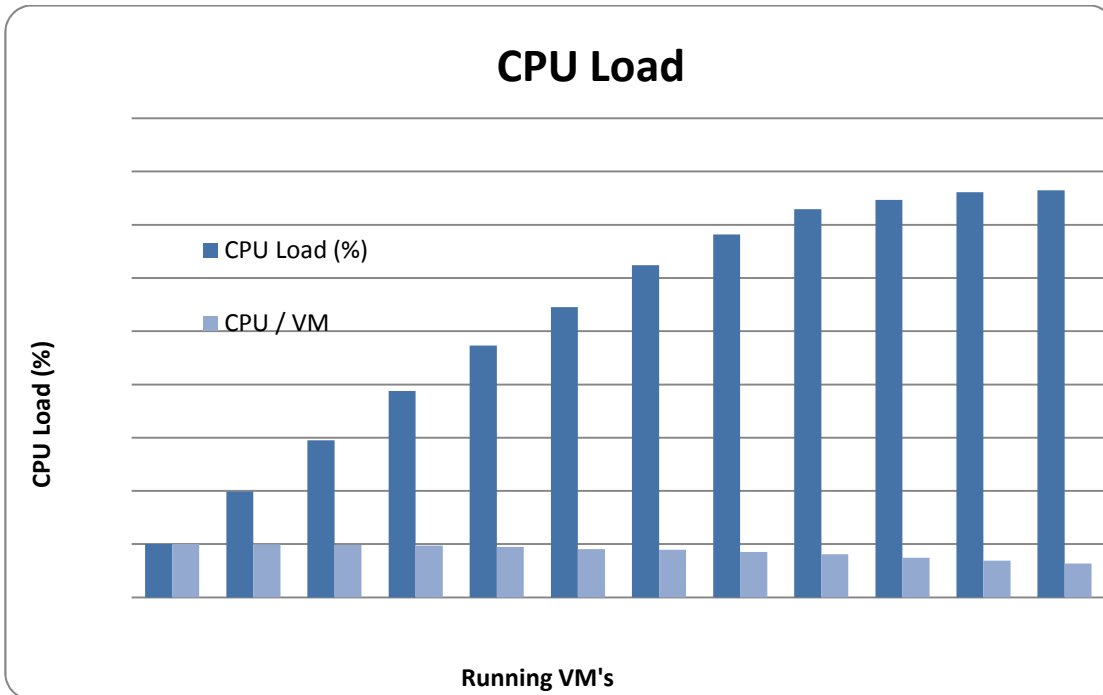


Figure 27: Dimensioning - CPU Load

We can quickly see how the VM's practically saturate the cores to 100% CPU load. The graph reflects two behaviors, which differ depending on the number of VM's in execution.:

$$\sum_i^n VM_i \leq cores$$

$$\sum_i^n VM_i > cores$$

At first, while the number of VM's is less than or equal to the number of cores in the host, every virtual CPU gets an entire physical core for execution. Now, once the number of virtual machines exceeds the number of cores, Virtualbox's VM balancing mechanisms kick in. At that point the CPU time each VM sees decreases proportionally to the number of VM's running. However, although execution times do increase for the factorial test program the QoE for the virtual machines was still quite acceptable even with 12 VM's saturating their "virtual" cores.

In an attempt to numerically evaluate the decrease in performance under heavy load (12 VM's running concurrently) we compared execution times for the CPU intensive factorial software at the host with respect to the VM's:



- Average execution time at host: 0.427s
- Average execution time at VM: 1.453s
  - ~3.4 factor performance degradation
  - In terms of performance that doesn't represent a very long leap to the past according to Moore's Law.

We would like to precise here the fact that Moore's Law refers to the volume of transistors per unit of area and not performance, so although the number of transistors arguably doubles every year this doesn't quite translate to performance. Now what to quantify how bad the performance drop is we can resort to Pollack's rule "*performance increases due to microarchitecture techniques are square root of of the number of transistors or the area of a processor*":

$$\Delta performance = \sqrt{\Delta transistors / area}$$

Therefore:

$$3.4^2 = \Delta transistors / area = 10.24$$

Since the transistor per unit of area doubles every year, we can quickly conclude the VM is performing like a machine from  $10.24/2 = 5.17$  years ago. Although that is obviously noticeable, it is definitely more than enough to perform quite well under typical workload conditions.

On the other hand, one of the first things we noticed when performing these tests, and it is the reason why we could not test more than 12 VM's concurrently, is the fact that the RAM is pre-allocated by VirtualBox. This implies that when the VM is started, it immediately uses up a large portion of its VM RAM (not quite 512 MB, but close) on the host. This helped us identify the first limiting, and critical factor in terms of infrastructure dimension: It is key to ensure the servers have enough RAM for the worst-case scenario: all users (or a statistical value with perhaps 99.99% of probability). Failure to provide enough RAM for running virtual machines will result in *memory thrashing*: at that stage performance decreases drastically.

## 5.5 Average Workload Test

Stressing the hosts CPU was obviously important to evaluate a worst-case scenario. But we wanted to also evaluate the server capacity to deal with average user workloads. We defined a user workload as the following:

- Email Client
- Web Browser (5 tabs, including flash site and a youtube video).
- Word Processor
- Presentation Software
- Audio Player
- Skype + IM client

Inspecting the CPU load while running such a workload yielded the following results:

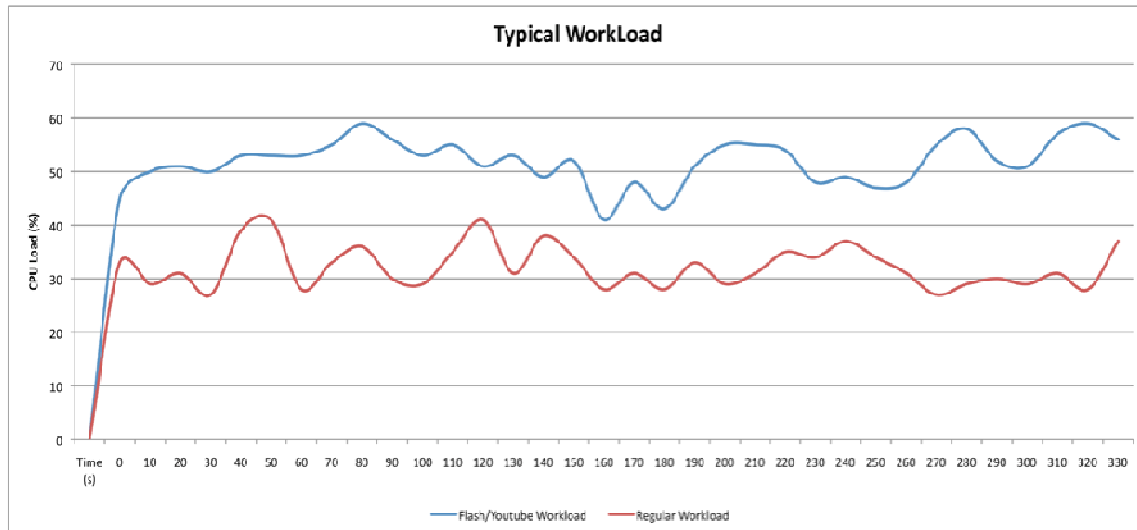


Figure 28: Typical Workload CPU load

The analysis of results, collected with the linux command *top* show the following for average loads:

- Youtube/Flash: ~51.9% cpu load.
- Regular Workload: ~32.2% cpu load.

These results suggest we can have the following number of VM's on an 8-core server before having **any** performance degradation whatsoever:

- Youtube/Flash workload: ~16 VM's
- Regular Workload: ~25 VM's

### 5.5.1 Scalability/Dimensioning Conclusions

With both the stress and typical workload tests, we were able to extract some conclusions. The stress test showed us how virtualbox was rather efficient at balancing VMs on the available cores. We can carry that knowledge and pair it with the typical workload results to try to make a rough estimate as to how many cores would be necessary. The fact that we were able to run 12 VM's on 4 real cores (with HT) without severe performance degradation implies saturating the physical cores is not that problematic. We also saw how we can run a typical workload 2 VM's per core (or 3 VM's per core if no flash is used) without even saturating the core. This combined suggests how we could consider running 3-5 VM's per core while still delivering an acceptable QoE.

The main conclusion extracted though was memory related. We need, without a shadow of a doubt, provide at least as much RAM as the sum of all executing VM's RAM and an additional 2GB RAM slack for the host. VirtualBox allows balancing and sharing CPU time with fine granularity. On the other hand, failure to provide enough RAM means we need to swap and memory thrashing has, there is no fine granularity in this case: you either have enough memory or your performance will be seriously hurt.

## **6 Conclusions**

This Deliverable described the first field trial of the MAINS project. This first set of testing validated UAM's Virtual PC Application Service over Primetel's optical network and using Intune's NXT Beta nodes.

This first field trial has achieved the validation and operability of the Virtual PC Application and MAINS' optimized transport protocol. This first set of tests was able to illustrate the performance of these two components in a real operator's network.

Performance Evaluation results were obtained both using several metrics like throughput support and CPU load but also with real users and Quality of Experience. Due to the ideal topology setup, the use of only optical links on the network transport and the not so large topological distances of the overall ring network (~16km) the Quality of Experience in using a Virtual Machine from a local server versus using a Virtual Machine from a remote server is unnoticeable to the user's eyes.

Nevertheless noticeable effects do occur when a user transfers his/her virtual machine from one server to another due to e.g. mobility or change of location and has to await for the time it takes to copy or duplicate the machine between the two servers involved in such a scenario. This time waiting however is mainly due to server limitations, hard disk drive writing speeds and TCP protocol constraints and not the underlying optical technologies proposed within MAINS.

A second round of testing will take place at the University of Essex Laboratory which will include the configuration of an integrated Optical Ring to Optical Mesh Subwavelength technologies, validation of the common GMPLS Control Plane and associated interoperability evaluation.

## 7 References

- [1] Juan Fernandez Palacios et. al. Metro Architectures Enabling Sub-Wavelengths: Rationale and Technical challenges, ICT Mobile Summit, Florence, Italy, June 2010.
- [2] EU FP7 ICT ETICS project; <https://www.ict-etics.eu> (Last Access Date: January 2012)
- [3] EU FP7 ICT STRONGEST project; <http://www.ict-strongest.eu> (Last Access Date: January 2012)
- [4] EU FP7 ICT GEYSERS project <http://www.geysers.eu> (Last Access Date: January 2012)
- [5] EU FP7 ICT BONE project [http://www.ict-bone.eu/portal/landing\\_pages/index.htm](http://www.ict-bone.eu/portal/landing_pages/index.htm)
- [6] J. Triay, G.S. Zervas, C. Cervello-Pastor, D. Simeonidou, "GMPLS/PCE/OBST Architectures for Guaranteed Sub-Wavelength Mesh Metro Network Services", to appear at Optical Fiber Communication Conference, Los Angeles, USA, March 2011
- [7] Michael Georgiades, Jaime Fullaondo, George Zervas, Juan P. Fernandez-Palacios, Mark Basham, A suite of Optical Network Testbeds for a Realistic Evaluation of MAINS' Sub-Wavelength Switched Metro Networks Architectures, ICT 2011 18th International Conference on Telecommunications Ayia Napa, Cyprus May 2011
- [8] Oracle/Sun. Deployment Guide <http://wikis.sun.com/display/VDI3dot2/Deployment+Guide>