Escola d'Enginyeria de Telecomunicació i Aeroespacial de Castelldefels

UNIVERSITAT POLITÈCNICA DE CATALUNYA

MASTER THESIS

TITLE: Contributions to a Green IT project: definitions of use cases and first steps towards a power model for routers

MASTER DEGREE: Master in Science in Telecommunication Engineering & Management

AUTHOR: Ismael Mateos Hernández

DIRECTOR: David Rincón Rivera

DATE: September 16th 2011

Title: Contributions to a Green IT project: definitions of use cases and first steps towards a power model for routers

Author: Ismael Mateos Hernández

Director: David Rincón Rivera

Date: September 16th 2011

Overview

One of the current research hot topics on IP networks is how to minimize energy-related communications. Currently, network technologies are responsible for 2% of global emissions of CO_2 and its reduction would mean a significant improvement of environmental conditions and a reduction in the rate of global warming.

There are several efforts related to Green IT such as the migration and consolidation of virtual machines, the impact of virtualization and cloud computing, and energy efficiency metrics. The proposed master thesis focuses on an exploration, in the sense of reviewing the state of the art (what other authors have done), to identify areas where contributions can be made, and suggest some sort of improvement.

In this context, the i2CAT Foundation is taking part in the Canadian project GSN, an innovative project focusing on the relationship between networks and green datacenters, powered by renewable energies and following the “follow the sun/follow the wind approach”, in order to migrate Green ICT services and therefore reducing the carbon footprint. This first part of this thesis describes some contributions to the GSN project, such as helping in the installation of a solar-powered node, and the definition of the GSN Use Case. The second part focuses on a specific task not included in GSN but of vital importance and of great interest: the development of a power model for routers, specifically for virtual routers as those that can be migrated in the GSN Project.

INDEX

INTRODUCTION

One of the current issues in research on IP networks is to minimize energy consumption related to communications. Currently, network technologies are responsible for 2% of global emissions of CO_2 [1], and its reduction would mean a significant improvement of environmental conditions and a reduction of the pace of global warming.

Power management for networks is nowadays considered a very important perspective and has begun to receive attention by researchers. The minimization of energy consumption of operations and environmental reasons is now a hot topic in the research community. See [2], where a series of practical solutions to reduce GHG (Greenhouse Gases) emissions are presented by Bill St. Arnaud or [3], where the advantages that Green IT can play in reducing CO_2 emissions are shown, for example.

Public concern about power consumption and carbon footprints is rising, and stands to affect network equipment via standards such as Energy Star [4], which was planned to promote and recognize energy efficiency in monitors, air conditioning equipment and other technologies. In fact, Energy Star standard proposals for 2009 discuss slower operation of network links to conserve energy when idle.

In this context, the i2CAT Foundation[1] is taking part in the Canadian project GreenStar Network (GSN) [5], whose objective is to initiate a Canadian consortium of industry, universities and government agencies with the common goal of reducing GHG emissions arising from information and communication technology (ICT) services. The project is innovative because it focuses on the relationship between networks and green datacenters, powered by renewable energies, in order to provide Green ICT services and therefore reducing the carbon footprint.

This thesis describes some contributions to the GSN project. The first part includes some tasks carried out in the project, such as collaborations in the installation of a solar-powered node, training on ISO 14064 [6] (studying and transferring to other team members), review of a series of documents related to the project, and collaboration in the definition of the GSN Use Case.

A specific task not included in GSN but of vital importance and of great interest for research is the development of a power model for routers, and specifically

[1] The i2CAT Foundation is a center of research and innovation, which focuses its activities on the development of the future Internet.

for virtual routers such as those that can be migrated in the GSN Project. That is why we undertook a measurement campaign that included a high-end router, Juniper MX480, and other low-end equipments such as Juniper J2320 and J2350. We did a set of current and power measurements in order to study the energy consumption of these devices.

This idea is given by the generic power model for router power consumption created by Chabarek and Barford in 2008 [7], where they applied this model in a set of target network configurations and used mixed integer optimization techniques to investigate power consumption, performance and robustness in static network design and in dynamic routing. Their results indicated the potential for significant power savings in operational networks by including power-awareness.

To carry out everything previously mentioned, the master thesis has been organized in the following sections:

Chapter 1 presents an overview of what is Green IT and explains a number of approaches that are currently underway.

Chapter 2 focuses on the GreenStar Network Project (GSN) as a Green IT alternative to promote the use of renewable energy sources, and thereby reduce emissions of CO_2. This chapter also summarizes the contributions of the thesis to the project.

Chapter 3 presents a set of current and power measures on the Juniper routers in order to create a generic model for virtual router power consumption.

Finally, we summarize the conclusions of this work and the problems found. It also shows a global vision of the future work that can be done in this research area.

Before going farther, we would like to acknowledge the contributions of other people involved in the project.

I would like to thank Sergio Jiménez (another student of EETAC working in his degree thesis) for helping me with the router configuration, really new for me, as well as to the i2CAT Foundation for giving me the opportunity to participate in the GSN Project, and specifically Pau Minoves and David González for giving me access to the MX480 router, and Cristina Cervelló for the Juniper J-series router.

CHAPTER 1. GREEN IT

This chapter offers a brief introduction in order to emphasize the importance of the emerging concept of Green IT taking into consideration worldwide power requirements to support IT infrastructures. In addition to this, an overview of activities and different technologies related to Green IT, which are aimed to achieve a better IT performance in terms of energy efficiency, is provided too.

1.1 Introduction

We live an age in which matters as serious as the shortage of energy, the global warming, or the greenhouse effect, have made environmental care a priority both for governments and companies, and the society as a whole. Information technologies (IT) cannot remain alien to this challenge, and all the agents involved in their development, implantation or utilization have started to consider all the possible efforts to mitigate the impact of IT in the environment.

This one is the subject on which Green IT focuses, and has been identified by Gartner as one of the ten strategic technologies for 2009 [8]. The term “green computing” was introduced for the first time after the beginning of the program Energy Star in 1992 (see section 1.2), promoted by the American government. It aimed to label electronic equipment (namely, computer screens) with improved energy efficiency. Nowadays the program Energy Star is the engine of the energy efficiency in the electronic systems (not only of information processing equipment, but also of the electronic domestic equipment).

Today, energy consumption is a critical question for the IT organizations, either to reduce costs, to preserve the environment or to support the data centers. Only in the United States, the data centers consumed 4,500 million dollars of electricity in 2006 [9]. Gartner estimated that during the next five years (2006-2011 period), the majority of centers of information of companies will spend as much money in energy (power and refrigeration) as in infrastructure hardware [9].

Figure 1.1 shows a comparative of the energy expense of data centers in United States from 1996 up to a projection for 2011:

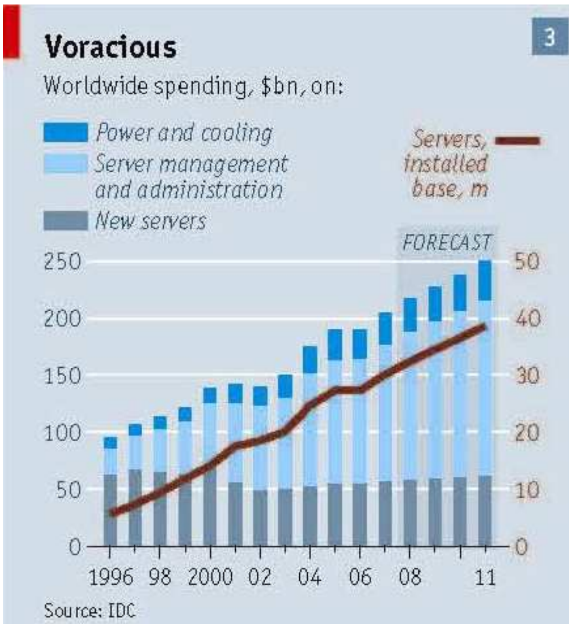

Fig 1.1 Energy expense estimation from 1996 until 2011 in United States [10]

Note how the number of servers in data centers increases at the same rate the energetic expense for the refrigeration and the management of these servers also grows. In addition to this, it is expected a linear growth of 12.5 per year.

Approximately, 40-50 % of the energy consumption of the companies goes to IT, and the energy cost of data centers has more than doubled in the last five years. Figure 1.2 shows the increase of energy consumption per rack (designed to house electronic equipment) in data centers.

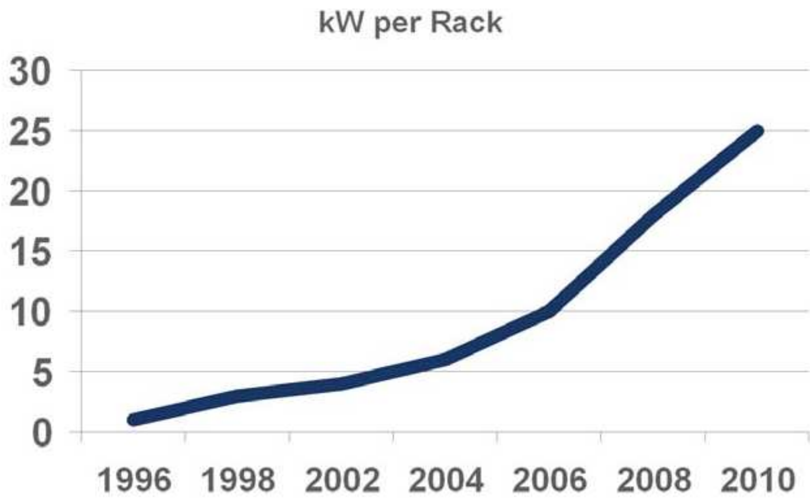

Fig 1.2 Increase of energy consumption in data centers [9]

As a solution to the aforementioned challenges, the adoption of more efficient products and approaches can enable more equipment within the same energy consumption (so-called energy footprint). Organizations should consider the established regulations and promote alternative plans for the growth of their data processing centers and capacity [11].

1.2 Energy Star

In 1992, the Environmental Protection Agency launched the Energy Star program, which was planned to promote and recognize energy efficiency in monitors, air conditioning equipment and other technologies [4].

The program was soon widely accepted, becoming a fact the presence of a sleep mode in consumer electronics. Nowadays, almost all the major suppliers of electronic equipment have joined the program, wearing the Energy Star logo (see Figure 1.3) in their systems.

Fig 1.3 Energy Star logo [4]

The program results have been very promising since the 2007 annual report estimated that United States energy savings amounted to 16 billion U.S. dollars, and prevented the emission of 40 million metric tons of greenhouse gases [4].

The Energy Star program plays a role as a source of credible and objective information so that consumers and businesses can make decisions based on useful information for themselves and the environment.

1.3 Renewable energy sources

A new focus of interest related to green technologies is the capture or collection of energy and its corresponding use, boosted mainly from renewable energy sources.

It refers to the process of extracting useful electrical energy from environmental sources (sun, heat, wind, kinetics, etc) using special equipment, called transducers, which have the ability to convert one form of energy into another.

There are three main renewable energy sources:

- Photovoltaics: Obtains energy from light, both external (solar) and indoor. It is by far the most advanced renewable energy source, and is present in many electronic devices.
- Thermal: Heat from the environment.
- Mechanics: Includes the wind, kinetics, vibrations, etc.

Figure 1.4 shows an expected impact on the daily consumption of these energy sources:

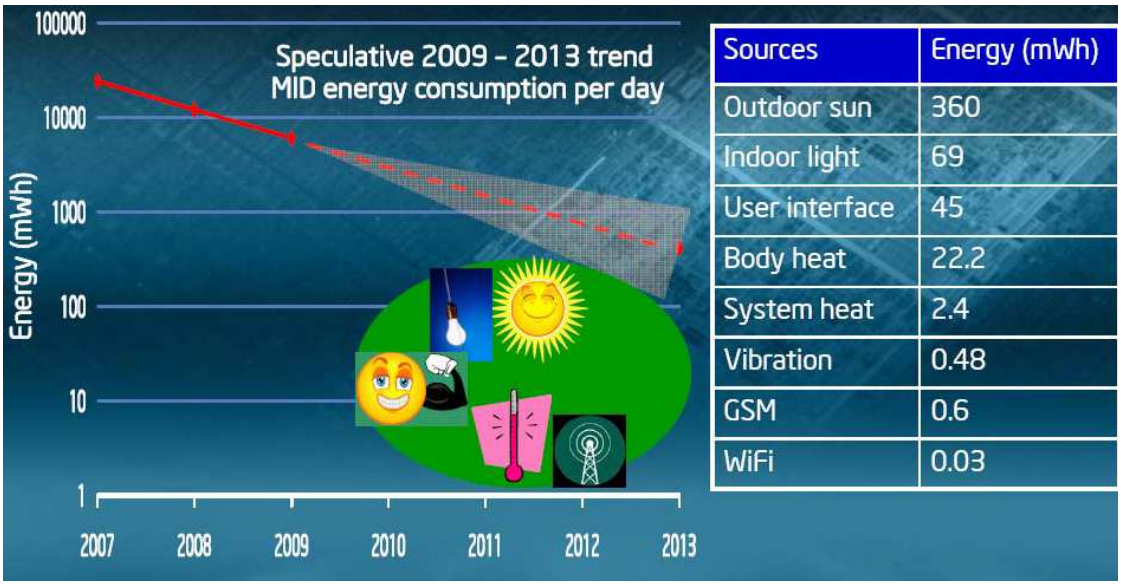

Sources	Energy (mWh)
Outdoor sun	360
Indoor light	69
User interface	45
Body heat	22.2
System heat	2.4
Vibration	0.48
GSM	0.6
WiFi	0.03

Fig 1.4 Energy capture from renewable sources [12]

The potential of these technologies is so great that provides a break in the current system of power generation based on fossil fuels. An example of this is the GSN project (GreenStar Network), described in Chapter 2 of this thesis.

1.4 Virtualization

Virtualization was also identified by Gartner as one of ten strategic technologies for 2009 [8]. It is a broad term that refers to the abstraction of the computer resources.

The common theme of all virtualization technologies is to implement an external interface by combining several physical resources to be used as if they were one, simplifying and hiding the control [13].

The energy demands of data centers can be reduced by proper sizing of the IT infrastructure through consolidation and dynamic management of computing power on a farm [14]. Server virtualization is the easiest way to do this, which according to VMware [15] allows:

- Consolidating servers: Consists on running several virtual machines in a physical server. It can reduce the number of servers in a ratio of up to 15:1, eliminating server sprawl and cutting costs (see Figure 1.5).

- Reduce energy consumption [16]: Each virtualized server saves 7,000 kW/h of electricity annually, or about 500 Euros in energy costs.
- Increasing the capacity of IT: The ratios of server utilization are improved from 5-15% to 60-80%. With fewer servers, but more used, it frees up space and power.
- Reduce emissions of CO_2: 4 tons of CO_2 are removed for every server virtualized, equivalent to taking 1.5 cars off the road.

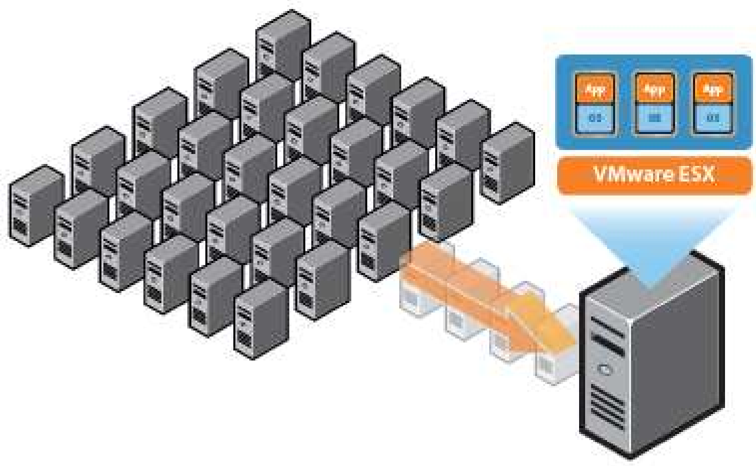

Fig 1.5 Server consolidation [15]

Virtualization is already having a positive impact on the environment. Gartner [17] estimated in 2007 that over one million workloads run in VMware virtual machines, representing an aggregate energy savings of around 8,500 million kW/h.

In addition to this, OpenNebula is the open-source alternative to the creation of a virtual infrastructure that enables dynamic deployment and relocation of virtual machines on a set of physical resources [18].

OpenNebula extends the benefits of virtualization platforms from a single physical resource to a set of them, decoupling the server not only from the physical infrastructure, but also from its physical location.

1.5 Cloud model

Cloud Computing, also identified by Gartner as one of ten strategic technologies for 2009, is a computing model in which vendors provide a variety of capacities, through IT, to consumers. The main features of Cloud Computing are:

- Providing capabilities as a service.
- Service provision in a highly scalable and elastic way.
- Use of Internet technologies to develop and deliver services.
- Design to cater external clients.

There are many providers who self-referred to as Cloud. The Cloud Computing Directory [19] classifies the existing suppliers in terms of what they offer:

- Software as a Service (SaaS): Software delivery model where a company performs the maintenance, support and operation that the client will use during the time that the service is contracted. The customer will use the system hosted by this company, which will keep the customer information in their systems and provide the necessary resources in order to exploit this information. Some examples are Google Apps [20], Oracle on Demand, and QAD on Demand.

- Platform as a Service (PaaS): This model provides everything needed to support the entire life cycle of construction and implementation of web applications and services fully available. Another important feature is that no downloading of software is installed on the computers of developers. PaaS offers multiple services, but all provisioned as a comprehensive solution on the web. Some examples are Google App Engine [21], Sun Project Caroline [22], and GigaSpaces [23].

- Infrastructure as a Service (IaaS): Distribution model of computing infrastructure as a service, usually through a virtualization platform (see Chapter 2). Instead of purchasing servers, space in a data center or network equipment, customers buy all these resources to a third party service provider. A fundamental difference with the virtual hosting is that the supply of these services is an integral way through the web. Some examples are Mantychore [24], Amazon Web Services [25], which includes EC2 (Elastic Computing Cloud), GoGrid [26], and Microsoft Azure [27].

Also there are already tools to create a cloud infrastructure, either within the organization or to serve other organizations:

- VMware vCloud.
- Sun xVM Ops Center.
- Citrix Cloud Center.
- Eucalyptus.
- OpenNebula.

As we mentioned in the previous section, OpenNebula provides interfaces for remote access to a virtual infrastructure using different virtualization technologies (Xen, KVM, etc) and enables dynamic provisioning of resources from different vendors such as Amazon EC2, as shown in Figure 1.6:

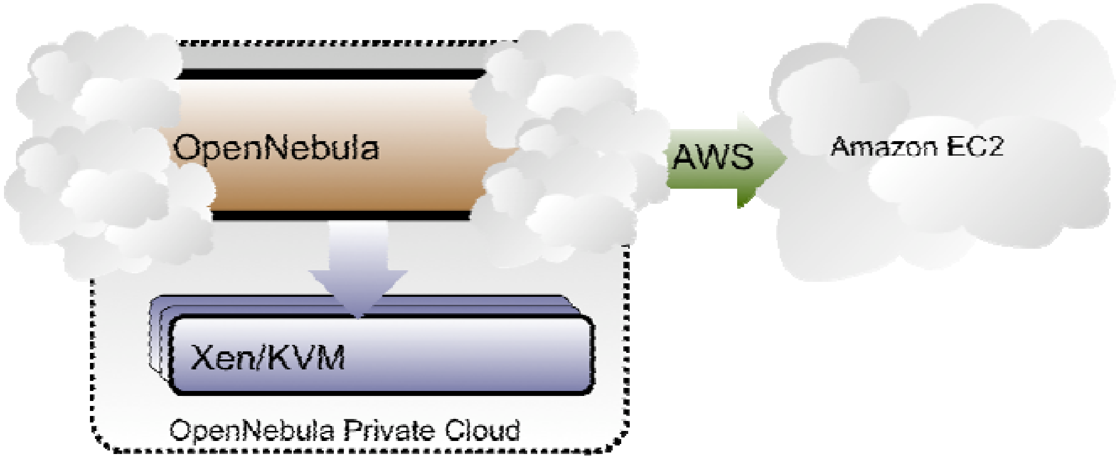

Fig 1.6 Dynamic provisioning of EC2 [18]

1.6 Energy efficiency metrics

When improving the energy efficiency of a component, system or installation is essential to understand how they really behave, that is to say, how they consume energy, dissipate heat, etc. It should be an assessment of the current situation, so you can set up monitoring of the improvements or a comparison with other alternatives.

An example of this is trying to accomplish in Chapter 3 of this thesis, where we describe the details a set of current and power measurements on a router, in this case the Juniper MX480, in order to create a generic model for virtual router power consumption.

On the other hand, it should be noted that the Green Grid consortium [28] has a standard set of metrics [29] in order to evaluate the efficiency of data centers. These metrics are: PUE (Power Usage Effectiveness) and DCiE (Data Center infrastructure Efficiency). They enable data center operators to quickly estimate the energy efficiency of their facilities, compare their results, and determine whether to make improvements in energy efficiency.

The PUE metric is defined as:

$$\text{PUE} = \text{Total Facility Power} / \text{IT Equipment Power} \qquad \textbf{(1.1)}$$

And DCiE metric, which is its reciprocal, is defined as:

$$\text{DCiE} = \text{IT Equipment Power} / \text{Total Facility Power} = 1 / \text{PUE} \qquad \textbf{(1.2)}$$

In Equations 1.1 and 1.2, IT Equipment Power includes the burden associated with all IT equipment, that is to say, servers, storage and networking, along with other additional equipment such as KVM (Keyboard-Video-Mouse), monitors, and fixed and portable equipment to monitor or control the data center.

In addition to this, Total Facility Power includes IT Equipment Power and everything you need to bear the burden:

- Power Distribution Components: UPS (Uninterruptible Power Systems), switches, generators, PDU (Power Distribution Unit), and batteries.
- Cooling system components: Refrigerators, air conditioning units for the computer rooms (Computer Room Air Conditioning, CRAC), direct expansion units, pumps, and cooling towers.
- Other component loads such as lighting.

The PUE metric is the original, however DCiE is usually being adopted, expressed as a percentage, which facilitates its understanding. Improvements in energy efficiency result in a close to 100% of DCiE, which is the ideal number.

CHAPTER 2. GREENSTAR NETWORK (GSN)

This chapter focuses on the GreenStar Network Project [30] as a Green IT alternative with the aim of reducing CO_2 emissions due to high consumption of electricity in datacenters. The project is innovative because it focuses on the relationship between networks and green datacenters (see section 1.3) in order to provide Green ICT (Information and Communication Technology) services.

2.1 Introduction

The goal of the GreenStar Network Project is to initiate a Canadian consortium of industry, universities and government agencies with the common goal of reducing greenhouse gas (GHG) emissions arising from ICT services.

ICT services are increasingly being provided by datacenters as a service, including innovative services such as cloud, grid, and utility computing, storage, and networking.

The urgent need to reduce greenhouse gas (GHG) emissions affects both the providers and customers of ICT services. Carbon emissions resulting from service implementation affect service providers' costs, as well as legislative contexts such as carbon taxes [31].

On the other hand, customers balance the environmental impact against price and service quality. For all that the GSN Project's aim is to develop a practical carbon footprint exchange standard for ICT services, and to create a testbed network (the GreenStar Network) for ICT service providers [31].

The GreenStar Network is applied to two Green ICT service provision scenarios:

- Sale of carbon credits based on the utilization of a green datacenter.
- Development of management and technical policies that leverage virtualization mobility to facilitate use of renewable energy within the GreenStar Network.

The expected result is the creation of tools, protocols, procedures, and use cases for a growing network of ICT service providers that offers customers the lowest price and greenest services.

2.2 Architecture

Services need to be able to move where computing or networking resources are available. This may be within a single domain of operation of an organization or across multiple domains. In order to achieve mobility, one of the key factors is the unified computing, which is the principle driving state-of-the-art next generation datacenters and networks [32].

All resources are linked together in a common architecture that can be virtualized. In other words, the compute and storage platform is architecturally unified with the network and the virtualization platform, making no distinction between the network and the edge devices connected to it [32].

It should be noted that this unification is realized through the use of middleware linking the different systems together.

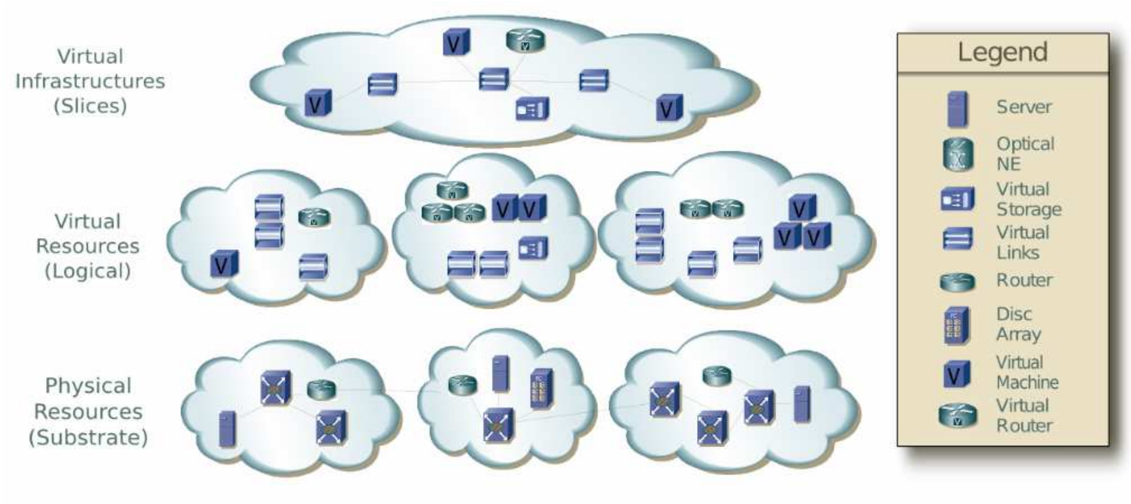

Fig 2.1 Unified Computing and Virtual Infrastructures [32]

Figure 2.1 shows how virtual infrastructures are created. They are assembled by creating "slices" of the physical substrates, which are then aggregated into a working virtual infrastructure from which services can be delivered.

The physical infrastructure consists mainly of servers, disc arrays and network elements such as switches and routers, whereas the virtual infrastructure consists of virtual machines, virtual storage, virtual routers and virtual switches.

In this project, the optical network elements, servers, routers and disc arrays are unified through the use of various middleware solutions and tools. In order to achieve the required level of mobility, software instances of are used [32].

2.2.1 Tools

To establish this unified computing environment at each GSN node, the tools listed below are used [32]. The IaaS framework [33], an open source framework for resource description and exchange, is used to wrap the different cloud solutions, tools and service offerings at the nodes.

- XORP [34] and Quagga [35] support software based on routers and they are running on dedicated virtual machines.
- Xen [36] and KVM [37] are the supported hypervisors, since most cloud computing middleware is based upon these. Puppet [38] is used for contextualization.
- VDE [39] and Linux Kernel are used for Ethernet Management.

- iSCSI and LUN, as well as network file systems (NFS, GFS, etc) are used too.
- Linux is the operating system used by hosts and guests.

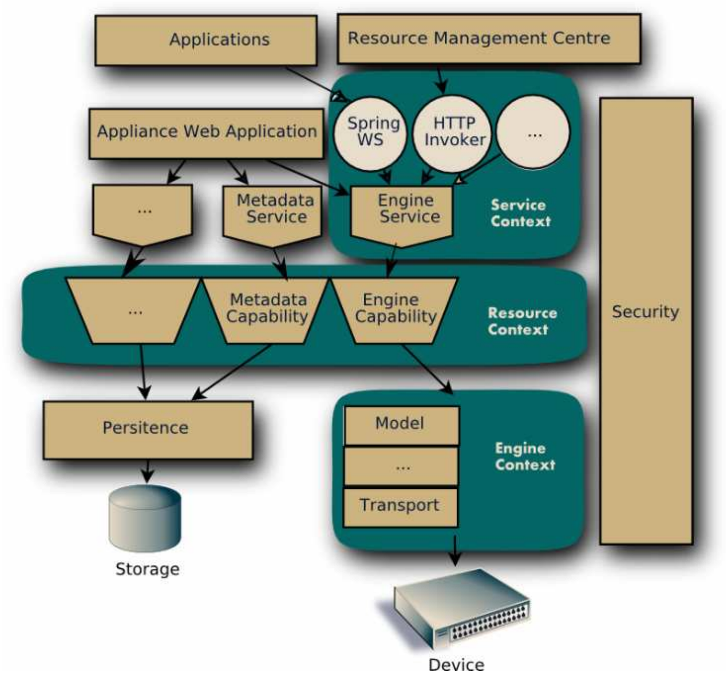

Fig 2.2 IaaS Framework Architecture [33]

The IaaS Framework is used as the integration middleware. Its architecture depicted in Figure 2.2 makes it suitable to represent any device or service as resources.

Moreover, while these tools offer the core technologies on which the GSN fabric is built, management middleware is needed. The following middleware solutions are used [32]:

- OpenNebula, Eucalyptus and Incus manage servers and virtual machines.
- Ether control VDE, Linux Kernel and the switches.
- Mantychore and Pike control routers and virtual routers.
- Walrus controls the storage system.

Note that Incus, Ether and Pike are commercial solutions and the rest are open source.

These diverse middleware solutions are integrated by translating the resources models in order to conform to IaaS Framework resources.

2.3 GSN Carbon Measurement Protocol

Although the goal of this project is to reduce carbon emissions caused by ICT, the key technical task is to measure such emissions. Reduction, with or without measurement, is certainly a positive development; however, quantification is required before performing any action, and is a hot topic in the emerging legislative and economic context. Furthermore, the measurement process must meet identified levels in order to be considered acceptable [40].

When quantifying carbon reduction for the purposes of selling this reduction to an organization needing to offset their own carbon emissions, the ISO 14064 [6] standard is the basis for guaranteeing that these reductions meet the aforementioned goals [40].

As far as we know, ISO 14064 has never been applied to measurement of carbon emissions related to ICT, although there is nothing inherent in the standard that prevents this. However, the standard is written at a very general level, and identifying the particular measurements appropriate to ICT is a central activity of the GSN Project. The resulting document is the GSN Carbon Measurement Protocol [40].

The protocol is more specific than ISO 14064, but more general than just the planned network implementations within the GSN project. It describes *what* has to be measured, but not *how* it must be measured, and indicates acceptable estimation methods for certain measurements [40].

Figure 2.3 shows the process to develop and promote a GHG quantification protocol related to greening information technologies in conformance to ISO 14064.

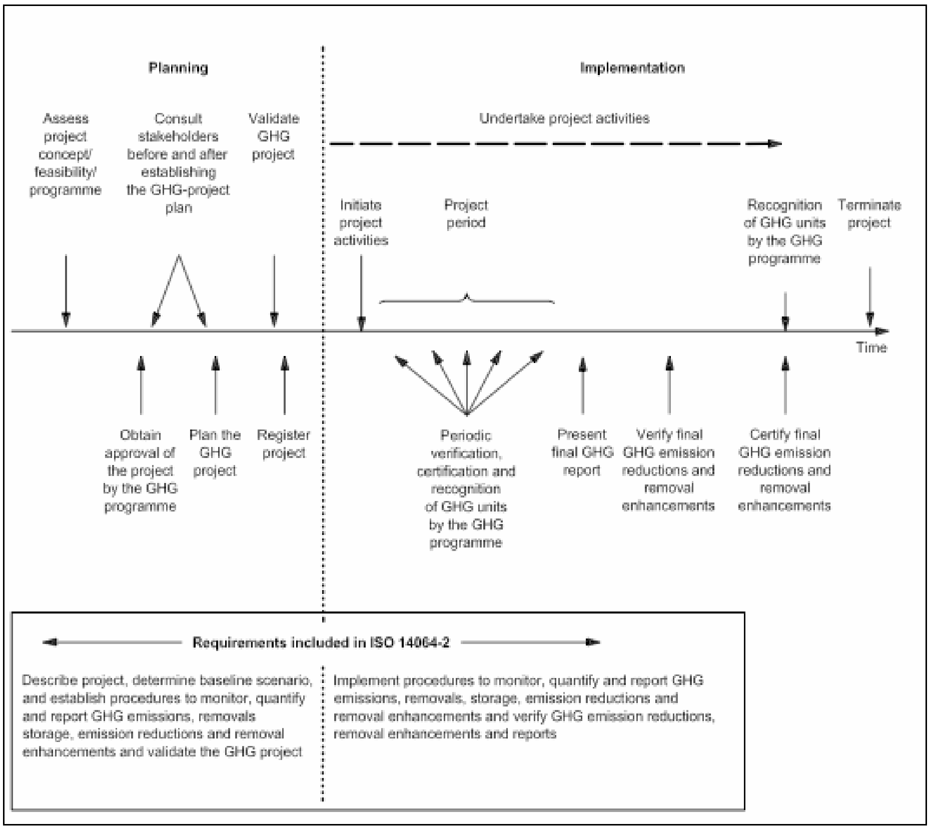

Fig 2.3 A typical GHG project cycle [40]

For the planning phase, ISO 14064-2 specifies requirements for establishing and documenting a GHG project. These are summarized as follows [41]:

- Description of the project.
- Identification of GHG sources, sinks and reservoirs relevant for the project.
- A baseline scenario determination.
- Procedures developments to quantify, monitor and report GHG emissions, removals, emission reductions and removal enhancements.

For the implementation phase, ISO 14064-2 specifies requirements for the selection and application of criteria and procedures for regular data quality management, monitoring, quantification and reporting of GHG emissions, removals, emission reductions and removal enhancements [41].

2.4 Technology

This section is a summary of [42].

GSN utilizes state of the art cloud computing and networking technology, thereby providing virtual machine mobility without interruption of user service, as explained in section 2.1. Moreover, because this is being done in a federated environment, the cloud resources form a single virtual infrastructure even when supplied by multiple infrastructure providers.

GSN's virtual infrastructure is structured into a hub and spokes topology, thereby inspiring the name Green Star.

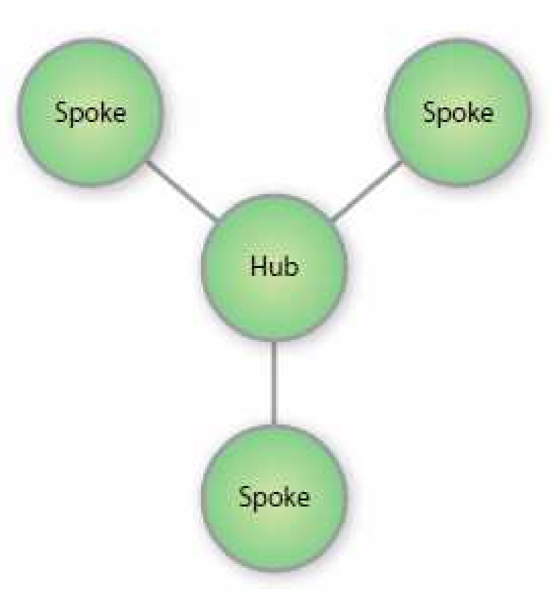

Fig 2.5 Hub and spoke topology [42]

As seen in Figure 2.5, a hub is a site where energy is always available (for example a physical router running several routers) including hydro power or windmill farms of sufficient size such that there are no chances of having downtime.

While there can be power management techniques at the hub that allows the manager to shutdown unused infrastructure, these nodes are expected to be reachable at all time. The spokes (physical routers running virtual routers), however, are expected to have limited availability because they are powered by renewable energy sources; they may have power grid backups, but are not expected to be always available for computing.

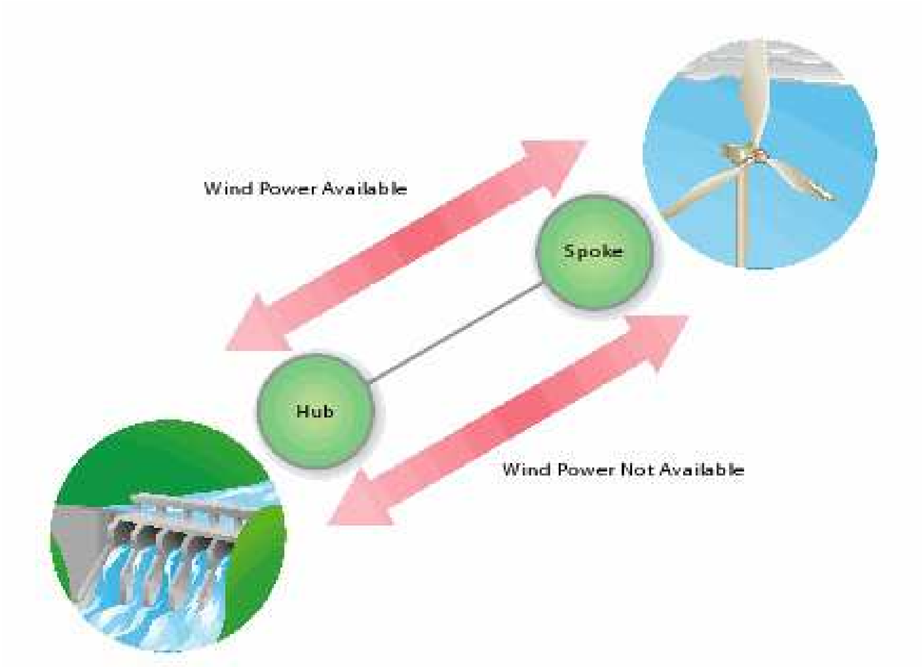

Fig 2.6 Migration of virtualized ICT services between hub and spokes [42]

As seen in Figure 2.6, a hub is connected to a spoke to which it can move jobs that need to be processed. The resources at the end of each spoke are powered by renewable energy; jobs are moved there only when the renewable energy is available. In addition to this, it is important to have a good model of power consumption of virtual machines and routers (see Chapter 3) in order to make these decisions of migration (in a way that is transparent to their users).

On the other hand, we have the Cloud Controller, which is a virtual infrastructure manager comprised of a simple optimization algorithm which tries to make the best use of the available compute spokes.

As seen in Figure 2.7, the storage is centralized at the hub to make sure data is always available. The spokes need to be connected via high speed optical networks; processing is dispatched by the centralized controller for this virtual infrastructure.

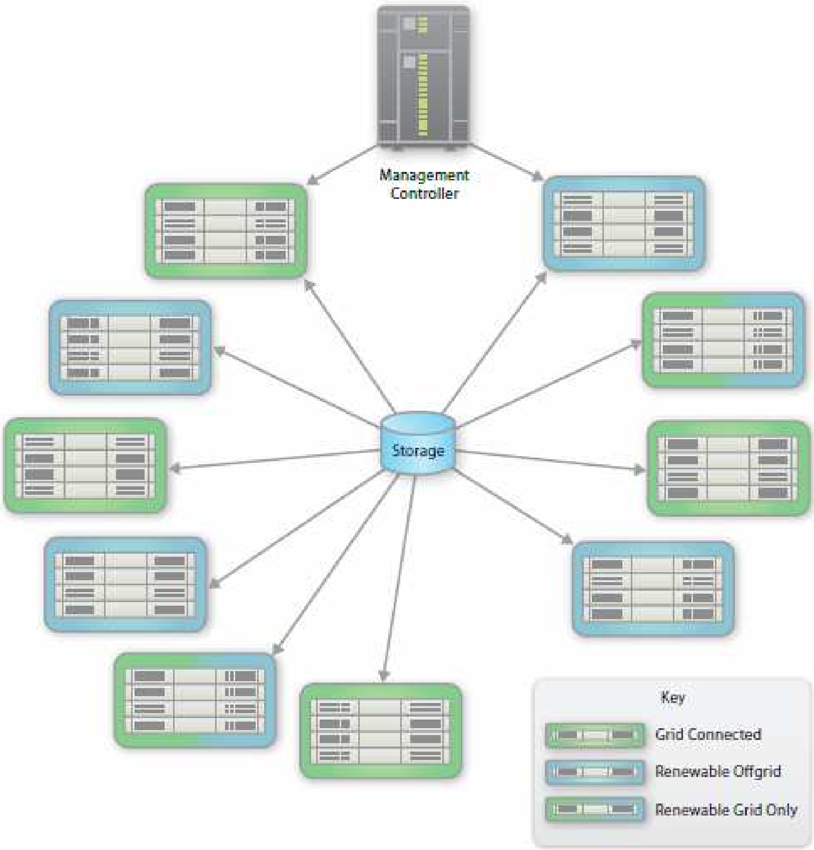

Fig 2.7 Management controller for renewable energy optimization [42]

2.5 Contributions

The main contribution of this thesis in the GSN Project was the installation of a solar-powered node in Riba Roja, a village near the province of Tarragona. As of September 1st 2011 the process is still underway because the optical link between Canada and the i2CAT node has not yet been established. For more information about the node description and its requirements see Annex D.

Other tasks performed are:

- Training on ISO 14064, which includes studying and transferring (to other team members) the knowledge included in:
    - ISO 14064-1: Specification with guidance at the organization level for quantification and reporting of greenhouse gas emissions and removals [43].
    - ISO 14064-2: Specification with guidance at the project level for quantification, monitoring and reporting of greenhouse gas emission reductions or removal enhancements [41].
    - ISO 14064-3: Specification with guidance for the validation and verification of greenhouse gas assertions [44].

- Review of document such as the Carbon Protocol Draft, collaboration on the writing of an abstract for the EUNIS [45] conference, and the GSN-Mantychore integration proposal.

- Presentation on the collaboration between Mantychore and GSN.

- Collaboration in the definition of the GSN Use Case, as explained below.

2.6 Summary of the GSN Use Case

This Use Case provides GSN with the possibility to provision connections between nodes in different domains and to migrate virtual machines.

Note that a node consists of one or more (physical) servers, one or more power distribution units (PDUs) and possibly a number of other devices with may or may not have network connectivity.

In addition to this, noteworthy that Mantychore [24] follows the Infrastructure as a Service (IaaS) paradigm, in the GSN Project, to enable National Research and Education Networks (NRENs) and other e-Infrastructure providers to enhance their service portfolio by building and deploying software and tools to provide infrastructure resources like routers, switches, optical devices, and IP networks as a service to virtual research communities.

This service follows the IaaS paradigm, consisting on offering remote access and control of infrastructure elements to third party organisations through software web services. By using IaaS services these organisations can control the remote infrastructure as if they owned it and be billed either per use or based on a monthly fee, promoting the reuse of existing infrastructure and avoiding the purchase of new devices on the provider and customer sites. IaaS also provides a new level of flexibility to the e-Infrastructure operators: their infrastructure can scale up or down following its customers' needs, therefore minimizing the cost of operating the infrastructure (both the capital and the operational expenditures).

2.6.1 Introduction

Recall that the GSN project focuses on live migration of virtual machines between a number of geographically separated physical datacenters, based on the availability of "green" energy at these sites. GSN employs a hub and spokes topology (see Figure 2.5), with a hub for control and central storage and a number of spokes for "green" processing.

The required Gigabit Ethernet (GE) circuits between the spokes and the hub are currently provisioned manually through third party carriers. Due to a limitation in the virtualization software, live migration of virtual machines can only be performed if these virtual machines and physical datacenters are all in the same subnet.

2.6.2 General goal

There are two main goals:

1. To provision permanent end-to-end circuits between spokes and the hub.
   - The end points are in different administrative domains[2].
   - The end points are switches or routers located at the datacenters hosting the nodes; local access is the responsibility of the network administrators, site administrators etc.
   - The underlying infrastructure should be transparent to the users; it should logically be a point-to-point connection.
   - IP connectivity should be tunneled between circuit end points.
     a. All the nodes have to be in the same subnet.
     b. A virtual router acts as end point / gateway for each physical node and any virtual machines residing on it (see Figure 2.8).

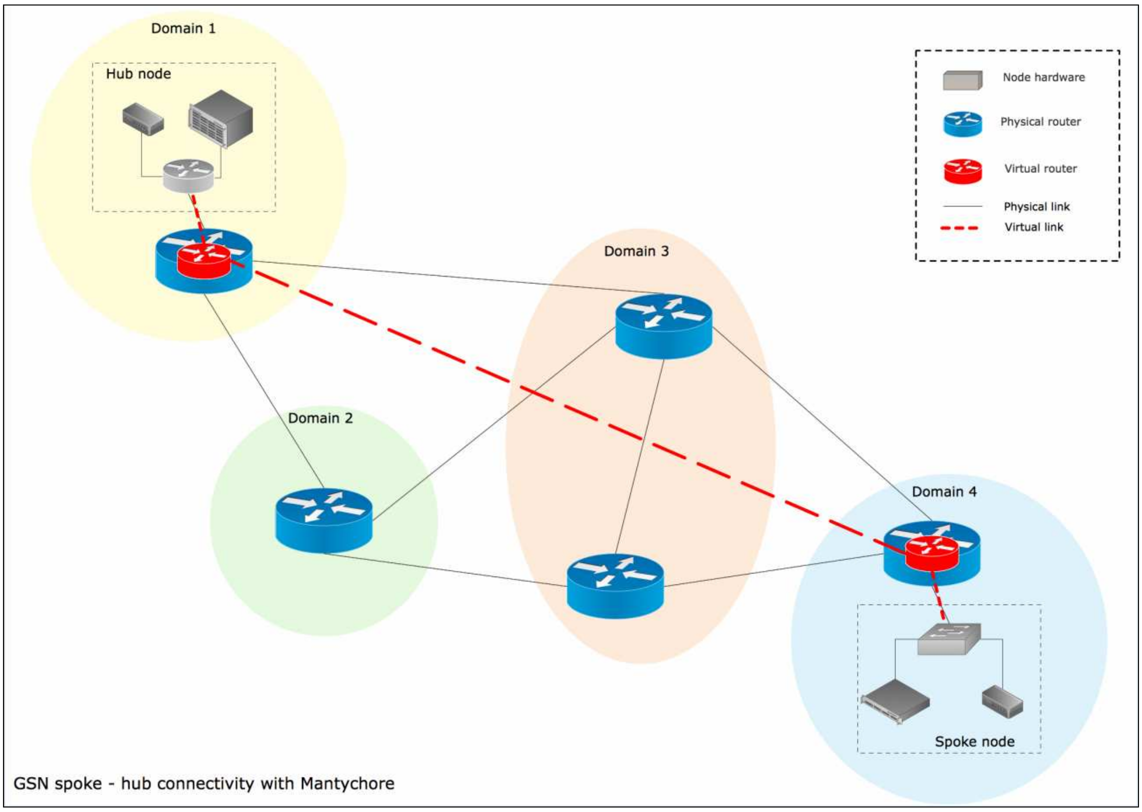

Figure 2.8 GSN spoke – hub connectivity with Mantychore [46]

[2] An administrative domain is a set of one or more nodes whose administrative servers share an administrative database.

2. To provision temporary circuits between nodes for migrations of virtual machines (see Figure 2.9).

   - The same requirements as for the permanent circuits.
   - The circuits are set up between two nodes and torn down after the migration is complete.
   - The virtual machines have to be in the same subnet and keep their IP addresses unchanged.
   - Enough bandwidth has to be available between the end points when a virtual machine is being migrated.

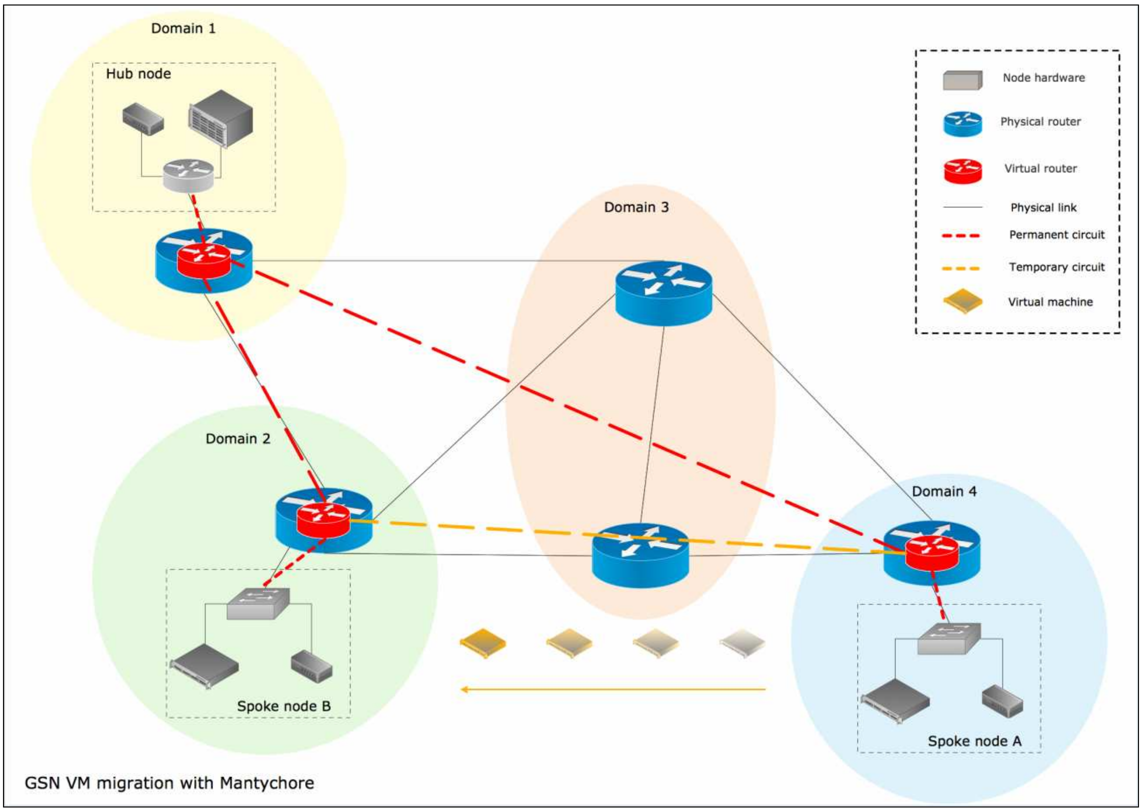

Figure 2.9 GSN virtual migration with Mantychore [46]

2.6.3 Environment

The GSN middleware is built around the IaaS Framework middleware platform and provides extensions to manage servers, virtual machines as well as power meters and PDUs. These exist as plugins for the platform and can be deployed in any IaaS container.

The software consists of the following components (see Figure 2.10):

- GSN Controller: Intelligence module for GSN infrastructure. Responsible for planning and executing the migration of virtual machines across GSN facilities according to the availability of renewable green energy.

- Network Manager: Responsible for control and management of network resources. Network services are exposed via the NSI Interface.

- Cloud Manager: Responsible for control and management of computing resources such as clouds, hosts, virtual machines, and storage.

- Facility Manager: Responsible for control and management of Facility Resources. A Facility Resource is an abstraction of a GSN facility or a node. It performs computations to predict availability of power based on current and long term weather forecast and historical data.

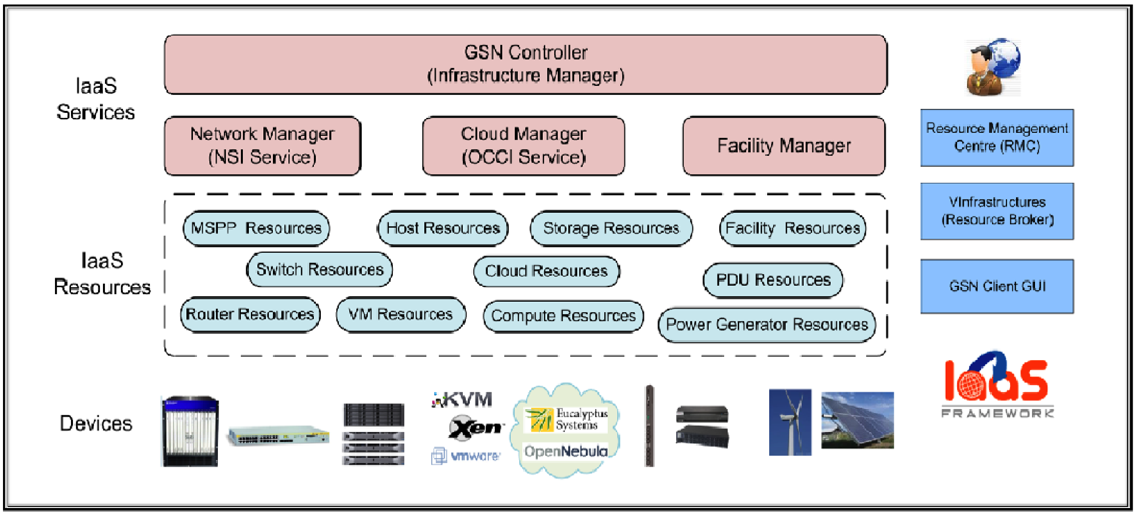

Figure 2.10 GSN software components [46]

The GSN controller performs the following tasks on a schedule (see Figure 2.11):

- Determine required migrations: Identifies virtual machines that must be migrated and the facilities to which they might be moved. When selecting receiving hosts, gives preference to hosts running in other green powered facilities having long expected number of operation hours. Must also consider hosting capacity on new hosts such as available operative systems and memory.

- Generate migration plan: Specifies pairs of hosts, together with a list of virtual machines to be migrated between them. Generates a migration plan that moves the virtual machines identified in the previous step.

- Optional migrations added to plan: If the plan generated in the previous step does not include moving virtual machines to hydro-powered facilities, then this step takes the opportunity to augment the plan to move virtual machines currently residing at hydro-powered facilities to sun and wind-powered facilities. Note that sun and wind-powered facilities take priority over hydro-powered facilities.

- Execute migration plan: Migration is executed in several parallel streams of serial migration. Migration is organized so as to not overload any link in the network guaranteeing a concrete bandwidth during the virtual machine migration.

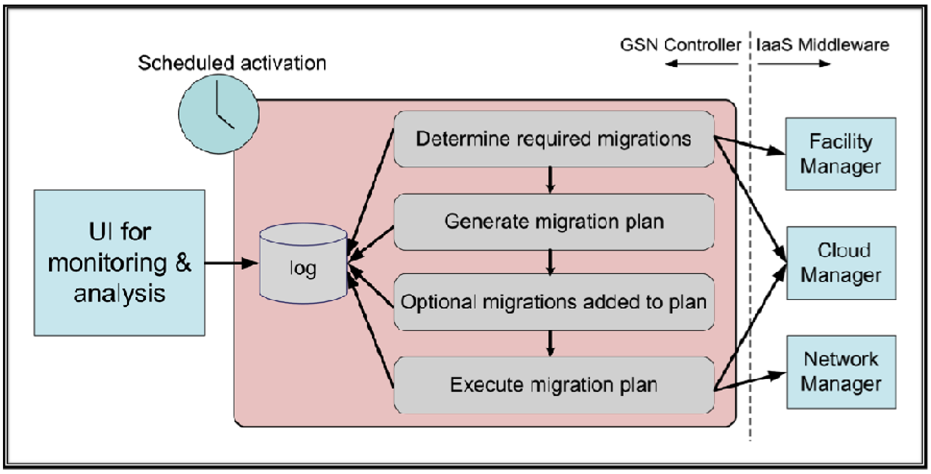

Figure 2.11 GSN controller tasks [46]

The interface to the GSN software is provided by a web-based client.

Other environmental constrains are:

- Multiple domains[3].
- Heterogenous hardware.
- Minimum bandwidth during virtual machine migration.

2.6.4 Actors involved

- Network administrators at each site.
- GSN node administrators at each site.
- Network administrators / NOC engineers responsible for the physical infrastructure.

2.6.5 Steps within Mantychore

To provision permanent spoke-hub connectivity:

1. Configure a virtual router at each end point; the virtual router at the hub end point can be reused, it just needs one interface per connected spoke.
2. Configure IP tunneling between the nodes by means of:
    a. Circuit cross-connects (CCC), practically turning the virtual routers into virtual switches (layer 2).
    b. MPLS (layer 2.5).
    c. IP-over-IP or GRE (layer 3).

[3] A domain is a set of one or more nodes whose servers share the same database.

To provision temporary migration circuits:

1. Source and destination nodes are determined outside of Mantychore.
2. Determine optimal path between source and destination, based on:
    - Available end-to-end bandwidth.
    - Available physical resources.
    - Resource efficiency.
    - Routing metrics.
3. Establish connection and inform client layer (GSN).
    - If connection cannot be established, try next best path.
    - If no connection can be established, inform client layer.
4. Virtual machines migration is performed outside of Mantychore.
5. On signal from client layer (GSN), tear down established connection.
    - Free up resources.

2.6.6 Manual steps

- Configure local connectivity and addressing for the nodes.
- Enable tunneling on the physical routers.

2.6.7 Success criteria

1. Establish permanent single-hop IP connectivity between a hub and a spoke node.
2. Establish temporary single-hop IP connectivity between two nodes, and then tear it down.
3. Migrate a live virtual machine between two nodes, without the risk of not having enough bandwidth during the migration.

2.6.8 Fail cases

- IP connectivity cannot be tunneled between the end points, through the physical infrastructure.
- Bandwidth cannot be guaranteed during virtual machine migration.

CHAPTER 3. MEASUREMENTS

As explained in Chapter 1, the main goal of this part is to measure the energy consumption on a Juniper MX480 router (see [47]) in order to obtain an analytical power model of a virtual-capable router. In this chapter, we will describe these measurements.

3.1 Previous works

To the best of our knowledge, [7] was the first study to develop a generic power model for router consumption created by Chabarek and Barford in 2008.

In their first set of experiments, they measured the power consumption of two idle router chassis with different combinations of line cards installed. These routers were a Cisco GSR 12008 and a Cisco 7507.

For each router they observed that the base system (chassis plus a router processor) consumes more than half the maximum observed power consumption for any configuration. These results lead directly to the conclusion that it is best to minimize the number of chassis that are powered on a router, and to maximize the number of line cards per chassis.

Regarding the second set of experiments, Chabarek and Barford set up a testbed consisting of 20 workstations for traffic generation, each of which had a Pentium 4 processor running at 2 GHz or better, at least 1 GB RAM, and an Intel Pro/1000 network interface card. Each host was configured to run either FreeBSD 5.4 or Linux 2.6. End host’s packet traffic was aggregated using Cisco 6500 routers and flowed through two Cisco GSR 12008s. Three parallel Gigabit Ethernet links connected the 6500s to the GSRs, and the GSRs were connected via a OC-48 link.

The focus of the evaluation was a 4-port Gigabit Ethernet engine 3 line card in the GSR chassis (Device Under Test). They established a series of configurations such as measuring the power consumption of the DUT with different size of packets. In addition to this, the difference between the idle state and the scenarios with packet traffic was about 20 W, that is to say, traffic has a little effect on power consumption.

These conclusions indicate the need to create a generic power model for router consumption in order to obtain a significant power savings in operational networks.

3.2 Description of the equipment

Initially we wanted to obtain a power model for a virtual router, but as we will describe this turned out to be impossible for several practical reasons, so we focused on measures of energy consumption which are divided into three groups: router with the minimal number of modules, router without traffic, and router with traffic.

The equipment we tried to model is a Juniper MX480 router with the following modules:

- Two SCB-MX.
- MX-MPC2-3D-RB.
- MS-DPC.

We follow an incremental approach, in which primarily these modules have been extracting one by one. They are described below.

The first two modules are the SCB-MX as shown in Figure 3.1:

Fig 3.1 Two SCB-MX integrated into SCB modules

The Switch Control Board (SCB) provides the following functions:

- Powers on and powers off Dense Port Concentrators (DPCs), Flexible PIC Concentrators (FPCs), and Modular Port Concentrators (MPCs).
- Controls clocking, system resets, and booting.
- Monitors and controls system functions, including fan speed, board power status, PDM status and control, and the craft interface.
- Provides interconnections to all the DPCs, FPCs, and MPCs within the chassis through the switch fabrics integrated into the SCB.

It should be noted that each SCB-MX module includes a Routing Engine, one active and one backup, which are integrated into the SCB modules of the chassis. For more information see [48].

Another module of the MX480 router is the one formed by the different ports. Figures 3.2, 3.3 and 3.4 show these ports:

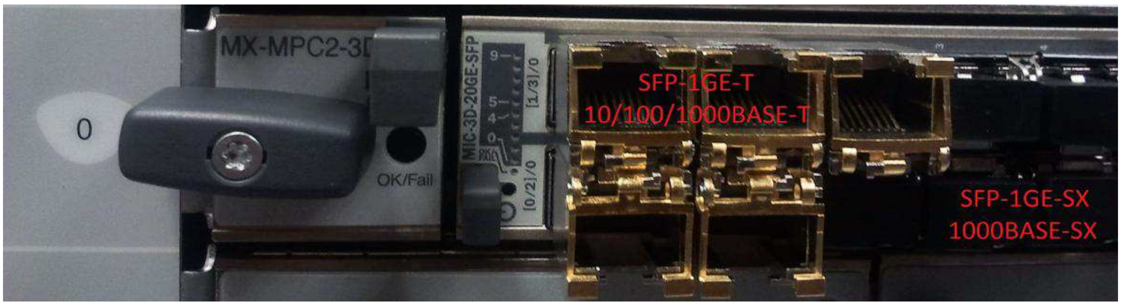

Fig 3.2 MIC-3D-20GE-SFP with 5 SFP-1GE-T and 7 SFP-1GE-SX

The MIC-3D-20GE-SFP comprises 20 Ethernet 10/100/1000 SFP ports and is located in the left half of the MX-MPC2-3D-RB in the module 0.

The SFP-1GE-T is a SFP module for Ethernet 1000BASE-T and consists of 5 units. In addition to this, the SFP-1GE-SX is a SFP module for Ethernet 1000BASE-SX and consists of 7 units.

On the other hand, in the same module MIC-3D-20GE-SFP we can find the remaining Ethernet ports. Figure 3.3 shows them below:

Fig 3.3 MIC-3D-20GE-SFP with 8 SFP-1GE-LX

The SFP-1GE-LX is a SFP module for Ethernet 1000BASE-LX and consists of 8 units.

In addition to this, within the MX-MPC2-3D-RB module but located on its right half we can find the MIC-3D-2XGE-XFP which consists of two 10Gigabit Ethernet XFP ports. Figure 3.4 depicts this module:

Fig 3.4 MIC-3D-2XGE-XFP with 1 XFP-10G-S and 1 XFP-10G-L-OC192-SR1

The XFP-10G-S is a XFP module for Ethernet 10GBASE-SR which works at 850 nm (300 meters), and the XFP-10G-L-OC192-SR1 is a XFP module for Ethernet 10GBASE-LR and STM-64/OC-192 which works at 1310 nm (10 km).

For more information about the MX-MPC2-3D-RB module and its corresponding ports see [49].

Finally, the last module which is part of the MX480 router is the MS-DPC as shown in Figure 3.5:

Fig 3.5 MS-DPC integrated into module 1

The Juniper Networks Multiservices DPC (MS-DPC) provides dedicated high-performance processing for flows and sessions, and integrates advanced security capabilities (Firewall, Deep Packet Inspection, Intrusion Prevention System, etc) that protect the network infrastructure as well as user data. For more information see [50].

Figure 3.6 shows the general appearance of the MX480 router with the 4 modules explained above. In this case, we made the second type of consumption measures (without traffic) to obtain a baseline scenario in order to compare it with the different scenarios with traffic (the last group of measures).

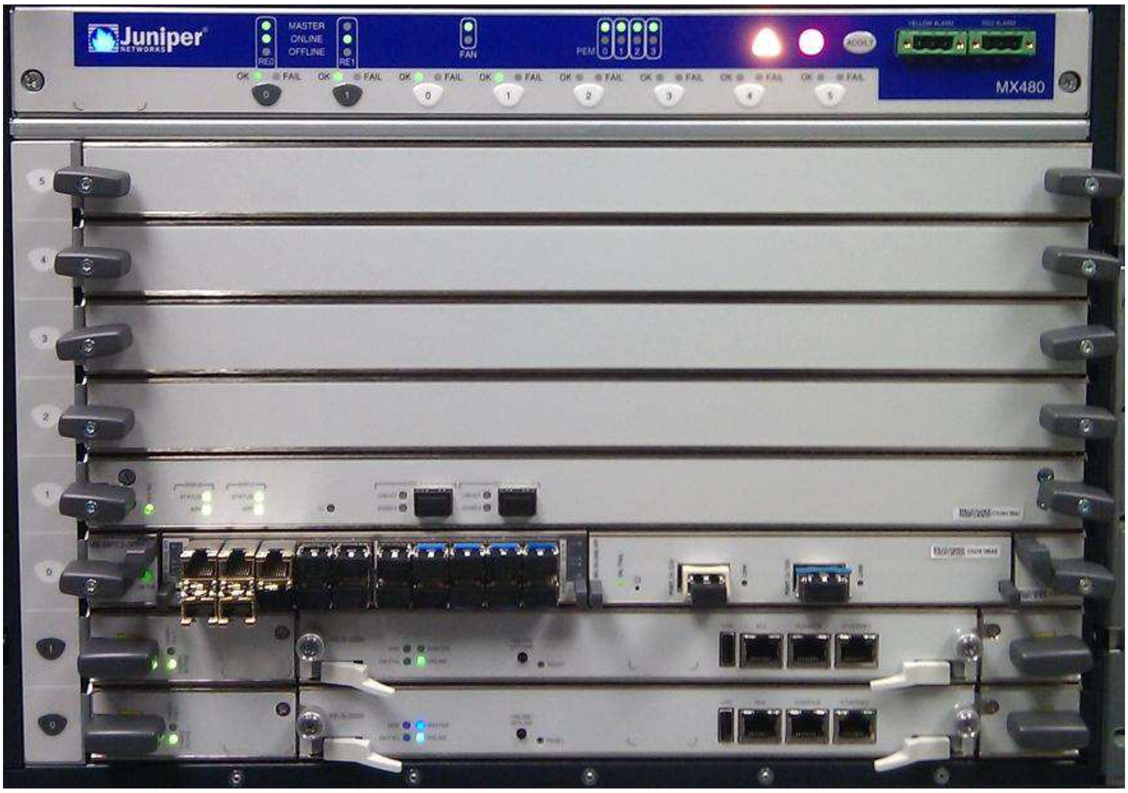

Fig 3.6 MX480 with the 4 modules bottom-up

The following subsections describe the scenarios established and their corresponding results in which two consumption measures have been made for each scenario: current and apparent power. The latter was taken into account because it is the total power consumed by all elements of the router (resistors, inductances, capacitors, etc).

In addition to this, a power distribution unit (PDU) is used in order to carry on this type of measures. A PDU is a device placed between the power line and the equipment under test, and is able to measure the voltage and current consumed by the device. The particular model used in our experiments is the Raritan DPXR-8-15 (see [51]). We were also able to obtain the total consumption of the router internally through the command `show chassis power`. For more information see [52].

It should also be noted that each sample was taken every 3 seconds (minimum time interval set by the PDU) for a time between 5 and 10 minutes (depending on the scenario), with a granularity of 55 mA for current measurements and 20 VA for measures of apparent power.

3.3 Measurements of isolated modules

As explained in the introduction of this chapter, this is the first group of consumption measures which is divided into four scenarios: router with a single SCB-MX module, router with a SCB-MX module and the MX-MPC2-3D-RB module, router with two modules SCB-MX and the MX-MPC2-3D-RB module, and finally the router with four modules (the MS-DPC module is added).

3.3.1 Scenario 0

Through this stage we want to calculate the energy consumption of the MX480 router with the minimum number of modules, in this case only with a SCB-MX module, because it needs at least one active Routing Engine to work. Figure 3.7 shows the schematic of this scenario:

Fig 3.7 MX480 only with one active SCB-MX module

In this case, the router consumption will be minimal because it is working with a single module.

3.3.2 Results

For each measure of current and apparent power we obtain two results, one for each power supply of the router. Note that the MX480 router supports the high-line (220 V) AC power configuration which requires two power supplies, with the third and fourth providing redundancy. For more information see [53].

Figure 3.8 shows the current consumed by each of the two power supplies of the router:

Fig 3.8 Current consumption Scenario 0 (Amps vs. time)

Taking into account the value of the different samples during the time interval of 5 minutes (with a total of 100 measurements approximately in intervals of 3 seconds), the average currents obtained with its corresponding standard deviation are 1.05 ± 0.27 A for the current outlet 3 and 1.14 ± 0.11 A for the current outlet 4 respectively.

On the other hand, Figure 3.9 shows the apparent power consumption for the two power supplies:

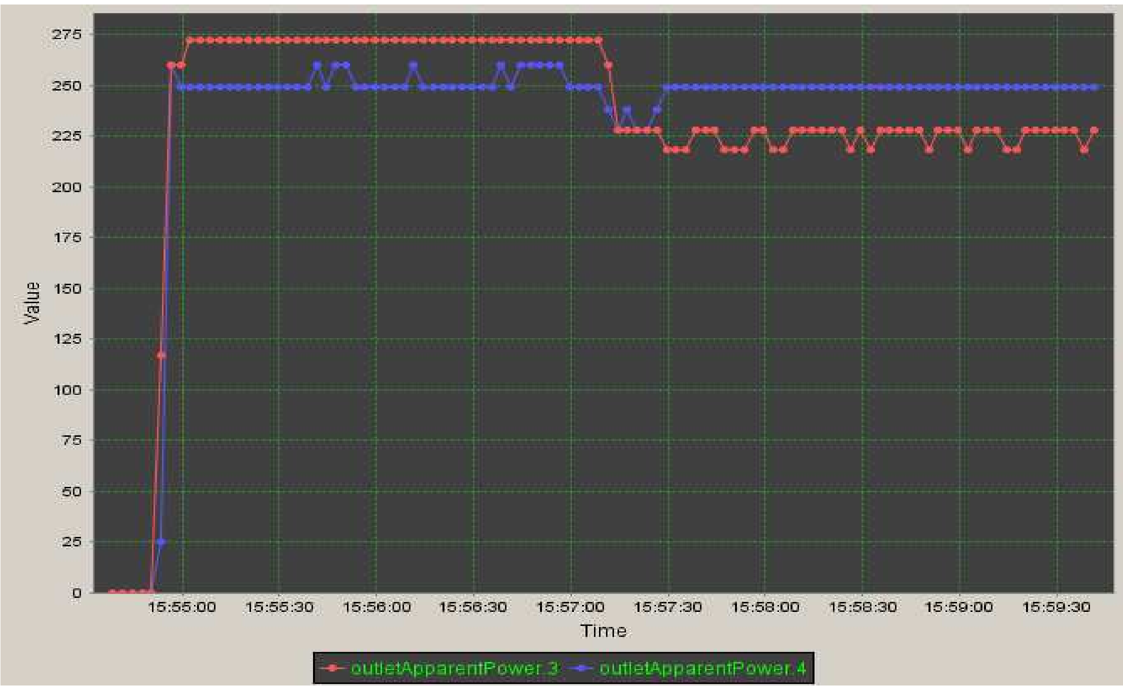

Fig 3.9 Apparent power consumption Scenario 0 (VA vs. time)

In this case, the time interval is the same as the previous because both current and apparent power samples have been taken at the same time. Note that this criterion is identical to that applied in all scenarios.

In addition to this, the results obtained are 234.24 ± 60 VA for the apparent power outlet 3 and 247.76 ± 24 VA for the apparent power outlet 4.

3.3.3 Scenario 1

Figure 3.10 shows the scenario:

Fig 3.10 MX480 with one SCB-MX module and the MX-MPC2-3D-RB module active

Since the difference of this scenario and the previous is the addition of the MX-MPC2-3D-RB module, we want to assess the potential increase of consumption caused by this module.

3.3.4 Results

We will check through the following graphics how energetic consumption has increased by adding another module, in this case which contains the various ports. Figure 3.11 shows the current consumption for the two power supplies:

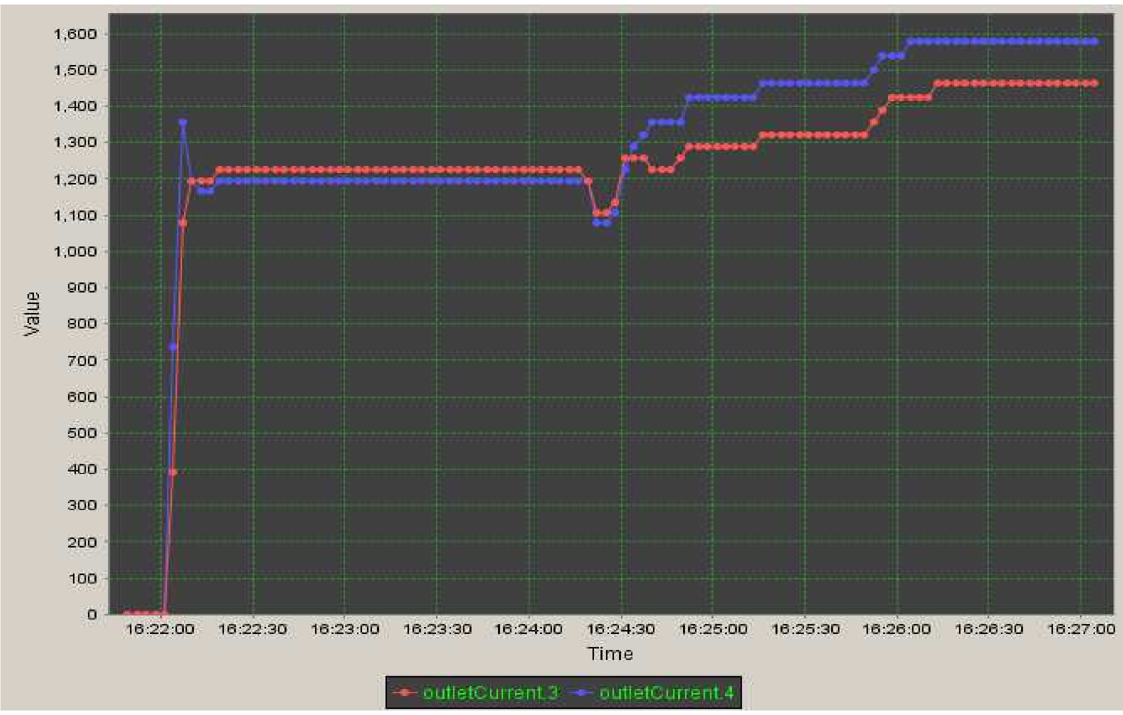

Fig 3.11 Current consumption Scenario 1 (Amps vs. time)

This scenario has taken the same time interval in comparison with the previous (5 minutes), and the current consumption results are the following: 1.21 ± 0.31 A for the current outlet 3 and 1.35 ± 0.18 A for the current outlet 4 respectively.

As we can see, adding the MX-MPC2-3D-RB module to the router, the current has increased by 12% for the two sources compared with the previous results.

On the other hand, the apparent power results are shown in Figure 3.12:

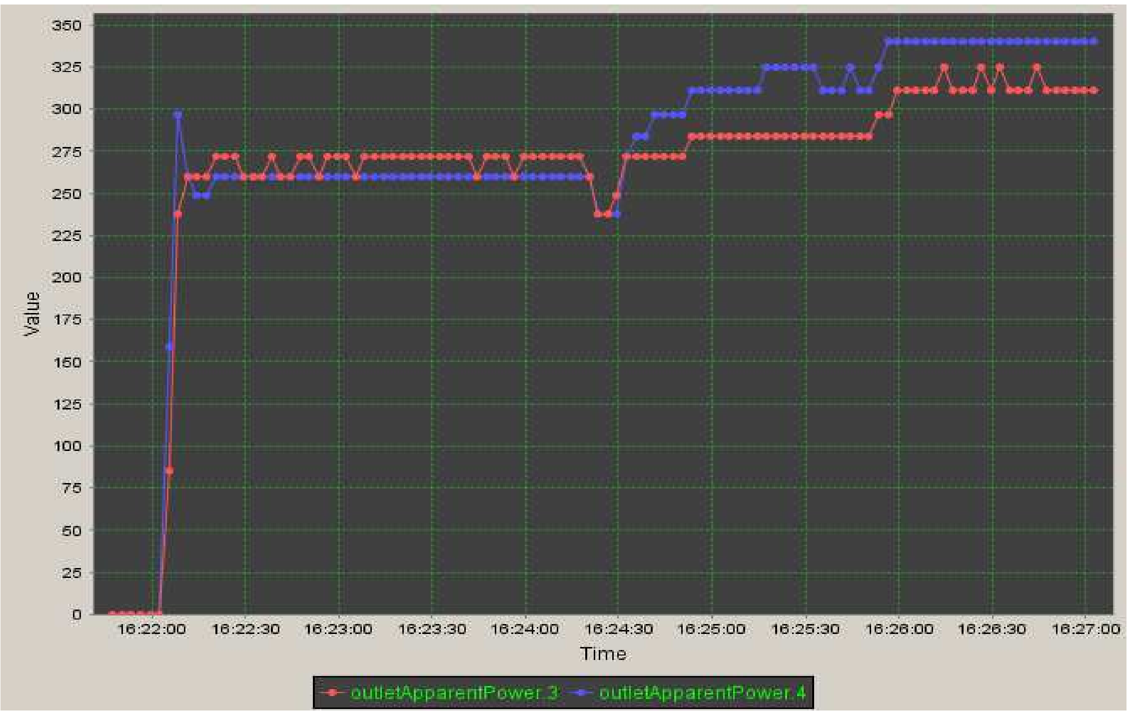

Fig 3.12 Apparent power consumption Scenario 1 (VA vs. time)

We have obtained 262.85 ± 71 VA for the apparent power outlet 3 and 291.41 ± 38 VA for the apparent power outlet 4.

Comparing these results with the previous scenario, the apparent power has increased by 29 VA for the first power supply and 44 VA for the second respectively, which is the same as saying that the total power consumption has increased by 12% approximately, the same increase as in the case of the current.

3.3.5 Scenario 2

Taking into account the increase in consumption that we previously achieved, in this scenario we will add another module in order to check how this increase varies. Figure 3.13 shows the scenario:

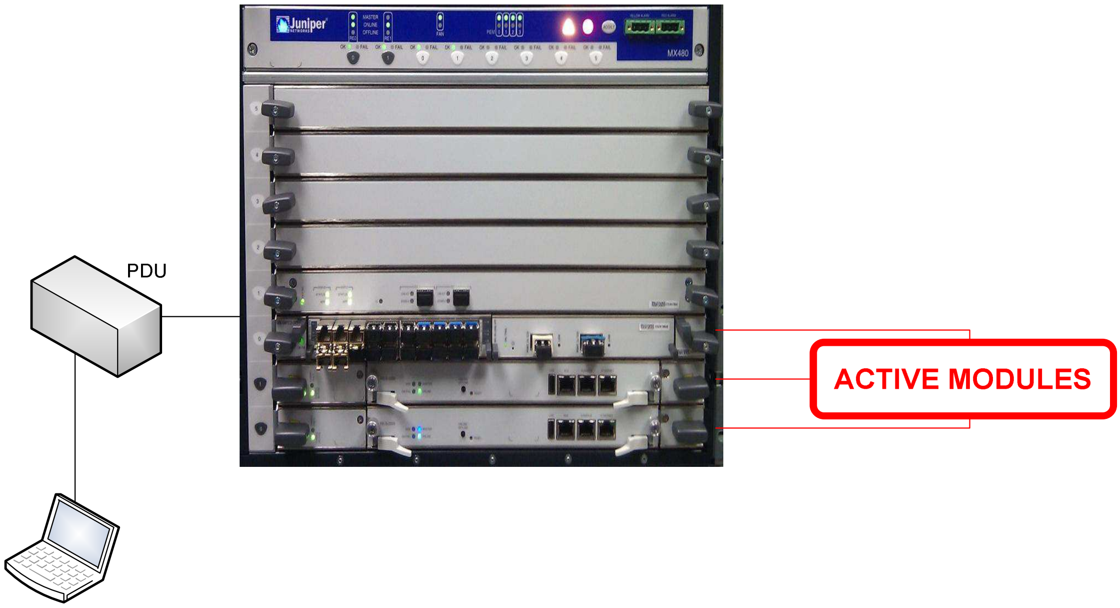

Fig 3.13 MX480 with 2 SCB-MX modules and the MX-MPC2-3D-RB module active

As we can see, the only difference with the above scenario is the addition of the other SCB-MX module (Routing Engine backup).

3.3.6 Results

Figure 3.14 shows the current drawn by the two power supplies when the router is working with the three modules mentioned above:

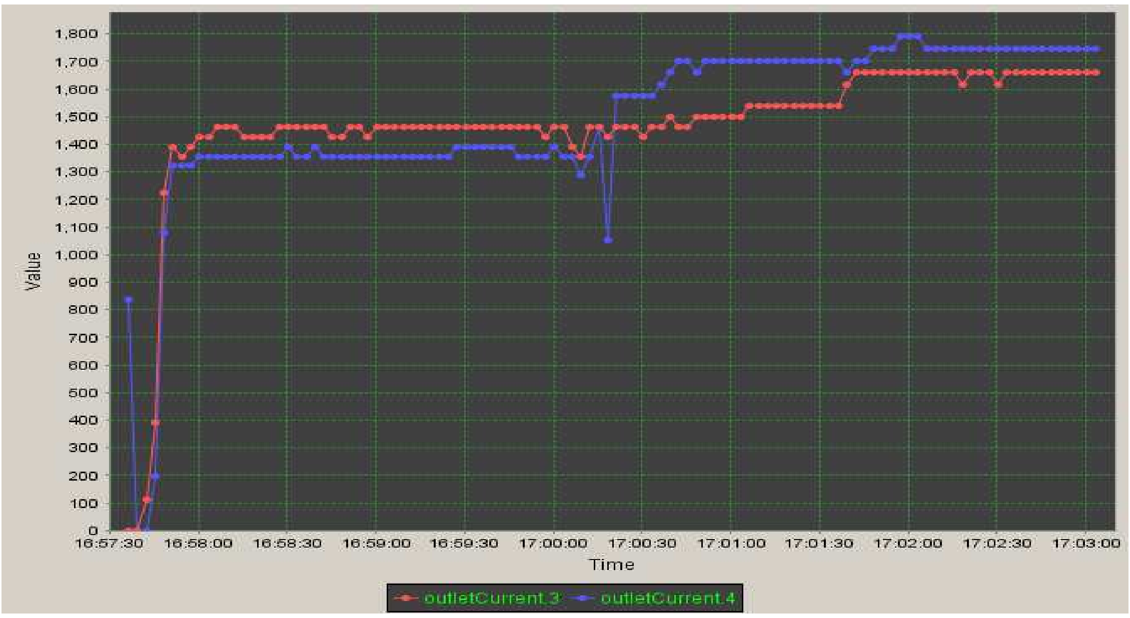

Fig 3.14 Current consumption Scenario 2 (Amps vs. time)

For this scenario, we took a time interval of five and a half minutes (with a total of 110 measurements approximately in intervals of 3 seconds) because it is the time that the samples have been tended to be constant.

In addition to this, the average currents obtained are 1.46 ± 0.29 A for the current outlet 3 and 1.5 ± 0.28 A for the current outlet 4, practically the same result.

If we compare these results with the previous scenario, the current has increased by 12% for the first power supply and by 11% for the second when we add the other SCB-MX module, approximately the same percentages.

Then, Figure 3.15 shows the results corresponding to the apparent power consumption of this scenario:

Fig 3.15 Apparent power consumption Scenario 2 (VA vs. time)

For the apparent power outlet 3, the consumption obtained is 313.22 ± 62 VA (12% more when it is compared to the previous scenario), and for the apparent power outlet 4 the result is 319.63 ± 62 VA (11% more in comparison with the previous scenario). If we consider these percentages, the increase is the same as in the case of the current.

In conclusion, Table 3.1 shows the total increase in consumption for the two power supplies and their corresponding measures at this point:

	Power supply 1 (current)	**Power supply 2 (current)**	**Power supply 1 (ap. power)**	**Power supply 2 (ap. power)**
Increase Scenario 0-1	12%	12%	12%	12%
Increase Scenario 1-2	12%	11%	12%	11%
Total	24%	23%	24%	23%

Table 3.1 Total increase in consumption in Scenario 2

As we can see, both the total current increase as the apparent power is practically the same for the two power supplies.

3.3.7 Scenario 3

This is the last scenario of this group of measures and is based on the router startup with all the modules active. Figure 3.16 shows the schematic:

Fig 3.16 MX480 with all the modules active

Since the difference of this scenario and the previous is the addition of the MS-DPC module, we want to assess the increase of power consumption caused by this module.

3.3.8 Results

Figure 3.17 shows the current consumed by the two power supplies of the router as their corresponding modules are starting:

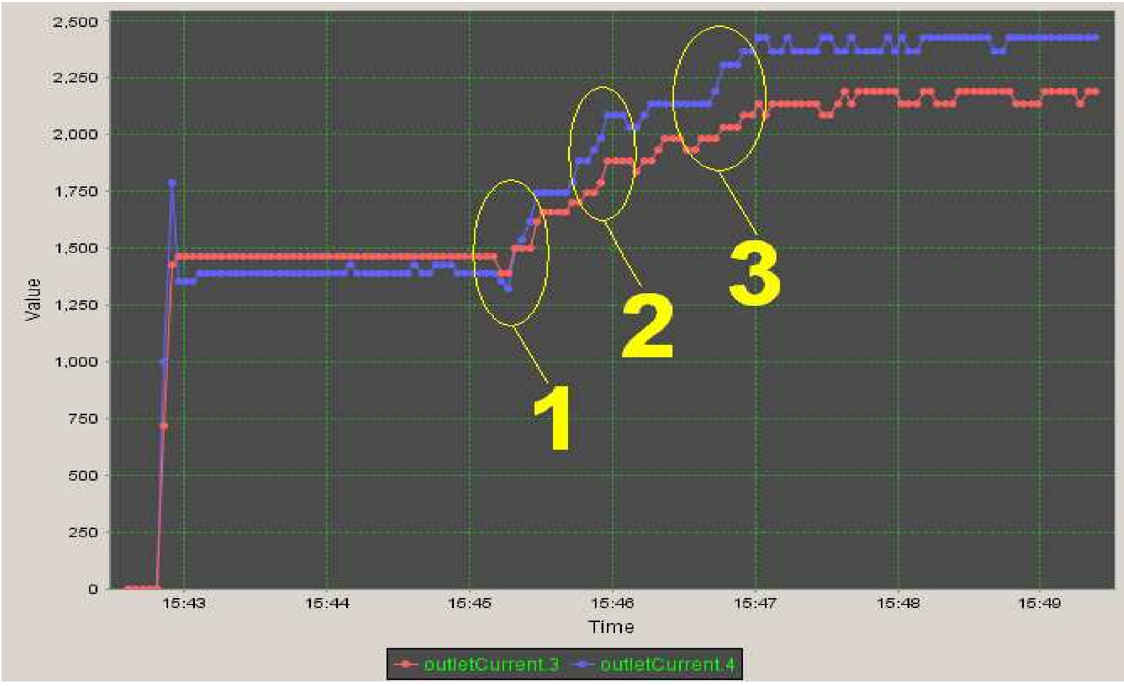

Fig 3.17 Current consumption Scenario 3 (Amps vs. time)

As we can see, this is the boot order that the four modules follow:

1. Two SCB-MX modules.

2. MS-DPC module.

3. MX-MPC2-3D-RB module.

Considering that the router takes 6 minutes and a half to boot (with a total of 130 measurements approximately in intervals of 3 seconds), the average currents for each power supply are the following: 1.61 ± 0.46 A and 1.76 ± 0.4 A respectively.

Taking into account these results, we can ensure that the current increases by 11% for the current outlet 3 and by 12% for the current outlet 4 when all the modules are active.

On the other hand, Figure 3.18 shows the apparent power consumed by the two power supplies:

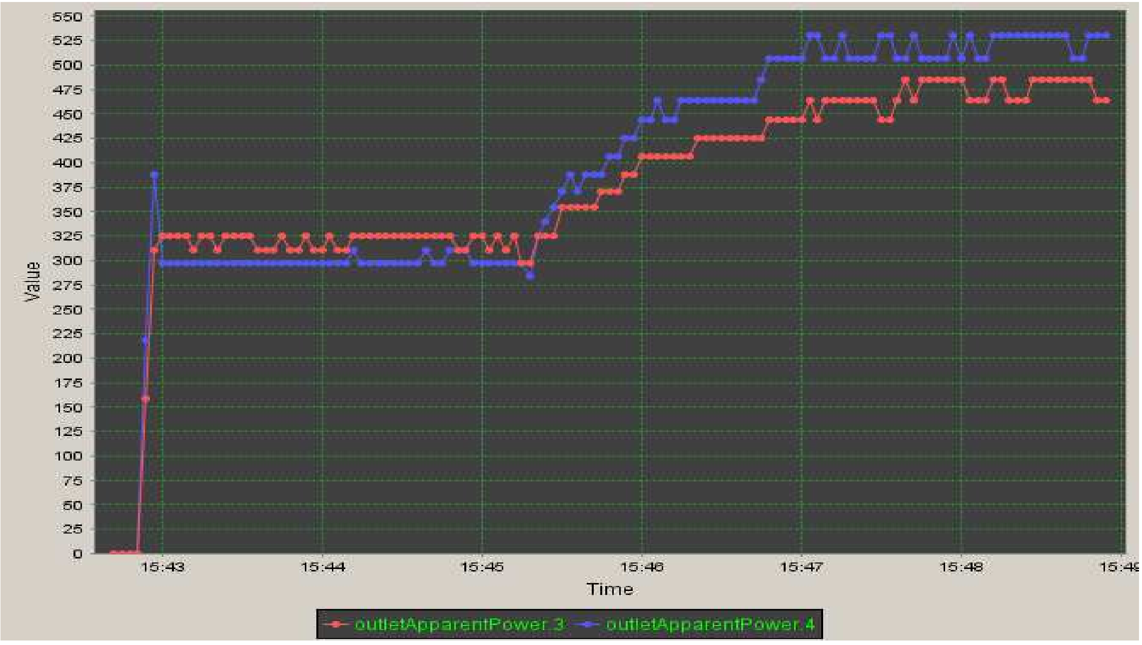

Fig 3.18 Apparent power consumption Scenario 3 (VA vs. time)

Considering the previous time interval, we have obtained 352.25 ± 94 VA for the first power supply and 387.3 ± 92 VA for the second. Hence the apparent power has increased by 13% for the two sources in comparison with the results of the previous scenario.

In conclusion, Table 3.2 shows the total increase in consumption for the two power supplies and their corresponding measures when the router is working with all the modules active:

	Power supply 1 (current)	**Power supply 2 (current)**	**Power supply 1 (ap. power)**	**Power supply 2 (ap. power)**
Increase Scenario 0-1	12%	12%	12%	12%
Increase Scenario 1-2	12%	11%	12%	11%
Increase Scenario 2-3	11%	12%	13%	13%
Total	35%	35%	37%	36%

Table 3.2 Total increase in consumption in Scenario 3

In the next section, we will explain the following group of measures where the MX480 router is running for a period of time and its corresponding energy consumption.

3.4 No traffic measures

This group of measures consists of a single scenario where we will calculate the router consumption without applying any traffic.

3.4.1 Scenario 0

This is the baseline/idle scenario, that is to say, two power supplies of the router directly connected to the PDU, with no other cards attached. There are no traffic flows through the router.

Figure 3.19 shows the scenario:

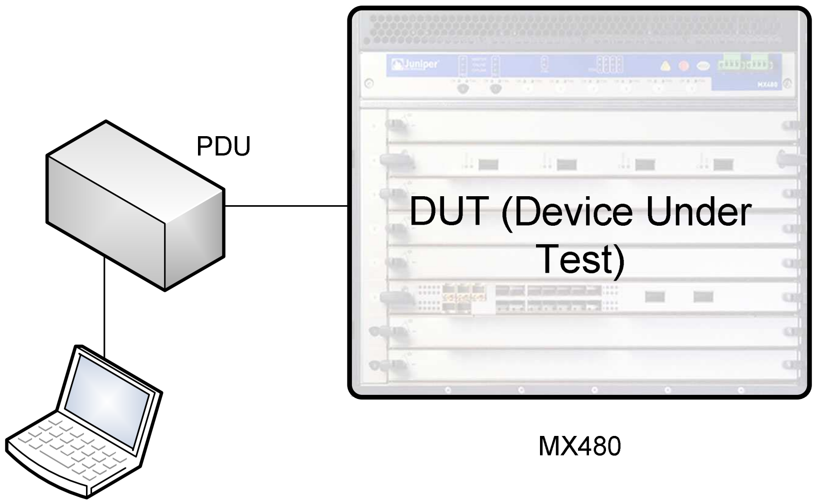

Fig 3.19 Baseline scenario

Since the difference of this scenario and the previous is the operation of the MX480 router with all its modules, we want to assess the potential increase of consumption caused by this activity.

3.4.2 Results

Figure 3.20 shows the current consumed by the two power supplies of the router:

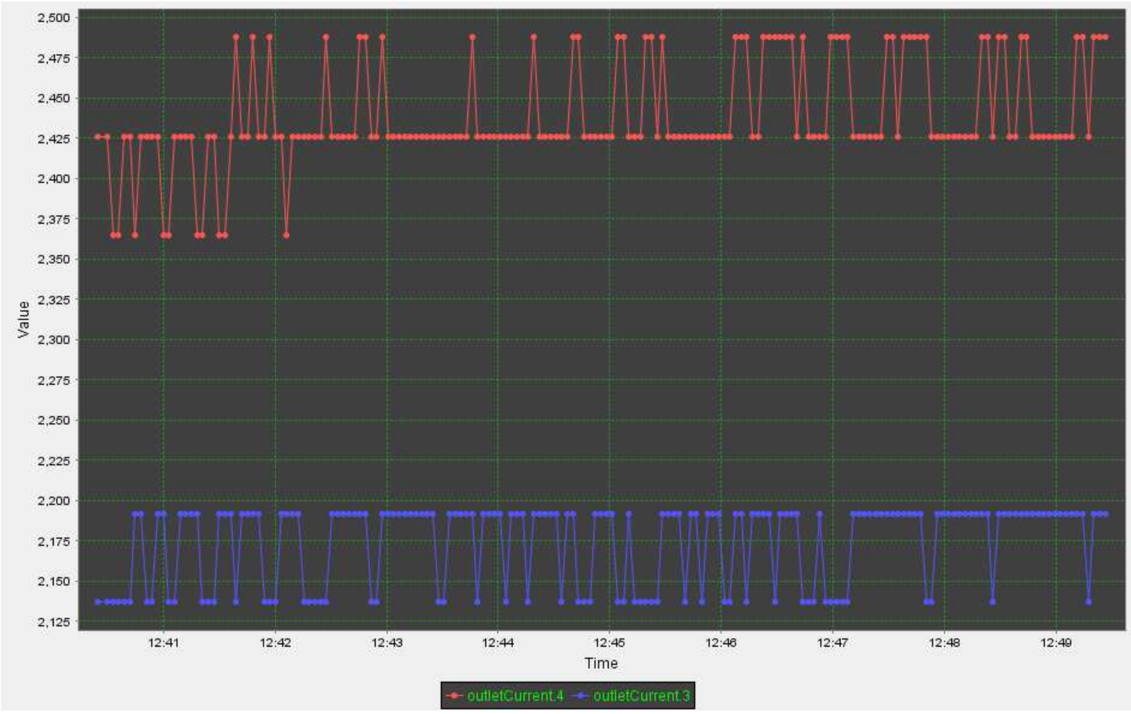

Fig 3.20 Current consumption Scenario 0 (Amps vs. time)

Taking into account the value of the different samples during the time interval of 10 minutes (with a total of 200 measurements approximately in intervals of 3 seconds), the average currents obtained are 2.14 A for the current outlet 3 and 2.43 A for the current outlet 4 respectively.

It should be noted that the jumps between samples are 55 mA (granularity) due to fluctuations. In addition to this, we can say that the router is so stable that does not generate abrupt changes as we can see in Figure 3.20.

With respect to the difference in consumption between the two power supplies, it seems the router does not perform load balancing and is using one of the sources in order to feed some modules and the other to feed the rest. This reasoning would explain the asymmetric consumption obtained because not all the modules and even the chassis consume the same as we have seen in previous scenarios.

On the other hand, Figure 3.21 shows the apparent power consumed:

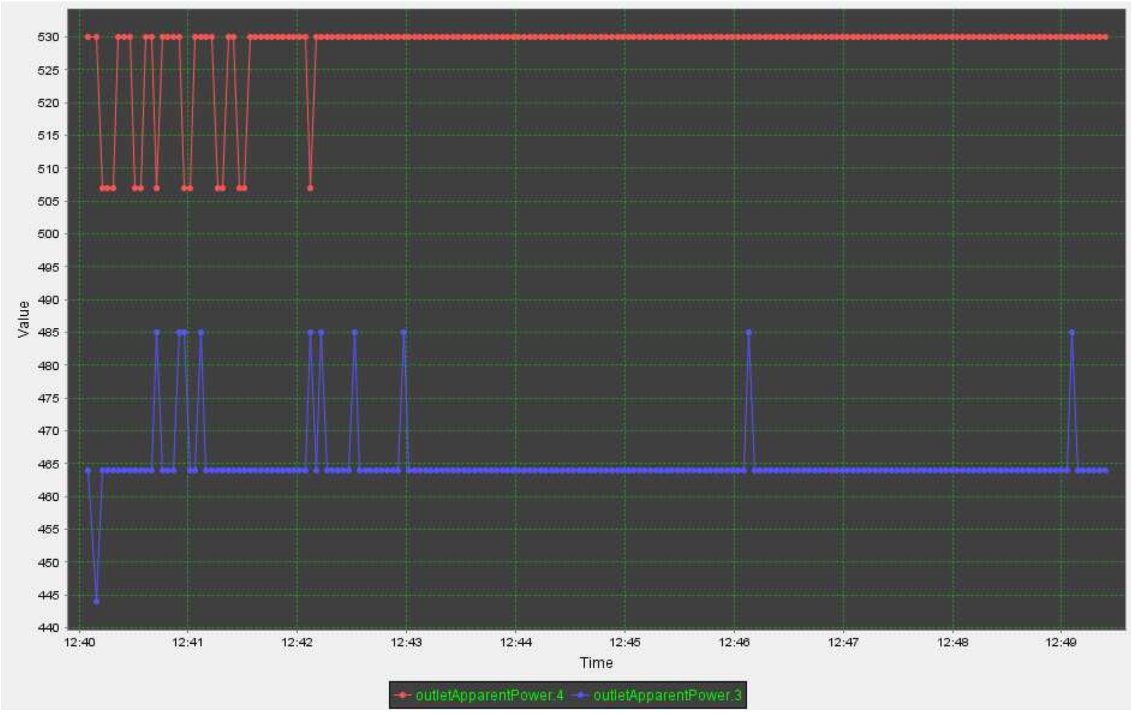

Fig 3.21 Apparent power consumption Scenario 0 (VA vs. time)

In this case, the time interval is the same as the previous because both current and apparent power samples have been taken at the same time. Therefore the results obtained are 465.02 VA for the apparent power outlet 3 and 528.39 VA for the apparent power outlet 4.

On the other hand, Figure 3.22 shows the power factor associated with each power supply:

Fig 3.22 Power factor of the two sources (º vs. time)

As we can see, the results obtained are 0.9 for the power factor outlet 3 and 0.92 for the power factor outlet 4. Multiplying the power consumed by each of the sources (465.02 VA and 528.39 VA respectively) for the corresponding power factor and adding them together, we obtain 1091.04 W, practically the same result obtained internally by the router, in this case 1093 W.

3.5 Traffic measures

This is the last group of consumption measures which is divided into five scenarios: router with 10 Gbit/s incoming traffic, router with traffic 17 Gbit/s incoming traffic, router with 100 Mbit/s incoming traffic and 2 logical systems, router with traffic 100 Mbit/s incoming traffic and 5 logical systems, and finally the router with 1 Gbit/s incoming traffic and 5 logical systems.

3.5.1 Scenario 0

The data plane scenario is composed of the router with all modules operating and includes a constant bit rate traffic of 10 Gbit/s, generated and processed by the Agilent N2X E7318A analyzer [54] connected to the XFP-10G-L-OC192-SR1 module [49] of the router via a second window SC-LC single mode fiber.

The Agilent N2X allows to perform detailed functional and performance testing on SONET/SDH interfaces. It is for this reason that the analyzer is connected to the XFP-10G-L-OC192-SR1 router module (10GBASE-LR and STM-64/OC-192).

Since the difference between this scenario and the previous (see section 3.4) is the addition of traffic, we want to assess the potential increase of consumption caused by this traffic, but finally we were unable to do this evaluation because the analyzer works with a wavelength of 1550 nm and the XFP-10G-L-OC192-SR1 module with a wavelength of 1310 nm.

3.5.2 Scenario 1

This second data plane scenario with traffic includes the router with all the modules, and various traffic flows: a 10 Gbit/s constant bit rate flow, and a link aggregation of 7 Gbit/s. For the traffic generation, Mgen [55] is used for creating constant bit rate UDP traffic with 9014 byte packets (Ethernet Jumbo frames).

In parallel, a SMC8824M switch [56] is used for connecting seven of its ports to the respective ports of the router in order to generate the link aggregation mentioned above. Finally, the aim of this experiment is to stress the router with 17 Gbit/s of traffic.

Figure 3.23 shows the schematic of this scenario:

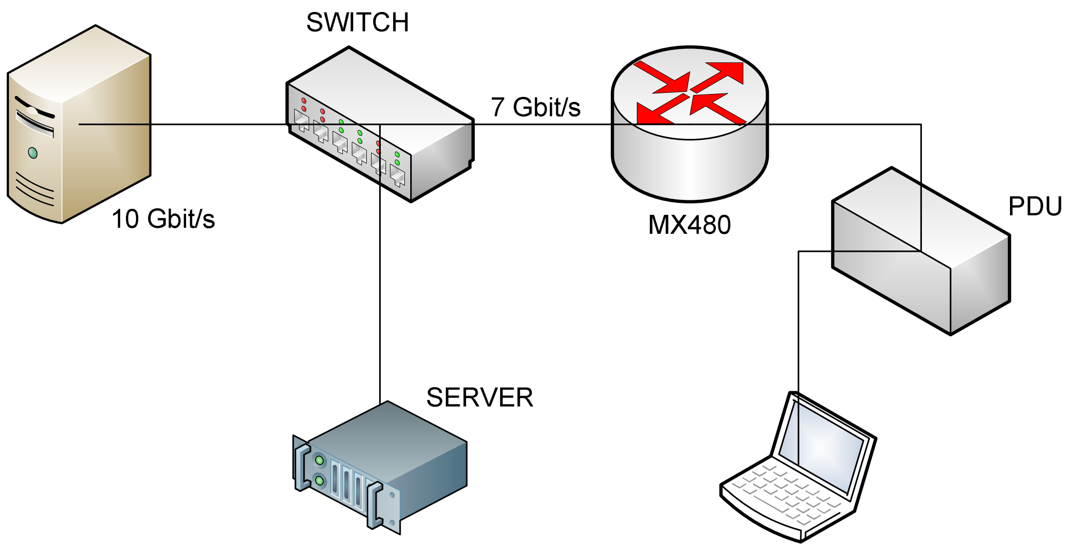

Fig 3.23 Data plane scenario with 10 Gbit/s + link aggregation of 7 Gbit/s

The configuration of this scenario is described below:

- The Mgen traffic generator is connected to the switch port number 25.
- Switch port 26 is connected to the XFP-10G-S router module (10GBASE-SR) [49]. In this point, the router sends and receives a 10 Gbit/s constant bit rate flow.
- Switch ports 13, 14, 16, 18 and 20 are connected to the respective 1000BASE-T ports of the router in the SFP-1G-T module (see introduction).
- Switch port 21 is connected to one of the 1000BASE-SX ports of the router in the SFP-1GE-SX module (see introduction).
- Switch port 22 is connected to one of the 1000BASE-LX ports of the router in the SFP-1GE-LX module (see introduction).
- A server is connected to the switch port 5 in order to check the amount of traffic flowing via Wireshark [57].

In addition to this, Figure 3.24 shows the assembly of the various links between the switch and router:

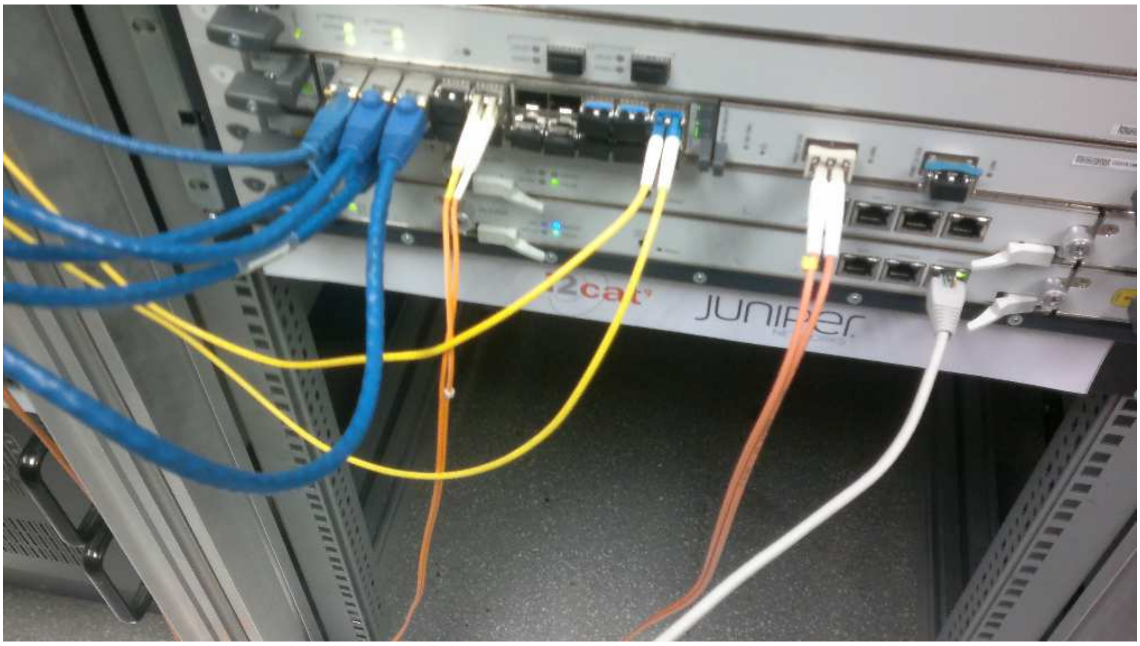

Fig 3.24 Scenario 1 assembly

As can be seen, five blue wires (Category 6 UTP cable) are used to connect the five 1000BASE-T ports of the router, two LC-LC multimode fibers in order to connect one 1000BASE-SX router port and the XFP-10G-S router module respectively, and finally, one LC-LC single mode fiber in order to connect one 1000BASE-LX router port.

3.5.3 Results

Figure 3.25 shows the current consumed by the two power supplies of the router:

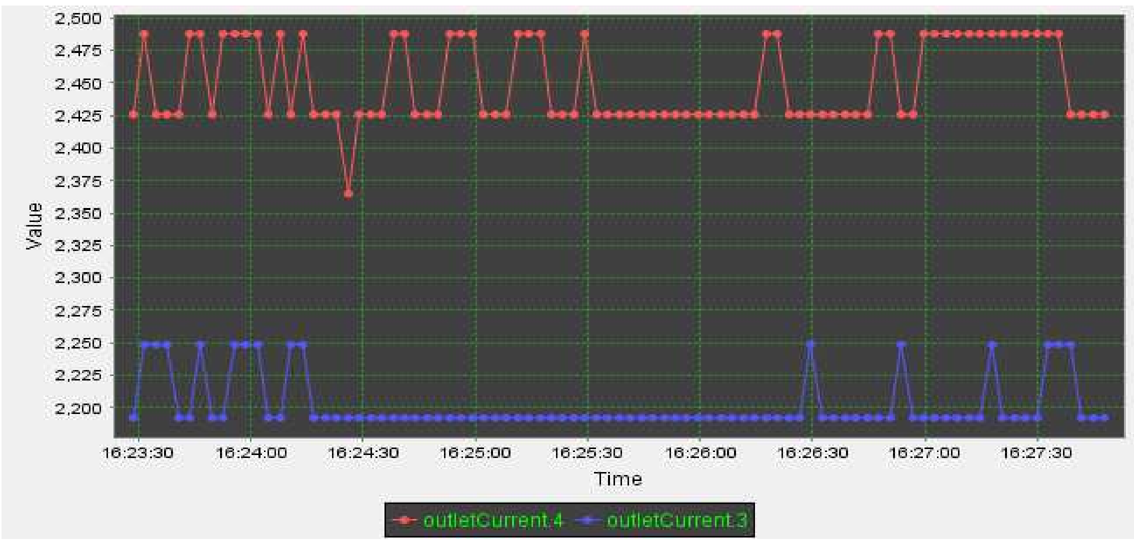

Fig 3.25 Current consumption Scenario 1 (Amps vs. time)

In this case, the time interval is 4 minutes and a half, approximately 90 measurements in intervals of 3 seconds, and the averages currents obtained are the following: 2.2 A and 2.45 A respectively.

Comparing these results with those obtained in the scenario without traffic (see section 3.3.2), practically there is no change between them because the consumption has increased by only 1%.

The conclusion is that the generated traffic is negligible and does not increase the energetic consumption of the power supplies because MX480 is very powerful (1.4 Tbps) and needs more traffic to notice some variation.

On the other hand, Figure 3.26 shows the apparent power results:

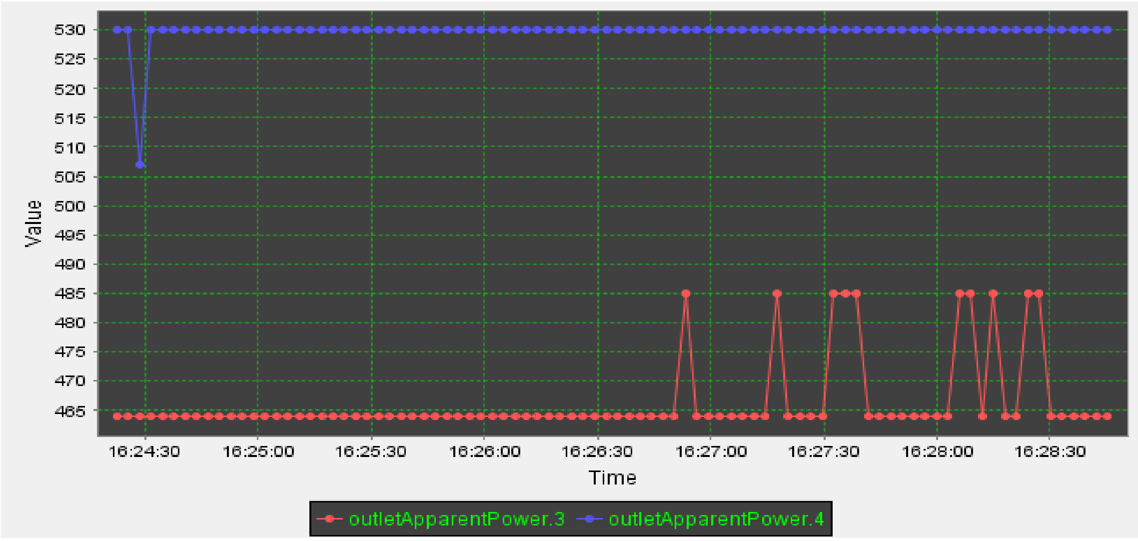

Fig 3.26 Apparent power consumption Scenario 1

Taking into account the value of the different samples during the interval of time established (4 minutes and a half), the average apparent powers obtained are 466.39 VA and 529.49 VA respectively.

There is a small variation of 1 VA for the two power supplies compared with the results obtained in section 3.3.2, practically negligible. The conclusion is the same as in the current-related results mentioned above.

3.5.4 Scenario 2

This scenario tries to study the influence of logical systems in the router power consumption. A logical system is basically a partition of a physical router that performs independent routing tasks and offers an effective way to maximize the use of a single router. For more information about the logical system configuration see [58]. Regarding traffic, a 100 Mbit/s constant bit rate traffic flow is injected with Mgen, and two logical systems are configured.

Figure 3.27 shows the schematic of this scenario:

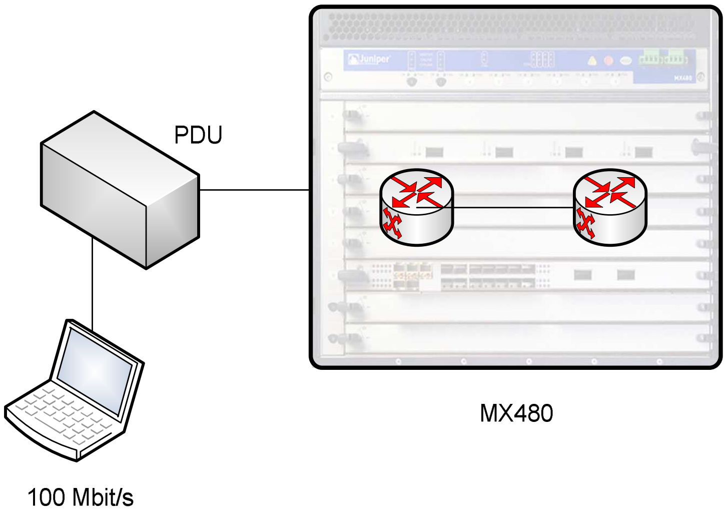

Fig 3.27 Data plane scenario with 100 Mbit/s + 2 logical systems

The main idea was to connect the minimum number of logical systems, in this case two, through a logical tunnel interface which allow us to connect logical routers. To connect two logical routers, we have to configure a logical tunnel interface on both logical routers. Then we configure a peer relationship between the logical tunnel interfaces, thus creating a point-to-point connection. For more information see [59].

Unfortunately we were not able to configure any logical tunnel interface because the router is not equipped with a special module (Tunnel Services PIC [60]) which is required in order to set up this type of interfaces.

As alternative, four 1000BASE-T ports of the MX480 router have been physically connected between them (through Category 6 UTP cable) in pairs, in order to emulate the logical tunnel by assigning each pair to a logical system. Figure 3.28 shows this scheme:

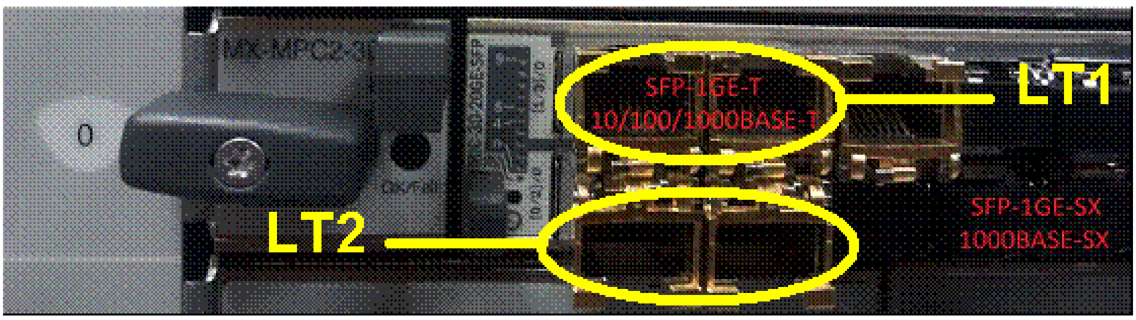

Fig 3.28 Two logical tunnels to connect the two logical systems

As we can see, the links marked in yellow are the two cables used to connect the ports of the router and thus create the two logical tunnels physically.

3.5.5 Results

The current consumed by the two power supplies of the router is shown in Figure 3.29:

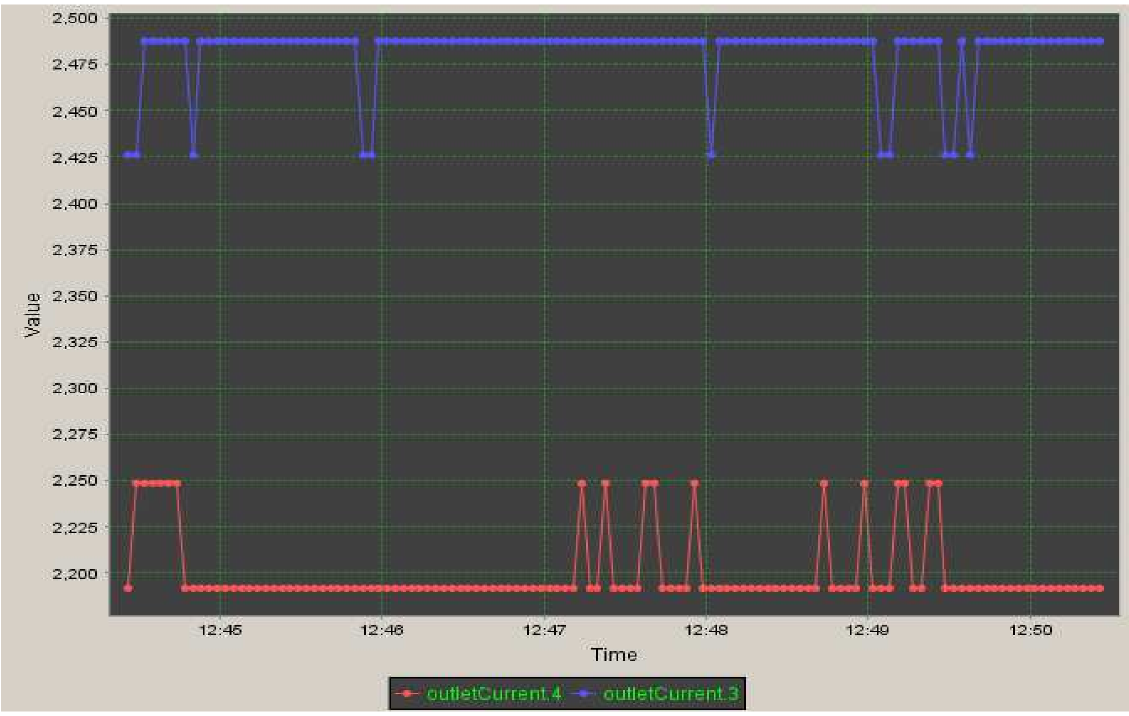

Fig 3.29 Current consumption Scenario 2

The average currents obtained are 2.2 A and 2.48 A respectively during the interval of time established, in this case 6 minutes approximately, with a total number of samples of 120 during intervals of 3 seconds.

The increased consumption compared to the previous scenario has been 1% for the second power supply because the first has not changed.

In conclusion, the two logical systems do not stress the MX480 router and, therefore, energetic consumption does not vary, such as in the above scenario.

On the other hand, Figure 3.30 shows the apparent power results:

Fig 3.30 Apparent power consumption Scenario 2

Taking into account the value of the different samples during the interval of time established, the average apparent powers obtained are 466.54 VA and 529.74 VA for the respective power supplies.

Comparing these results with those obtained in Scenario 1, there is no change between them. The conclusion is the same as in current results mentioned above.

3.5.6 Scenario 3

This is the same data plane that Scenario 2, but the only difference is the number of logical systems, in this case five. Note that this is the maximum number of logical systems that we have created since we did not have more optical fibers to connect the other ports.

Figure 3.31 shows the assembly of this scenario:

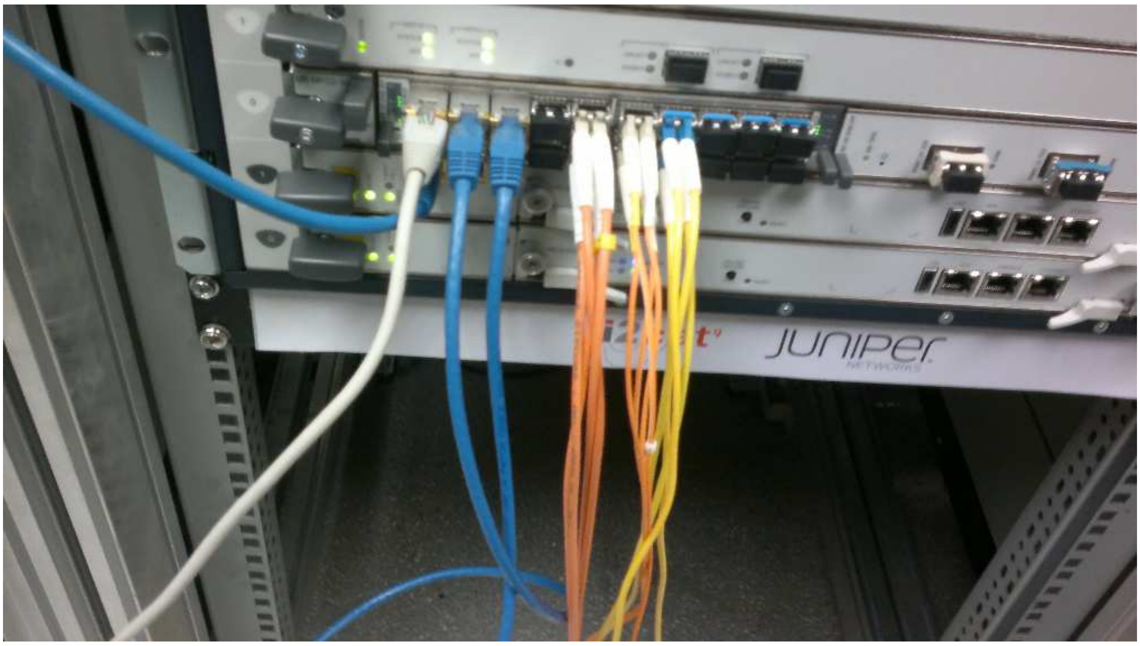

Fig 3.31 Scenario 3 assembly

This experiment is designed to measure power use when the MX480 router has five logical systems configured and the Mgen generator creates 100 Mbit/s of UDP traffic.

As can be seen in Figure 3.31, four 1000BASE-T ports of the router have physically connected between them through Category 6 UTP cable and in pairs. Each pair is assigned a logical system, as in Scenario 2.

In addition to this, four 1000BASE-SX ports have physically connected between them through two pairs of LC-LC multimode fiber. Finally, two 1000BASE-LX ports have also physically connected between them through one LC-LC single mode fiber.

Finally, there are five logical systems in total, one per connection between two ports of the router.

3.5.7 Results

Figure 3.32 shows the current consumed by the two power supplies of the router:

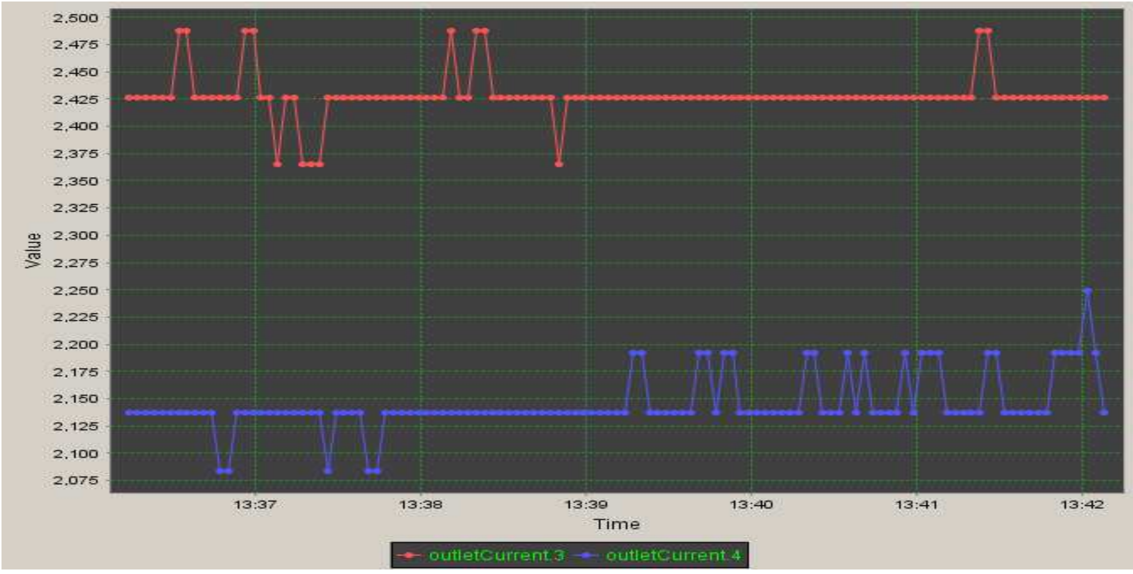

Fig 3.32 Current consumption Scenario 3

In this case, the average currents obtained are 2.16 A for the current outlet 4 and 2.44 A for the current outlet 3 during the time interval of 6 minutes approximately.

In comparison with the results obtained in Scenario 2, there is a small variation between them. Current consumption decreases 40 mA and 60 mA respectively for the two power supplies.

Surprisingly, increasing the number of logical systems, current consumption slightly decreases. In fact, this consumption should have been higher, increasing the number of logical systems.

The apparent power consumed is shown in Figure 3.33:

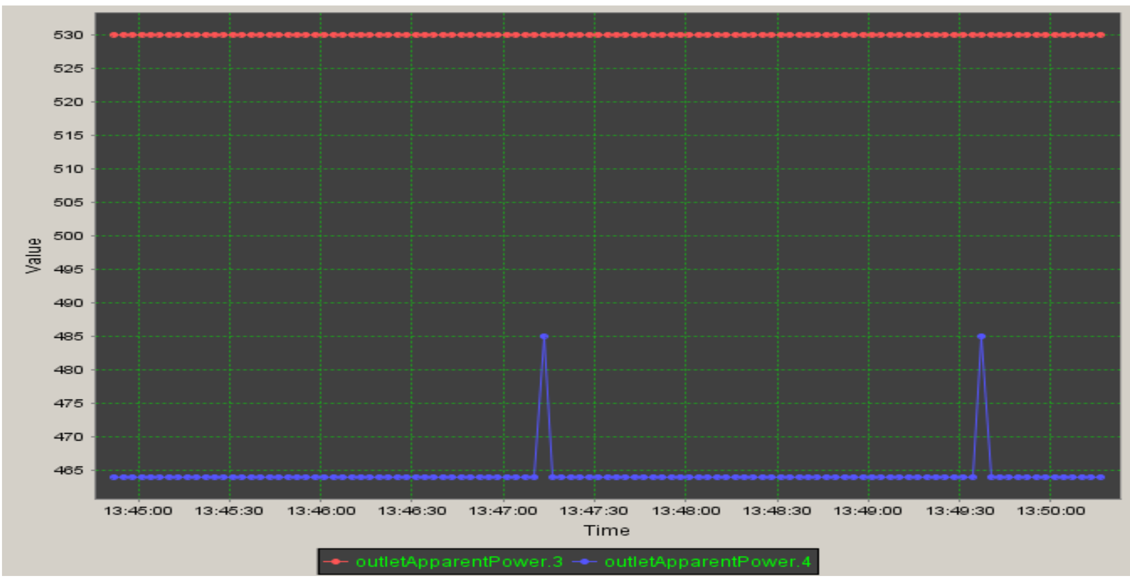

Fig 3.33 Apparent power consumption Scenario 3

Taking into account the value of the different samples during the interval of time established, the average apparent powers obtained are 464.38 VA and 530 VA for the respective power supplies.

Comparing these results with those obtained in Scenario 2, there is a small variation of 2 VA in one of the power supplies, but it is not conclusive. In this case, we could use the same argument as in current results mentioned above, but for the case of apparent power.

3.5.8 Scenario 4

This is the same data plane that Scenario 3, but the only difference is the amount of traffic, in this case 1 Gbit/s.

Figure 3.34 shows the schematic of this scenario:

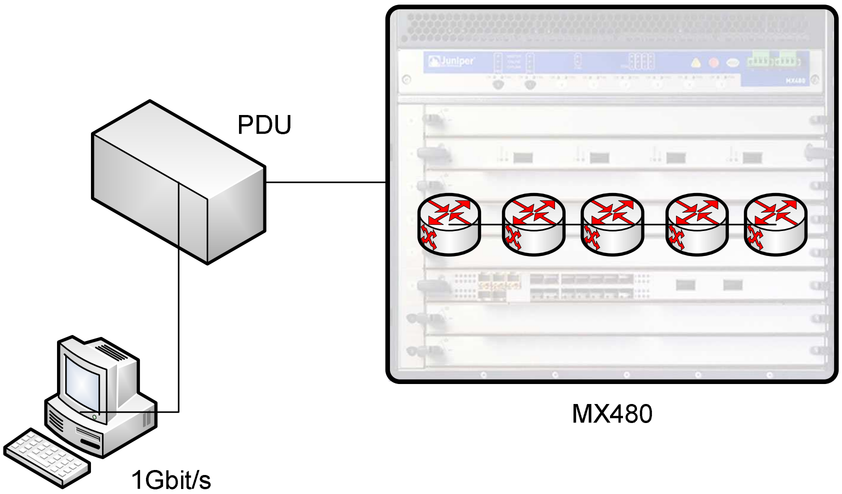

Fig 3.34 Data plane scenario with 1 Gbit/s + 5 logical systems

This experiment is designed to measure power use when the MX480 router has five logical systems configured and the Mgen generator creates 1 Gbit/s of UDP traffic.

The connections between the various ports are the same as shown in Figure 3.31.

3.5.9 Results

Figure 3.35 shows the current consumed by the two power supplies of the router:

Fig 3.35 Current consumption Scenario 4

During the time interval of 7 minutes, taking about 140 samples in intervals of 3 seconds, the average currents obtained are 2.14 A for the current outlet 4 and 2.43 A for the current outlet 3.

If we compare these results with those obtained in Scenario 3, there is a small variation between them. Current consumption decreases 10 mA and 20 mA respectively for the two power supplies.

Surprisingly, increasing both amount of traffic and the number of logical systems, current consumption slightly decreases. In fact, this consumption should have been higher, increasing these two parameters.

On the other hand, Figure 3.36 shows the apparent power consumed:

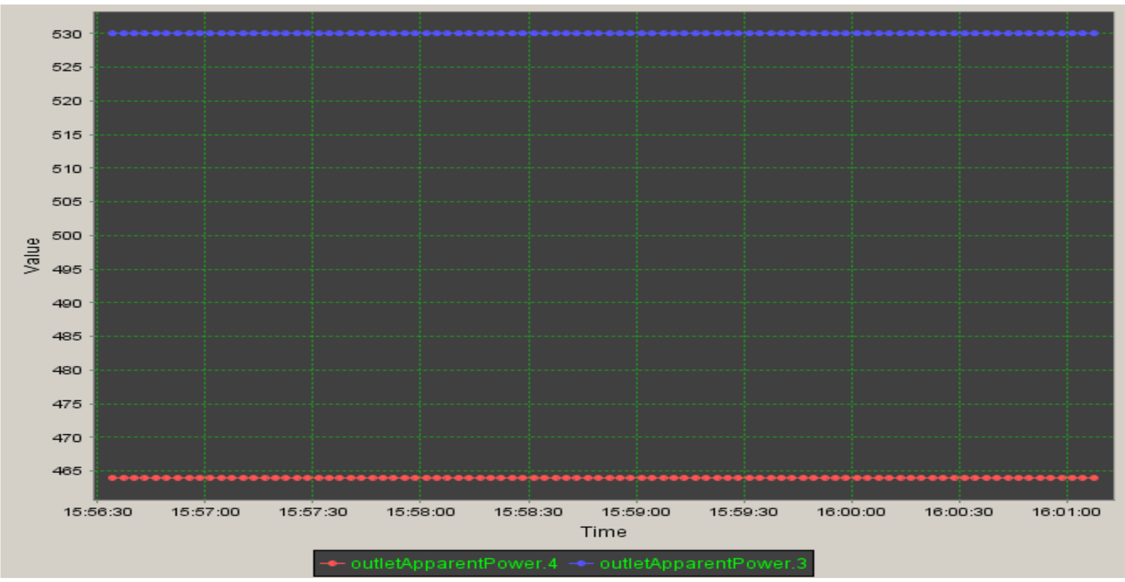

Fig 3.36 Apparent power consumption Scenario 4

Taking into account the value of the different samples during the interval of time established, the average apparent powers obtained are 530 VA and 464 VA for the respective power supplies.

In comparison with the results obtained in Scenario 3, there is no change between them. In this case, we could use the same argument as in current results mentioned above, but for the case of apparent power.

CHAPTER 4. CONCLUSIONS AND FUTURE DEVELOPMENTS

4.1 Conclusions

We now describe the conclusions obtained from the work carried out. This thesis has been divided into three main parts.

Firstly, an introduction to Green IT has been presented together with a series of activities and technologies that are currently being undertaken in order to reduce the current 2% of global CO_2 emissions created by IT, such as: Energy Star, renewable energy sources, virtualization, cloud model and energy efficiency metrics. These definitions were important in order to understand the sections that followed.

Then, the GSN Project has been described as a Green IT approach which includes many of the activities and technologies mentioned above. The main contribution of this thesis in the GSN Project was the installation of a solar-powered node in Riba Roja, a village near the province of Tarragona. As of September 1st 2011 the process is still underway because the optical link between Canada and the i2CAT node has not yet been established.

Other tasks performed were:

- Training on ISO 14064, which includes studying and transferring to other team members.
- Review of document such as the Carbon Protocol Draft, EUNIS abstract, and the GSN-Mantychore integration proposal.
- Presentation on the collaboration between Mantychore and GSN.
- Collaboration in the definition of the GSN Use Case.

Finally, the second focus of this project was the development of a general power model for routers with virtualization capabilities through a series of power consumption measures, in this case on the Juniper MX480 that we could not complete due to several reasons, including lack of resources such as fibers and computers.

The main idea was to carry out a series of power consumption measures in different scenarios, including full traffic and idle situations, and measure the power consumption of the MX480 router and its associated virtual routers in order to establish a general power model.
The conclusions that we have obtained from the results are:

1. The chassis of the i2CAT's MX480 router consumes nearly 50% of the total power required by the device, a similar result to that obtained by Chabarek and Barford in 2008 [7] with similar carrier–class routers. In addition to this, the other 50% is distributed among the four card modules (routing engines, Gigabit and 10Gigabit Ethernet cards).

2. We have tried to stress the router with 17 Gbit/s of maximum traffic and five logical systems (virtual routers), and we have only obtained an increase in consumption of 1%. This also coincides with the results of Barford [46], where both traffic and routing do not affect too much on consumption, compared to having the chassis or cards active. Note that the router is so powerful (1.4 Tbps power backplane) that we did not have enough resources to stress it in terms of traffic, such as fibers to configure more logical tunnels in order to connect more logical systems (up to 16), and computers to generate more traffic. These resources have not been able to get basically for lack of time and availability.

3. Another alternative was to use J-Flow [61] in order to obtain traffic flow statistics, and thus being able to increase CPU, memory and (hopefully) power consumption, but we had the drawback that the MX480 router did not have this license installed.

Although we did not succeed in the development of the power model, the work carried out provides some preliminary insight on the router characteristics, and we hope that the scenarios and testing designed in this project serve in the future to meet the ultimate goal of creating the power model in order to take this into account in routing optimization studies, load balancing or to make decisions on migration.

4.2 Future developments

The current and power measurements carried out in this project were designed to obtain a general power model for routers with virtualization. This model could not be achieved by lack of resources, but the proposed measures remain open for their development in a future. The project can be extended in several ways.

The future measures should be done with more fibers and computers in order to stress the MX480 router, and thus get better results. Other related ideas are as follows:

- Acquire the Tunnel Services PIC module in order to set up logical tunnels with the aim of connecting up to 16 logical systems.

- Obtain the J-Flow license to have statistics on the traffic flowing through the router, and thus being able to increase CPU, memory and (possibly) power consumption.

- Replicate traffic with the help of switches in order to obtain a greater amount of it, and thus increase load.

- Generating broadcast traffic within the router, and route it across the logical systems.

On the other hand, it is also expected to finish by September 2012 the installation of the node powered by solar panels here in Catalonia, and the connection between it and Canada, as an example of green datacenter.

GLOSSARY

CCC	Circuit Cross-Connects
DCiE	Data Center Infrastructure Efficiency
DPC	Dense Port Concentrator
FPC	Flexible PIC Concentrator
GHG	Greenhouse Gases
GSN	GreenStar Network
IaaS	Infrastructure as a Service
ICT	Information and Communication Technology
IP	Internet Protocol
KVM	Keyboard-Video-Mouse
LP	Light Path
MPC	Modular Port Concentrator
MPLS	Multiprotocol Label Switching
NREN	National Research and Education Networks
PaaS	Platform as a Service
PDU	Power Distribution Unit
PUE	Power Usage Effectiveness
SaaS	Software as a Service
SCB	Switch Control Board
SFP	Small Form-factor Pluggable
SNMP	Simple Network Management Protocol
SPE	Special Purpose Equipment
SPS	Solar Powered System
UPS	Uninterruptible Power System
WPS	Wind Powered System

BIBLIOGRAPHY

[1] Christy Pettey, Gartner Estimates ICT Industry Accounts for 2 Percent of Global CO_2 Emissions, April 2007: http://www.gartner.com/it/page.jsp?id=503867

[2] Bill St. Arnaud blog: http://billstarnaud.blogspot.com/

[3] Bill St. Arnaud, Green IT/Broadband and Cyber-Infrastructure, March 2011: http://green-broadband.blogspot.com/

[4] Energy Star homepage: http://www.energystar.gov/index.cfm?c=about.ab_index

[5] GSN homepage: http://www.greenstarnetwork.com/

[6] ISO 14064: http://www.global-greenhouse-warming.com/ISO-14064.html

[7] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright. Power Awareness in Network Design and Routing. In IEEE INFOCOM, April 2008.

[8] David Cearly, Gartner's Top 10 Strategic Techniques for 2009, October 2008: http://blogs.gartner.com/david_cearley/2008/10/14/gartner%E2%80%99s-top-10-strategic-technologies-for-2009/

[9] VMware homepage: http://www.vmware.com/solutions/green-it/

[10] The Economist, Where the cloud meets the ground, October 2008: http://www.economist.com/node/12411920

[11] Hari Srinivas, What is Energy Footprint? http://www.gdrc.org/uem/footprints/energy-footprint.html

[12] Intel homepage: http://www.intel.com

[13] Jack Lo, VMware and CPU Virtualization Technology: http://download3.vmware.com/vmworld/2005/pac346.pdf

[14] VMware, Understanding Full Virtualization, Paravirtualization, and Hardware Assist, September 2007: http://www.vmware.com/files/pdf/VMware_paravirtualization.pdf

[15] VMware, Reduce Energy Costs and Go Green: http://www.vmware.com/files/pdf/VMware-GREEN-IT-OVERVIEW-SB-EN.pdf

[16] VMware, How VMware Virtualization Right-sizes IT Infrastructure to Reduce Power Consumption: http://www.vmware.com/files/pdf/WhitePaper_ReducePowerConsumption.pdf

[17] Holly Stevens, Gartner Says Agility Will Become the Primary Measure of Data Centre Excellence by 2012, October 2007: http://www.gartner.com/it/page.jsp?id=535714

[18] OpenNebula homepage: http://www.opennebula.org

[19] Cloud Computing homepage: http://www.cloudcomputing.org

[20] Google Apps homepage: http://www.google.com/apps/intl/en/business/index.html

[21] Google App Engine: http://code.google.com/appengine

[22] Project Caroline Home: http://research.sun.com/projects/caroline

[23] GigaSpaces homepage: http://www.gigaspaces.com

[24] Mantychore homepage: http://www.mantychore.eu/

[25] Amazon homepage: http://aws.amazon.com

[26] GoGrid homepage: http://www.gogrid.com

[27] Windows Azure: http://www.microsoft.com/azure

[28] The Green Grid homepage: http://www.thegreengrid.org

[29] Green Grid Metrics: Describing Datacenter Power Efficiency, February 2007: http://www.esolutions.com.co/index.php?view=article&catid=37%3Acpd&id=92%3Aindicadores-green-grid1&format=pdf&option=com_content&Itemid=89

[30] Cybera, GreenStar Network: http://www.cybera.ca/projects/greenstar-network

[31] GSN, Mobility Principle: http://www.greenstarnetwork.com/drupal6/node/5

[32] GSN, Fabric Computing: http://www.greenstarnetwork.com/fr/node/8

[33] IaaS Framework: http://www.iaasframework.com/

[34] XORP homepage: http://www.xorp.org/

[35] Quagga homepage: http://www.quagga.net/

[36] Xen homepage: http://xen.org/

[37] KVM homepage: http://www.linux-kvm.org/page/Main_Page

[38] Puppet labs: http://puppetlabs.com/

[39] VDE: http://vde.sourceforge.net/

[40] GSN, GSN Carbon Measurement Protocol: http://www.greenstarnetwork.com/node/18

[41] ISO 14064-2:2006, Greenhouse gases - Part 2: Specification with guidance at the project level for quantification, monitoring and reporting of greenhouse gas emission reductions or removal enhancements.

[42] GSN, Technology: http://www.greenstarnetwork.com/node/6

[43] ISO 14064-1:2006, Greenhouse gases - Part 1: Specification with guidance at the organization level for quantification and reporting of greenhouse gas emissions and removals.

[44] ISO 14064-3:2006, Greenhouse gases - Part 3: Specification with guidance for the validation and verification of greenhouse gas assertions.

[45] EUNIS homepage: http://www.eunis.org/

[46] GSN Use Case: http://jira.i2cat.net:8090/display/MANTECH/GSN+Use+Case

[47] MX480 Hardware Guide: http://www.juniper.net/techpubs/en_US/release-independent/junos/information-products/topic-collections/hardware/mx-series/mx480/hwguide/mx480-hwguide.pdf

[48] MX480 Switch Control Board (SCB) Description: http://www.juniper.net/techpubs/en_US/release-independent/junos/topics/concept/scb-mx480-description.html

[49] Modular port concentrators for the MX Series: http://www.juniper.net/us/en/local/pdf/datasheets/1000294-en.pdf

[50] Multiservices DPC: http://www.juniper.net/us/en/local/pdf/datasheets/1000258-en.pdf

[51] Raritan PDU Online Help: http://web1.raritan.adxstudio.com/help/px/v1.4.1/en/#18615

[52] Show chassis power command: http://www.juniper.net/techpubs/en_US/junos/topics/reference/command-summary/show-chassis-power.html

[53] MX480 AC power supply description: http://www.juniper.net/techpubs/en_US/release-independent/junos/topics/concept/power-supply-mx480-ac.html

[54] Agilent N2X technical datasheet: http://www.ixiacom.com/pdfs/ixn2x/legacy_products/5988-9337EN.pdf

[55] Mgen User's Guide: http://www-hera-b.desy.de/subgroup/network/mgen/UserGuide.html

[56] SMC8824M switch datasheet:
http://www.smc.com/files/AT/ds_8824M_8848M.pdf

[57] Wireshark homepage: http://www.wireshark.org/

[58] Logical System Configuration Overview:
http://www.juniper.net/techpubs/en_US/junos10.4/topics/concept/junos-logical-interfaces-logical-system-configuration-overview.html

[59] Configuring Logical Tunnel Interfaces:
http://www.juniper.net/techpubs/en_US/junos9.5/information-products/topic-collections/config-guide-services/services-configuring-logical-tunnel-interfaces.html

[60] Tunnel Services PIC:
http://www.juniper.net/us/en/local/pdf/datasheets/1000092-en.pdf

[61] J-Flow description: http://www.hardware.com/store/juniper/JX-JFLOW-LTU

[62] GSN Node Description and Requirements, GSN Infrastructure Working Group (CRC), September 2010.

APPENDIX A. Configuring the Raritan PDU

This appendix is a summary of [51].

To configure the Raritan PDU, you will require the serial adaptor that is provided when you purchase the Raritan PDU. Using a terminal emulator application like HyperTerminal from Microsoft Windows, you can connect to the Raritan PDU to configure the IP address for remote access and monitoring of the PDU. Set the terminal setting to the following:

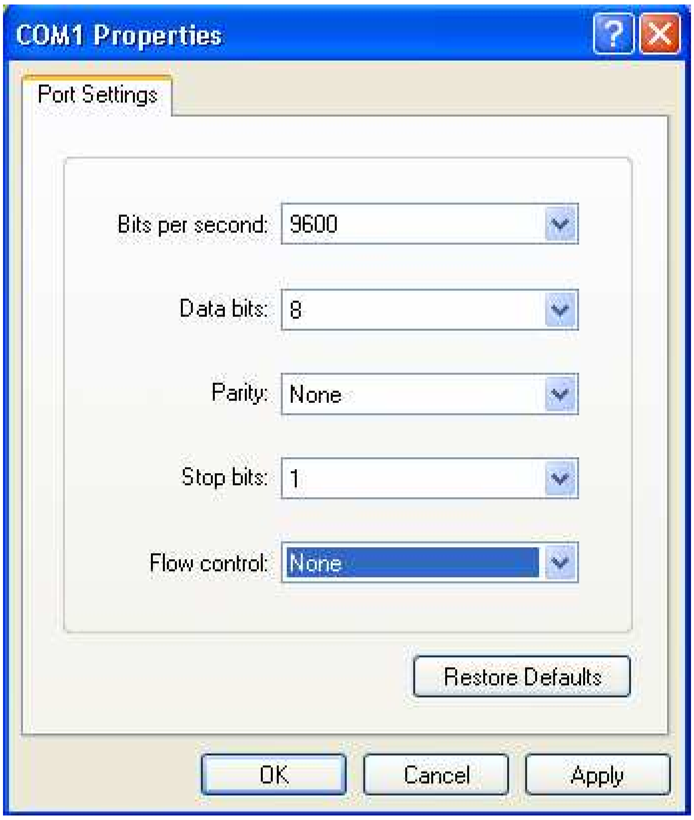

When you are connected to the Raritan PDU via the serial interface, you will see the following CLI commands.

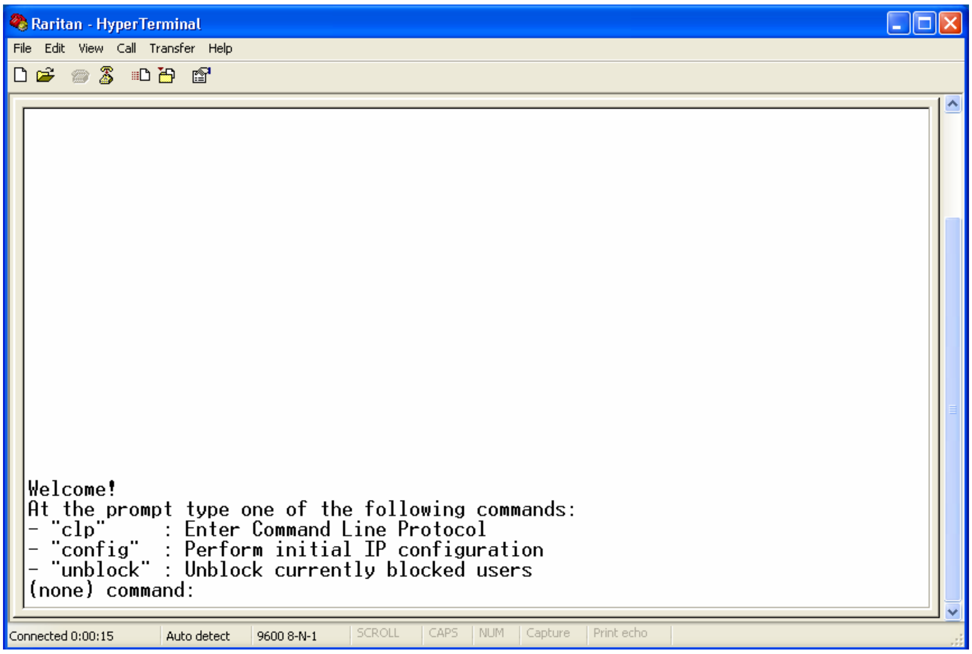

Type "config". This will allow you to configure the networking parameters for the PDU.

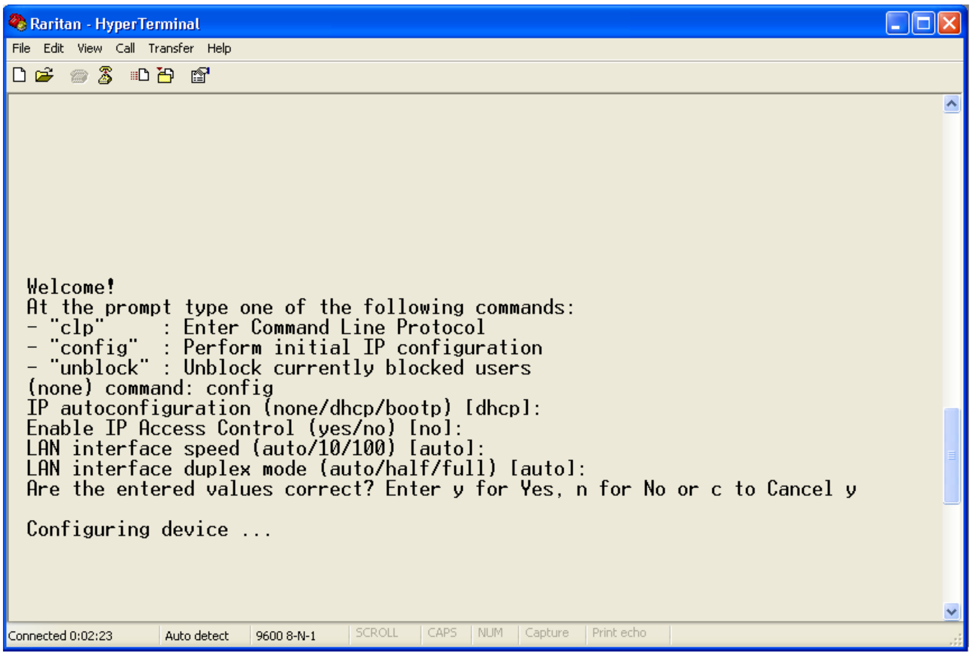

Enter "none" to manually configure the networking parameters. Enable IP Access Control to "yes" to allow access via the web, telnet or SSH. Save configuration will upload the new parameters to the Raritan PDU. Connect the LAN port on the PDU to the network switch to allow remote access of the PDU.

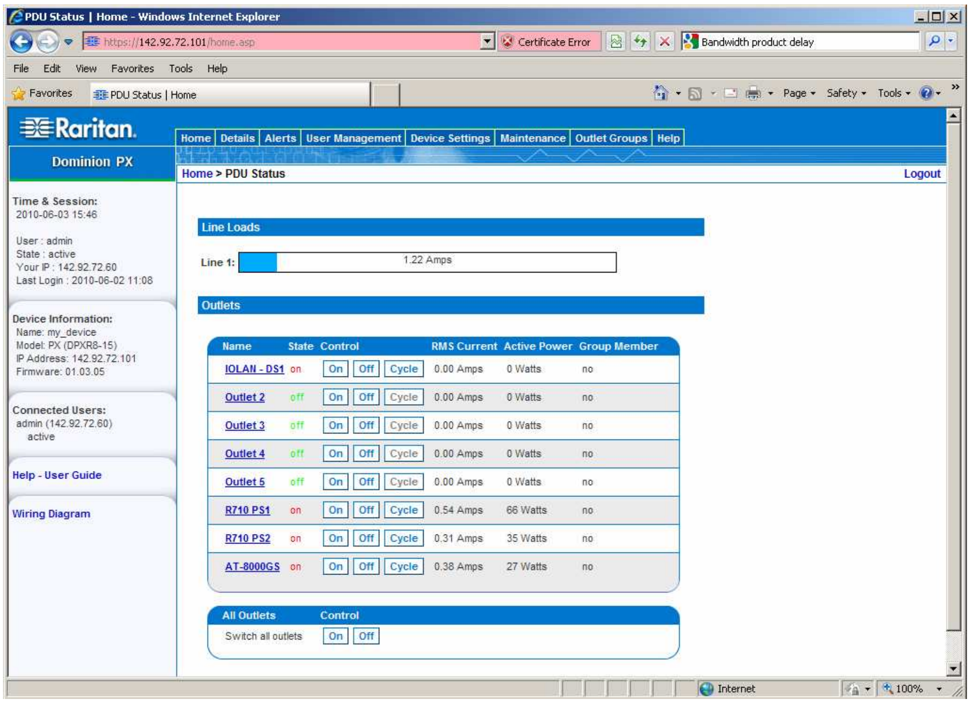

The picture above is a screen shot showing the HTTP interface of the Raritan PDU.

APPENDIX B. Configuring the SNMP Settings

This appendix is a summary of [51].

You can enable or disable SNMP communication between an SNMP manager and the PDU.

To configure the SNMP communication:

1. Choose Device Settings > SNMP Settings. The SNMP Settings page opens.

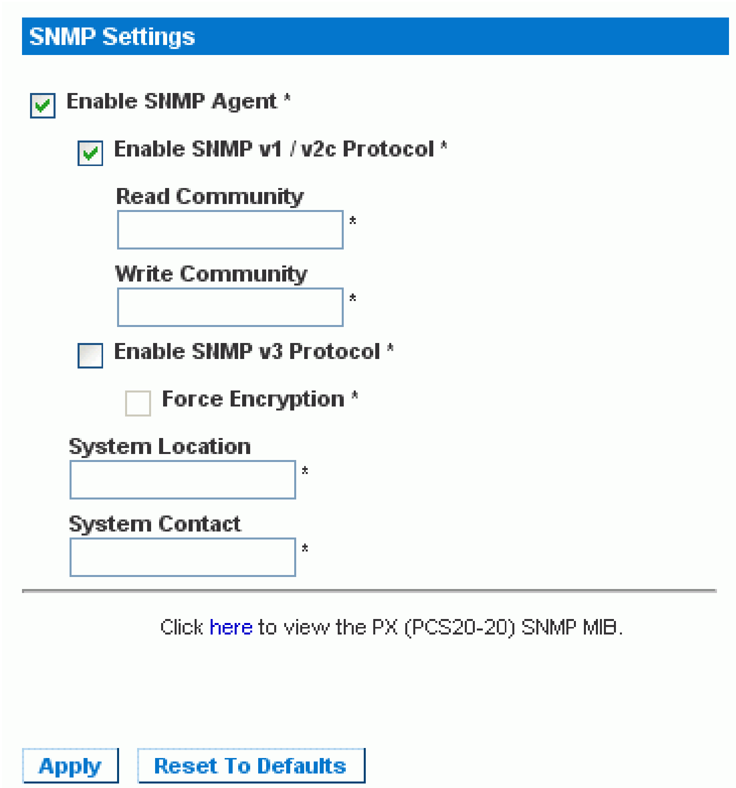

2. Select the Enable SNMP Agent checkbox to enable PDU to communicate with external SNMP managers. A number of options become available.

3. Select the Enable SNMP v1 / v2c Protocol checkbox to enable communication with an SNMP manager using SNMP v1 or v2c protocol. Type the SNMP read-only community string in the Read Community field and the read/write community string in the Write Community field.

4. Select the Enable SNMP v3 Protocol checkbox to enable communication with an SNMP manager using SNMP v3 protocol. Additionally, select the Force Encryption checkbox to force using encrypted SNMP communication.

5. Type the SNMP MIBII sysLocation value in the System Location field.

6. Type the SNMP MIBII sysContact value in the System Contact field.

7. Click on the link at the bottom of the page to download an SNMP MIB for your PDU to use with your SNMP manager.

8. Click Apply. The SNMP configuration is set.

APPENDIX C. Enabling Data Retrieval

This appendix is a summary of [51].

The data retrieval feature allows the retrieval of PDU data by an SNMP manager, such as the data of outlet, line, and circuit breaker. When enabled, PDU measures all sensor data at regular intervals and stores these data samples for access over SNMP.

PDU stores up to the last 120 measurements (samples) in the data log buffer.

Configuring the delay between samples adjusts how often the sample measurements are made and stored for retrieval. The default delay is 300 seconds. Delays must be entered as multiples of 3 seconds.

PDU's SNMP agent must be enabled for this feature to work. In addition, using an NTP time server ensures accurately time-stamped measurements.

To configure the data sample delay:

1. Choose Device Settings > PDU Setup. The PDU Setup page opens.

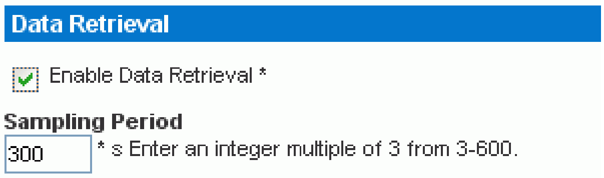

2. By default, Data Retrieval is disabled. Select the Enable Data Retrieval checkbox, and the Sampling Period field becomes configurable.

3. Type a number in the Sampling Period field, indicating how often (in seconds) PDU stores data samples. Values in this field are restricted to multiples of 3 seconds, ranging from 3 to 600 seconds (10 minutes).

4. When you finish, click Apply. The retrieved data samples are stored immediately once this feature is enabled and the delay between samples is configured.

After data retrieval is enabled, an external manager or application (such as Raritan's Power IQ) can access the stored data using SNMP. Download the Dominion PX MIB file to assist you in configuring third-party managers.

APPENDIX D. GSN Node Description and Requirements

This appendix is a summary of [62].

A GSN node should consist of the following Special Purpose Equipment (SPE):

- A Layer 2 (L2) Switch.
- A Server based on the Intel E5500 or Later Chip Set.
- A Power Distribution Unit (PDU).
- A Solar Powered System (SPS) or Wind Powered System (WPS).

The SPS should include solar panels, a charge controller, an inverter and battery banks. For the WPS, the solar panels would be replaced by a wind turbine. Note that remote control and monitoring is required for either an SPS or WPS.

A typical GSN node installation is depicted in Figure 1.1.

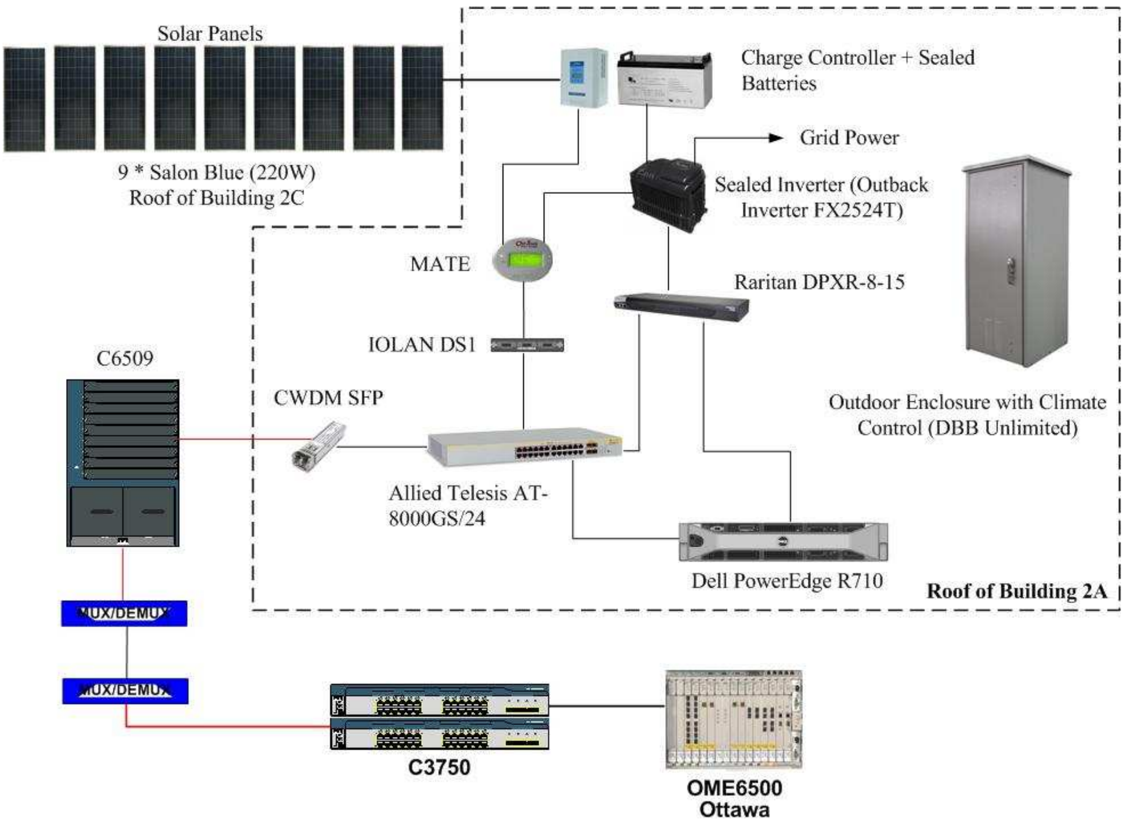

Fig. 1.1 CRC GSN Node [62]

A GSN node is expected to connect to the GreenStar Network using a 1 Gbps LightPath (LP) across CANARIE that would be terminated at the ÉTS Hub Node in Montréal using a 1 GbE interface. The preferred exchange point with CANARIE is the MANLAN in New York.

The Special Purpose Equipment (SPE) that constitutes a GSN Node includes an L2 Switch to connect the equipment to it and to terminate the LP onto it.

Due to the budget constraints within the GSN project, the recommended L2 Switch is the Allied Telesis AT-8000GS/24 model.

The required optics for the Small Form-factor Pluggable (SFP) transceiver will depend on the distance and type of fiber used to establish the connection between the L2 Switch and the network end point. CWDM equipment is used to the matching.

For ÉTS, a second and more powerful L2 switch, an Arista Network 7124S L2/L3 switch, is required to terminate all the LPs associated with the remote GSN nodes.

APPENDIX E. Juniper J-series Measurements

Table 1.1 summarizes the results obtained for the different Juniper J-series routers and their corresponding scenarios:

	J2320 (current)	J2520 (current)	J2320 (ap. power)	J2520 (ap. power)
Idle state	270 mA	340 mA	57 W	75 W
1 Gbit/s incoming traffic	391 mA	467 mA	85 W	102 W

Table 1.1 Juniper J-series energy consumption

In conclusion, energy consumption has increased by 14% both in the case of the current and the apparent power.