



GreenDataNet

D3.10 – Summary of Virtualization of IT tasks

Final

EATON

Rev 0.1

EATON: Aurelien Begou, Maria Laura Corallini, Emilien Kia, Christian Nerriere

2. CONTENTS

REVISION SHEET	3
KEY REFERENCES AND SUPPORTING DOCUMENTATIONS	4
3. INTRODUCTION	5
3.1 Document Purpose.....	6
3.2 Definition, acronyms and abbreviations	7
3.2.1 Key Acronyms and Abbreviations.....	7
3.3 Document overview	8
4. RACK CONTROLLER: SOFTWARE PLATFORM.....	9
4.1 Communication model of the legacy system	9
4.2 Control Algorithm.....	9
5. RACK CONTROLLER: GREENDATANET DEVELOPMENT	11
5.1 Communication model of the Global GDN System	11
5.2 Architecture of the vCenter connector	12
5.2.1 Architecture.....	12
5.2.2 Configuration	13
5.2.3 Restful API	13
6. EXEMPLE OF EXECUTION OF A VIRTUAL MACHINE MOVE AND CONCLUSIONS.....	14
6.1 Virtual Machine move	14
6.2 Conclusions	16

REVISION SHEET

Revision Number	Date	Brief summary of changes
Rev 0.1	7/11/2015	Baseline document
Rev0.2	2/2/2016	EC comments on non vCenter hypervisors

KEY REFERENCES AND SUPPORTING DOCUMENTATIONS

- [1] VMware, vCenter Server Website.
<http://www.vmware.com/uk/products/vcenter-server/>
- [2] Microsoft, SCVMM Website.
<https://technet.microsoft.com/uk/technet-techcenter-system-center-virtual-machine-manager.aspx>
- [3] Citrix, XenCenter Website.
<https://www.citrix.com/products.html>
- [4] C Clark, K Fraser & al., Live migration of virtual machines, Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation, 2005.
- [5] M Nelson, BH Lim & al., Fast Transparent Migration for Virtual Machines, USENIX '05 Paper, 2005.
- [6] **D2.3, GreenDataNet project.** *D2.3 Server Multi-Level Software Management.* s.l. : Seventh Framework Program, 2015.
- [7] VMware Sphere5 Survey, InformationWeek Website, August 2011
<http://www.informationweek.com/mobile/vmware-pricing-controversy-exclusive-user-research/d/d-id/1099595>
- [8] H Tardy, Power Management for Server Virtualization, EATON White Paper, April 2012
- [9] H Estrada, S Almuabe, M Farid, Survival Guide of Data Centre Shutdown, IBM Blog, July 2013
https://www.ibm.com/developerworks/community/blogs/0587adbc-8477-431f-8c68-9226adea11ed/entry/survival_guide_for_a_data_center_power_shutdown?lang=en
- [10] REST API. R Fielding , Architectural styles and the Design of Network-based Software Architectures, PhD Dissertation, 2000
http://www.ics.uci.edu/~fielding/pubs/dissertation/rest_arch_style.htm
- [11] Common information mode DMTF Website.
<http://www.dmtf.org/standards/cim>

3. INTRODUCTION

One of the main goals of the GreenDataNet project is to investigate and develop new technologies to decrease the power consumption of urban Data Centres. *Deliverable 2.3 – Server Multi-Level Software Management* provided a first step on the achievement of this goal by delivering a first solution, composed by a Rack Controller and a Server Manager, able to monitor the power elements of the Data Centre and to suggest the corrective actions to be done in order to decrease the IT power consumption. The corrective actions to be applied on the servers consist in reconfiguring or shutting off the Virtual Machines, moving them to another part of the IT infrastructure or changing their operational frequency. In order to do that, the solution needs to be able to interact with a virtualisation management application.

The virtualisation management applications (such as vCenter by VMware, SCVMM by Microsoft, XenCenter by Citrix) allow to manage the problems issued by the virtualized infrastructure, such as processors saturation, corruption of the hard drive or network saturation [1][2][3]. In particular, they allow to set up processes to prevent or to fix those problems by applying actions as the ones proposed by the Server Manager [4][5][6]. To fulfill this task, Eaton has adapted an existing VM management solution developed internally to be integrated with the Server Manager tool developed by EPFL. The Software developed by Eaton aims to manage the virtualization infrastructure by taking into account the power requests.

The legacy Eaton Software has the capacity to manage the problems related to the bad quality of the power supply, by providing an automated process that alert and eventually switch off the IT equipment to protect the IT infrastructure in case of power losses and poor quality of the supply current. The software developed for GreenDataNet focuses on the optimization of the IT power consumption on servers, in particular, it integrates the functionalities needed to interact with the virtualisation management application, to collect IT data and to control the move of virtual Machines. The existing algorithm implemented in the Rack Controller corresponds to a reactive approach: once the power failure is detected, IT equipment is softly powered off, while the GreenDataNet approach is pro-active: the system is continuously supervised to optimize the overall power consumption.

3.1 DOCUMENT PURPOSE

This document is a summary of deliverable *D3.9 – Virtualisation of IT tasks*, that explains how the interaction between the Rack Controller and the virtualisation management application is performed, in particular, the way the Rack controller collects the server performance data from the virtualisation management application and how the VM move control is deployed.

As the time for development was limited, it has been decided to focus on one specific virtual management application, and vSphere has been chosen as it is one of the most popular virtualization platform. More specifically, the SW has been developed to be used with VMware **vSphere** 5.5.

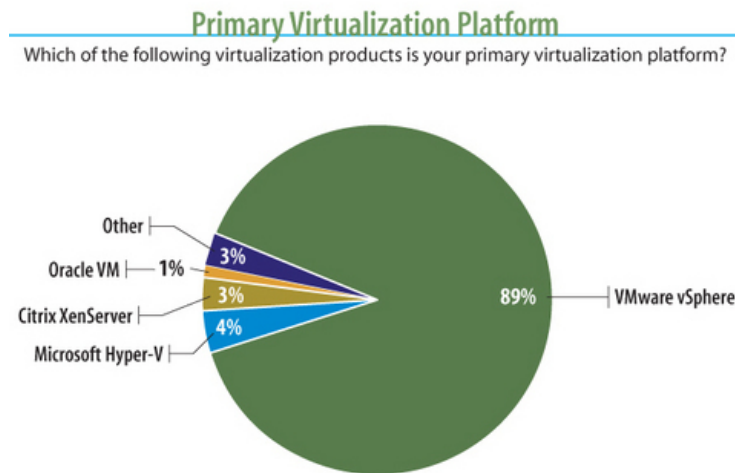


Figure 3-1 Market share of main virtualization products [7]

This deliverable provides an overview of how the software developed within the GreenDataNet scope communicates and controls the virtualisation manager, without specifying how the Server Manager algorithm is working.

3.2 DEFINITION, ACRONYMS AND ABBREVIATIONS

3.2.1 KEY ACRONYMS AND ABBREVIATIONS

AEMS	Aggregated Energy Management System
AMQP	Advanced Message Queuing Protocol Standard
BMS	Battery Management System
CIM	Common Information Model
CPU	Central Processing Unit
DCIM	Data Centre Infrastructure Management
DOW	Description of Work
EC	European Commission
EV	Electric Vehicle
EXSi	Hypervisor
GC	GreenDataNet Controller
ICT	Information and Communication Technology
IPR	Intellectual Property Rights
IT	Information Technology
ITEE	IT equipment Energy Efficiency
KPI	Key Performance Indicator
LEED	Leadership in Energy and Environmental Design
MoM	Message oriented Middleware
NIC	Network Interface Controller
NoC	Network-on-Chip
NUT	Network UPS Tools
OOB	Out Of Band (control channel)
PCC	Project Coordination Committee
PDU	Power Distribution Unit
ePDU	Rack Power Distribution Unit
PSB	Project Steering Board
PUE	Power Usage Effectiveness
RC	Rack Controller, synonym of GC – GreenDataNet Controller
SDK	Software Development Kit
SEMS	Smart Energy Management System
SLA	Service Level Agreement
SME	Small or Medium size Enterprise
SNMP	Simple Network Management Protocol
SOA	Service Oriented Architecture
SoC	System-on-Chip
SW	Software
UPS	Uninterruptible Power Supply
VM	Virtual Machine
WP	Work Package

3.3 DOCUMENT OVERVIEW

The document is divided into four main sections, including the introduction. Section 4 provides an understanding of how the algorithms of the EATON platform works, Section 5 describes the modifications applied to the SW platform to collect the right data for the Server Manager algorithm and apply the VMs move. Section 6 provides an example of VM move.

4. RACK CONTROLLER: SOFTWARE PLATFORM

As explained in the description of work, an existing Software platform have been used by EATON to develop the IT virtualization tool necessary to integrate the IT power optimisation algorithm described in deliverable *D2.4 – Racks Multi-objective Energy Management Specification and Implementation*. It is important to describe the behaviour of the legacy system to highlight the improvement that has been provided through the GreenDataNet project.

4.1 COMMUNICATION MODEL OF THE LEGACY SYSTEM

Figure 3-1 shows the communication flows between the different parts of the Software. The Rack controller is composed by a Power controller module, which is in charge of monitoring the Uninterruptible Power Supply, and a vCenter connector, whose function is to push orders of graceful shut down of the Virtual Machines or of the servers when the battery capacity level is below a certain threshold configured by the Data Centre manager.




Figure 3-4-1 - Communication model overview

4.2 CONTROL ALGORITHM

In a first step, after that the vCenter connector has recognized the list of the VMs and application running in the Data Centre, the user is asked to prioritize the application according to their mission critical level, this means that the user is asked to identify the activities that can be interrupted during a power outage without impacting the business. Based on the workload priorities, the Data Centre manager would define the business continuity policies [8].

Generally, In term of priority, the customer will proceed in this order to shutdown [9]:

1. Applications
2. Servers
3. Storage
4. Network

After that, there may be some high priority applications, but the rest of the applications, which are not mission critical, could be gracefully shut down in the event of a power failure. Doing so can enable the user to consolidate workloads further, reducing power draw and extending battery runtime, giving more minutes to restore normal operations. The fact is that, reducing the power load by half, effectively extends the runtime by twice. Those additional minutes may be the difference between a minor power or environmental event or major down time.

The Power controller module of the Rack controller collects the information regarding the status of the UPS and the Power data such as the Battery capacity. The power data are compared with the thresholds set by the user and if those values are within the critical range defined, a specific event notification is generated by the Power controller and sent to the vCenter connector. The event can be of different levels of criticality, from a simple notification to a business critical alarm.

In the case of the most critical alarms, the vCenter connector would react automatically, by sending to the virtualisation management application an order to proceed with the graceful shut off of the non-business critical Virtual Machines, following the prioritization list provided by the user.

5. RACK CONTROLLER: GREENDATANET DEVELOPMENT

5.1 COMMUNICATION MODEL OF THE GLOBAL GDN SYSTEM

In deliverable *D2.3 – Server Multi-Level Software Management*, the IT power optimization algorithm, integrated in the Server Manager SW developed by EPFL has been applied to real power data coming from a more complete version of the power controller. The power controller is able to monitor the UPS, each plug of the Power Distribution units and some additional environmental sensors.

The vCenter connector is able to interact with the virtualization management application, by collecting information regarding the footprint of the Virtual Machines and the frequency of the exchanges between them. Moreover the vCenter connector is able to order the VM moves from one server to another one, in addition to the existing capacity to control the graceful shutdown of the servers, the VMs or the applications. This data are collected with a timestamp of 1 minute and they are provided to the Server manager once in an hour: this means that each hour the Rack controller provide the historical data of the previous hour with a granularity of 1 minute.

The Server Manager receive this information, then it runs the optimization algorithm and it sends back to the vCenter connector the list of the virtual machines to be moved. The vCenter connector orders to the virtualization management application to execute the VMs move.

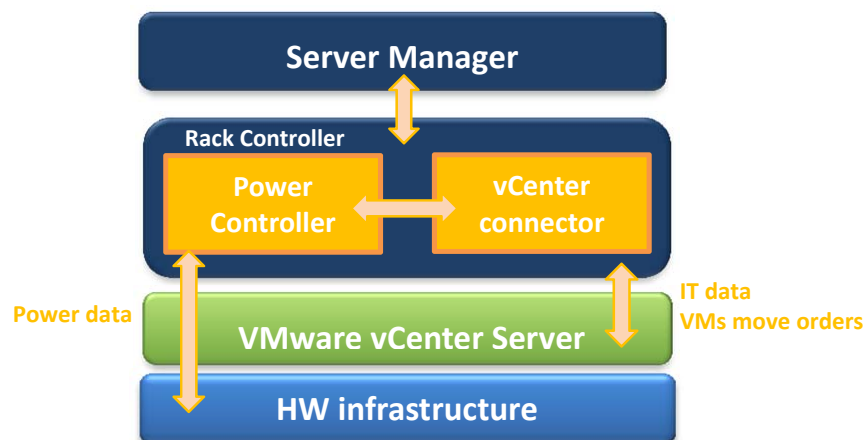


Figure 5-1 - Communication model of the global solution overview

The table 4-1 list the monitored data necessary to have the optimization algorithm work.

Monitored Power data	Monitored IT data
Power consumed by each server [W]	CPU allocation per virtual machine per server [Hz]
	Data communication between VMs [MBs]
	Disk size of the VMs (Total provisioned one) [GB]
	Time when the VM has been started
	VM status

Table 5-1 List of monitored Data

5.2.1 ARCHITECTURE

The Figure 4-2 describes the architecture of the vCenter connector. It composed by an Element manager, which is a sort of “dictionary” that collects the data sent by the virtualization provider and organize them in a way which is easily understandable and reachable from the outside: the information can be required by an external software via a REST API[10].

Element manager is the base mechanism to manage all the elements of the Data Centre, their features and functionalities, their properties and their links, its design have been influenced by DMTF CIM [11], so it should be easy to be mapped in an CIM end-point. The goal of this is to represent **elements** in a coherent, unified and structured way. **Elements** provides **functionalities** by implementing **facets** and having **links** to other elements.

The Virtualization provider is the software layer which interacts with the vCenter, it is able to send the VM moves orders and collects the monitoring data from the servers. The virtualization provider implements the Provider interface of the ElementManager bundle. When starting, the virtualization provider reads a configuration file where virtualization server properties are described. This configuration file is managed by bundle "common-properties". Connections and data from virtualization servers are managed by the Virtualization bundle.

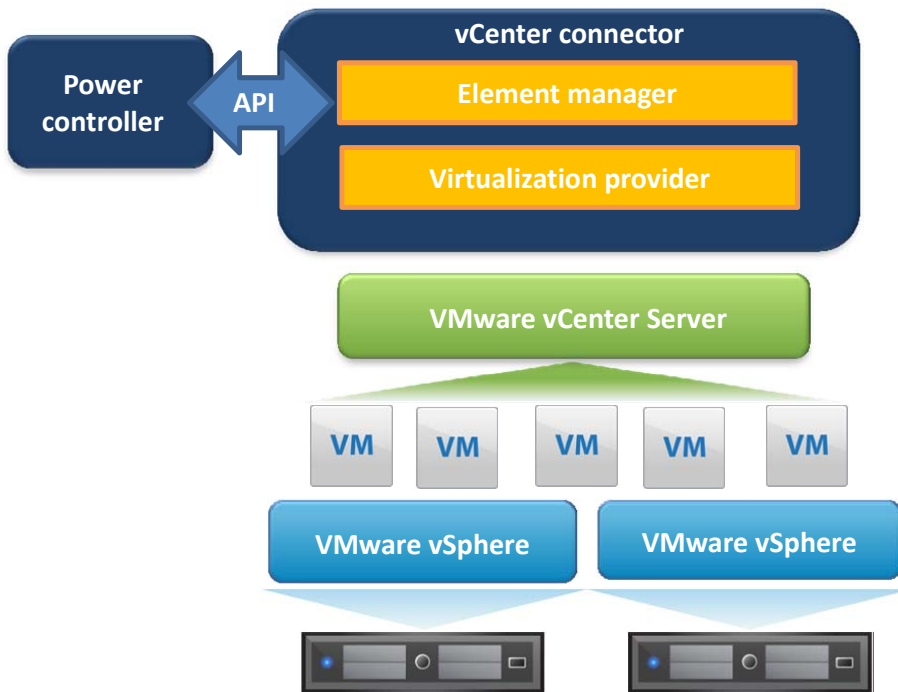


Figure 5-2 Architecture of the vCenter connector

5.2.2 CONFIGURATION

To connect to one or more VMware vCenters, the Virtualization provider needs a configuration file where are described the vCenter hostname and credentials. This file must contain a "vserver.list" property followed by a comma-separated list of vCenter names. For each name, a couple of properties must be defined to describe hostname and credentials used to connect:

```
vserver.list=vcenter04  
vcenter04.address=vcenter04.mbt.lab.etn.com  
vcenter04.port=443  
vcenter04.username=admin  
vcenter04.password=admin
```

In this case, one virtualization server is defined and named "vcenter04", its address is "vcenter04.mbt.lab.etn.com", its connection port is 443 (SSL) and the Virtualization provider can use "admin/admin" as credentials.

5.2.3 RESTFUL API

The VP (Virtualization Provider) Restful API relies on the ElementManager API.

For each virtualization manager configured, the VP creates a new related element, then it tries to connect to the configured virtualization manager and starts listening to virtualization bundle events. When a virtualization manager discovers a hypervisor or a virtual machine, the VP is notified and creates the equivalent element, set some of their properties and register it to the ElementManager.

6. EXAMPLE OF EXECUTION OF A VIRTUAL MACHINE MOVE AND CONCLUSIONS

6.1 VIRTUAL MACHINE MOVE

The aim of the following example is to show that the vCenter connector whose structure has been described in the previous chapters is working properly.

In the following example Virtual Machine *vm11 (vm-011)* located in server *esxi09 (host-11)* will be moved to server *esxi10(host-12)*.

Step 1: localization of the VM to move

```
GET http://127.0.0.1:8181/rest/etnrs/0.2/links/vm-011-hostlnk
{
  "dependent": {
    "instanceID": "vm-011",
    "ID": "vm-011",
    "type": "virtualization.host",
    "interfaces": [
      "logical.service",
      "logical.enabled",
      "logical.element",
      "virtualization.host",
      "core.element"
    ]
  },
  "antecedent": {
    "instanceID": "host-11",
    "ID": "host-11",
    "type": "virtualization.machine",
    "interfaces": [
      "virtualization.machine",
      "logical.enabled",
      "logical.element",
      "core.element",
      "logical.device"
    ]
  },
  "linkID": " vm-011-hostlnk ",
  "ID": "vm-011-hostlnk",
  "type": "virtualization.hosts",
  "interfaces": [
    "core.hosted",
    "virtualization.hosts",
    "core.dependency"
  ]
}
```

Step2: Move vm11 from esxi09 to esxi10

```
POST http://127.0.0.1:8181/rest/etnrs/0.2/instances/vm-011/actions/move
Payload ["host-12"]
Content-type application/json
Return : "OK"
```

Step3: Check VM position after move

GET <http://127.0.0.1:8181/rest/etnrs/0.2/links/vm-011-hostlnk>

```
{
  "dependent": {
    "instanceID": "vm-011",
    "ID": "vm-011",
    "type": "virtualization.host",
    "interfaces": [
      "logical.service",
      "logical.enabled",
      "logical.element",
      "virtualization.host",
      "core.element"
    ]
  },
  "antecedent": {
    "instanceID": "host-12",
    "ID": "host-12",
    "type": "virtualization.machine",
    "interfaces": [
      "virtualization.machine",
      "logical.enabled",
      "logical.element",
      "core.element",
      "logical.device"
    ]
  },
  "linkID": " vm-011-hostlnk ",
  "ID": "vm-011-hostlnk",
  "type": "virtualization.hosts",
  "interfaces": [
    "core.hosted",
    "virtualization.hosts",
    "core.dependency"
  ]
}
```

6.2 CONCLUSIONS

This deliverable describes the structure and the organization of the vCenter connector, the tool developed by EATON whose functionality is to interact with virtualization management applications. This tool allows to automatize the control procedure to move Virtual Machine, to switch ON and OFF the servers and, more in general to monitor the state of health of the virtualized system. This tool is part of a wider SW called Rack Controller and it pushes data related to the Virtualized Infrastructure performances to the Server Manager, the tool running the IT power consumption optimization algorithm.

This deliverable focuses only on the description of the vCenter connector, it is in deliverable D2.4 that the interaction between the Server Manager and the Rack Controller is explained.

In the last part of the document, an example of Virtual Machine move is provided, to show to the reader a practical example of how the software work.

The connector has been developed to interact with vCenter because, as explained in Paragraph 3.1 VMware is one of the main actors in the market and also the Hypervisor best known inside the consortium, but the Software is developed in a modular way, so that it would be easy to change the connector module (viJava library), to be able to interact with other Hypervisors such as Citrix or Hyper-V.

