# Teaching Penetration Testing and Malware Analysis within a Cloud-based Environment

Prof William J Buchanan (1), Bruce Ramsay, Richard Macfarlane, Adrian Smales, Dr Gordon Russell, Bill Buchanan (Jr)
The Cyber Academy, Edinburgh Napier University
Edinburgh, UK
(1) w.buchanan@napier.ac.uk

Eamonn Keane (2), Cormac Callahan (3), Borka Jerman Blazic (4), Oliver Popov (5)
Police Scotland, (2), Aconite Internet Solutions, Ireland (3)
JSI, Slovenia (4) University of Stockholm, Sweden (5)

***Abstract*—This paper outlines the design, implementation and evaluation of a private Cloud-based system over two semesters at Edinburgh Napier University for three modules. Overall, over the two semesters, these modules supported over 200 students within an isolated Cloud-based environment for penetration testing and malware analysis. The infrastructure allowed for a wide range of server and desktop instances to be used in a wide variety of network configuration, including for the latest version of Linux Kali and for Microsoft Windows server integration. The Cloud supported both on-campus students and distance learning ones. The main conclusions are that the Cloud coped well, but struggled within a few hours of a hand-in for a Penetration Testing coursework, where the overall utilizatiation was 25%. The issue was traced to the VMware vCenter controller which ran out of resources, and thus caused a slow down.**

***Index Terms*—Cyber Security, VMware vSphere 5.5**

## I. INTRODUCTION

There is an increasing demand for computer security and networking students with hands-on experience of a range of systems including for penetration testing and offensive security methods. Unfortunately many of these techniques cannot be taught on public clouds, such as for Amazon EC2 and Microsoft Azure, as it would breach contractual arrangements. The only solution is thus to use private Clouds which are isolated from the Internet and from corporate networks. This paper outlines the usage of a private cloud within Edinburgh Napier University as part of the EU-funded DFET (Digital Forensics Evaluation and Training) research project, and which has been used to teach advanced methods in computer security and networking.

The DFET Cloud creates new training methods/techniques to support judicial authorities, law enforcement agencies and associated stakeholders in the fight against cybercrime. It has development of a virtual (cloud-based) cybercrime training environment to include real life simulation and scenario analysis [5]. Currently it is partly funded by DG Home – Prevention Of and Fight against Crime, and aims to improve crime detection rates by providing scenario-based training in line with the dynamic nature of cybercrime.

Overall DFET aims to create a training infrastructure which can share cyber training across Europe, and allow access to hands-on environments, no matter the physical location of trainer. Its core partners are Edinburgh Napier University, Joseph Stefan Institute (JSI), Stockholm University, Police Scotland, and Aconite Internet Solutions.

The modules which used the DFET Cloud for 2014/2015 CSN09112 (Network Security and Cryptography) [1], CSN10107 (Security Testing) [2] and CSN11123/4 (Advanced Cloud and Network Forensics) [3]. These modules used a Cloud-based environment for every lab, including the usage of Linux Kali, Windows XP/20032008, Metasploit, Metasploitable and Web assessment systems.

For the evaluation, the paper focuses on the performance analysis of a large-scale penetration testing coursework. The results highlights a key breakpoint in the Cloud which occurred when over 80 students where completing their Penetration Testing coursework and working for a 12pm deadline, and where the resources within the infrastructure became stressed where there was a considerably lag in remotely accessing the DFET Cloud. Overall the Cloud coped by increasing memory allocation for vCenter, and in stopping running instances which had been dormant for several weeks.

## II. SECURITY TESTING MODULE CONTENT

The DFET project aims to integrate practical evaluation metrics within cyber security training and academic study, and to give credit for professional practice with academic study. As Figure 1 outlines professional practice often has four levels: leader; expert; professional; and associate. The challenge for academia is to map these onto academic levels of achievement which range through undergraduate and postgraduate study.

In a typical academic module the levels of achievement are often measured with **evaluate** at the top and **knowledge** as the lowest level. In computer security the practical elements of the work are also important to show the application of the technologies. Thus the phychomotor domain highlights areas such as **originality** and **adaptation** at the top and set and perception at the lowest level.

The key focus for the application of the DFET Cloud is for Penetration Testing and Malware Analysis was the CSN10107 [2], which runs at BEng (Hons) level. Table 1 outlines the range of lectures and labs used in the module. There were supported with full on-line lectures and labs that were run within the DFET Cloud environment.

The module aims to examine fundamental areas such as within password cracking, network protocols and SIEM (Security Incident and Event Management) and has two tests worth 25% and a Penetration Testing coursework worth 50%. In the tests students are provided with real-life network traces and logs for them to analyse.

The first two labs get students accustomed to the virtual environment and where the setup Windows and Linux servers to sit within a firewall environment with a private, DMZ and public network infrastructure (Figure 4). The end of these labs results in a fully working network infrastructure, and where students can firewall certain ports for different networks. This infrastructure is then used for the rest of the labs, where the students setup their environment each week from scratch with a new range of IP addresses. Overall this re-enforces the learning processes. See [2] for all the on-line material, including lectures, lab demos and associated material.

From discussions with local industry the module has been crafted with key skills around: network architectures, an in-depth understanding of network protocols; SIEM/IDS; command line tools; and vulnerability analysis. The team have found that these provide graduates with a wide range of skills which can be instantly applied into a commercial environment.

Table 1: Lecture and Lab schedule

|  | Lecture | Lab |
|---|---|---|
| 1 | Introduction | Vyatta Firewall Integration |
| 2 | Threats | Secure Architectures |
| 3 | Metasploit | Vulnerability Analysis |
| 4 | Network Forensics | Introduction to Metasploit |
| 5 | Advanced Network Forensics | Metasploit (Enumeration and Scanning) |
| 6 | SIEM |  |
| 7 | Malware Analysis | Web Attacks |
| 8 | Test 1 |  |
| 9 | DLP | Backdoors and Weak passwords |
| 10 | C/W Setup | SIEM/Splunk |
| 11 | IoT | Malware Detection |
| 12 | Steganography | Armitage |
|  | Test 2 | Coursework |

The modules run a range of subjects in each, and have differing requirements. There are two main aspects to the labs:

- **Labs which connect to the Internet.** In order to run servers within a firewall architecture the setup uses a VLAN which connects to the Internet through a pfSense firewall [4], or which can be directly connect to the Internet (in this case using VLAN 200 for a direct connection). With this hosts setup on different networks on the firewall can connect to the Internet, and provide their connectivity. For each VM, the gateway is set at 10.200.0.1/24 (Figure 4).
- **Isolated labs.** Several labs, including for Penetration Testing and malware analysis were run on isolated network, using VLAN which could not be routed onto the Internet (such as for VLAN 201, VLAN 202 and VLAN 203). The traffic on these networks will not reach the public address spaces.

## III. Pen Testing Coursework

The Penetration Testing coursework allowed for students to setup their own Pen Testing company and brand it. Overall 20% of the mark was allocated to the interaction that they had with the company (ApplesRUs) running the target Web server. This included sending penetration testing specifications and test timetables, along with corresponding with the company on a regular basis on the status of their tests. If the students crashed the target server which they were testing against, they were expected to inform the company (who would reset it back to its original state). Also if they were performing a major test, they were expected to get approval from the company before going ahead with it.

The marking schedule was structured so that the students could showcase both their academic and practical skills (with a 20-page limit, not counting appendices which often contained screen shots of the penetration test):

- **Research and Methodology** [20%]. This should show research into testing methodologies and tools for Web-based infrastructures, and in the use of virtualized infrastructures. It should also outline your own methodology for analysing the Web-based infrastructure, identifying the tools that will be used, and how the results will analysed.
- **Findings** [30%]. This should outline the main vulnerabilities found, and outline the methods that could be used to mitigate these weaknesses.
- **Conclusions** [20%]. This should reflect on the main findings, highlighting the key findings, evaluate the testing methods used, and critically appraise the infrastructure under test.
- **References/Presentation** [10%]. All references must be listed at the end of the report in APA/Harvard format, and all sources should be cited in the body report.

The key element of the business aspect of the Penetration Testing is then covered with:

- **Project Management** [20%]. The test management and communications that you have with the target system administrator will also be assessed, with a special focus on the nature of your communications and in how you have managed the project.

The company itself, ApplesRUs, was setup with an email address that students sent their emails, and this was received by the academics running the module, and where they took the role of the company contact (Figure 6). Each student was given their own Kali instance (and Windows 2008 with Nessus, if they required it), and they would then use this to test against the vulnerable server. In this case the ApplesRUs server was created as a Web site which was hosted on Windows 2003, and which had many intentional vulnerabilities. The key pieces of evidence that students needed to uncover was:

- Structure of infrastructure, with IP addresses, Operating Systems, MAC addresses, and so on.
- TCP ports open.
- Anonymous access on FTP.
- SQL injection possible on Web site.
- Weak user password on domain.
- FTP server crackable with dictionary attack.

Higher-level assessment involved:

- On-line tortoise selling from hidden links.
- Secret messages stored on site within guest FTP folder and in Web files, which are cracked by brute force.
- Hidden messages are contained within images and other content. The evidence leads to an intruder uploading an illegal Web site.

## IV. Malware Analysis

The DFET Cloud also allowed students to develop skills in malware analysis, and this part of the module involved students investigating:

- **Back-doors.** This used Metasploit to generate the PUTTY.EXE program with back-doors, and then students used an IDS system to detect the presence of the back-door. This was done within an isolated environment, which did not connect to the Internet. A demo link is contained in [6].
- **Investigation of a real-life malware package.** With this students started from a fresh unpatched Windows XP instance, and were told to make sure they were disconnected from any network (along with being isolated on a separate VLAN). The Worm.Win32 .Dorkbot malware was then installed on the instance and student investigated how it created a hook in the system in the Windows registry and how it hid itself on the system. They then had the task of manually deleting the malware, which was challenging as the malware uses a range of methods to hide itself, and stop itself from being deleted. Students also then analysed the network footprint of the malware, in order that they can setup network probes to detect it. A demo link is contained in [7]. At the end of the lab, students reset the VM back to its original state.

## V. Cloud Setup

The current DFET Cloud contains four main cluster nodes (see Figure 2), where each cluster node runs VMware vSphere 5.5 with VMware vCenter used to manage the instances. This gives a total of 119 GHz CPU, 511 GB of physical memory, and 18 TB of disk space. Overall DFET Cloud contains two main clusters:

- **DFET**. This contains four cluster nodes with a total 119GHz of CPU, 511GB of memory and 18.14TB of disk storage. In a normal running mode there are over 800 virtual machines and templates.
- **SOC Cluster**. This is a cluster with older cluster nodes, but which is used as a back-up if the main cluster was to fail.

Figure 3 shows the details for the setup, and the networks it connects to, along with disk storage elements. The disk storage into three main storage areas: DFET01, DFET02 and DFET03, with around 5TB each, running over RAID-10, in order to support failure of a disk in the array.

In order to provide priority for management tasks, there are two Resource Pool allocations:

- **Management**. This is allocated to the main management VMs, such as vCenter and Windows Domain controller.
- **DFET Lab**. This is allocated to most of the student VMs, and includes a limit on the CPU burst and memory allocation.

The networking of the cluster involves setting up a number of VLANs, which are used to isolate each of the networks created. Overall the mapping to the VLAN is done dynamically and depends on the lab. In order to make sure the management traffic for the VMs are kept separate from the lab traffic, a VLAN of 100 has been created for management traffic.

The VMs are created from single templates, with selected templates of Microsoft Windows XP (unpatched), Microsoft Windows 2003, Microsoft Windows 2008, Linux Kali, Ubuntu, Metasploit, Metasploitable, pfSense, and Vyatta. These have been selected as they give students a good range of systems in which to configure. Each module contains around eight labs covering a range of topics, where each lab is undertaken by a group of two students, and where each group fits into an overall network architecture, with differing networking requirements each work. Most labs about at least two VMs, so that each student will take control of at least one of the instances. This approach provides increased interaction between students.

Along with the VMs running in the DFET Cloud, all of the VMs were available on a shared OneDrive in an OVA format, which allowed students to download their own versions of the instances that they used in the lab, and install on their home system.

## VI. Evaluation

The DFET Cloud ran well over two semesters (2014/2015) with students especially appreciative of the lab demos on YouTube, and which allowed them to either play them back in the lab watch a complete solution of the lab, or playback from home if they had missed the labs. Along with this students especially liked the ability to go home and repeat the lab, especially in resetting it to its initial state.

The greatest challenge for the DFET Cloud was the CSN10107 module which included the coursework for a full penetration test, and with a hand-in date of 28 April 2015. At its peak it had over 80 students performing scanning, DoS evaluation and adversarial roles. For this a bank of test servers were setup, and students could scan these, and evaluate their security infrastructure. As students could change the passwords on the servers, these were rolled-back on a regular basis to a defined state so that students could work on clean versions of the Web targets.

The peak of the CPU utilization occurred on the evening of the coursework submission, where the CPU utilization reached over 25% overall (22.5 GHz). Figure 5 more clearly shows the activity over the month where the Penetration Testing coursework was undertaken, with most variation around the week before the hand-in. At 6pm on the evening before the coursework was to be handed-in, students reported a slow-down in their pen testing, and the source of the problem was not an exhaustion of resources on the cluster nodes, but an exhaustion of resources on the vCentre VM, as it was running on only 16GB of memory, with a reboot and reallocation of more memory, the Cloud returned back to its normal performance.

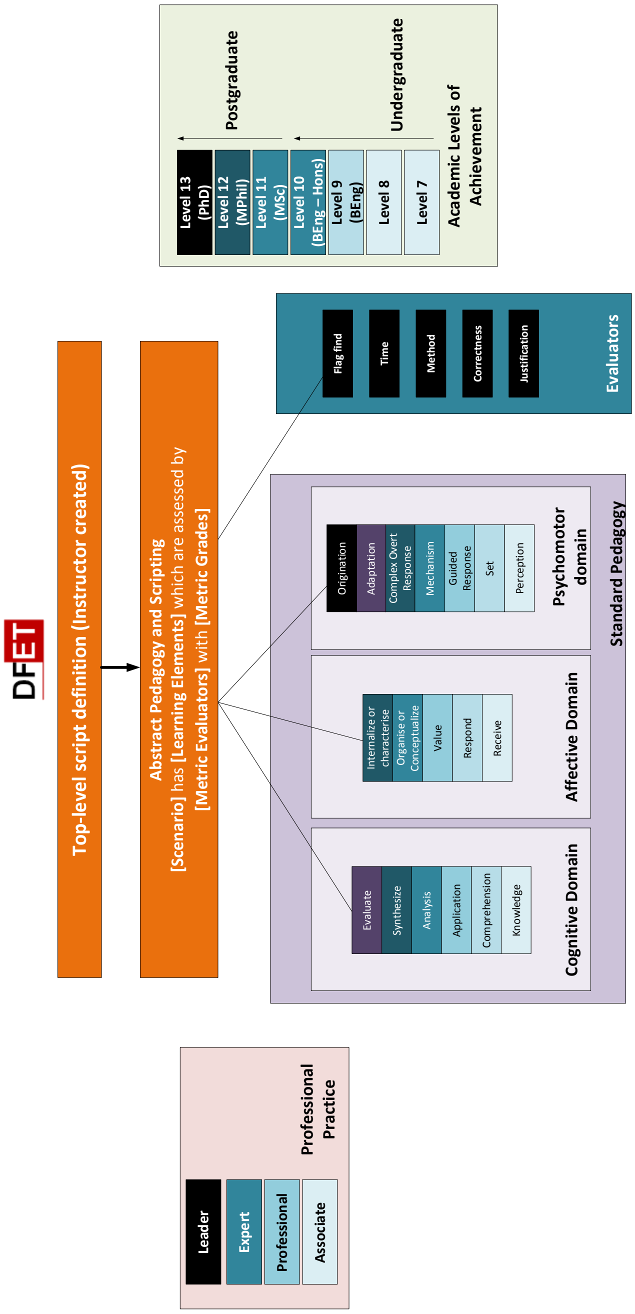

Figure 1: Pedagogy approach

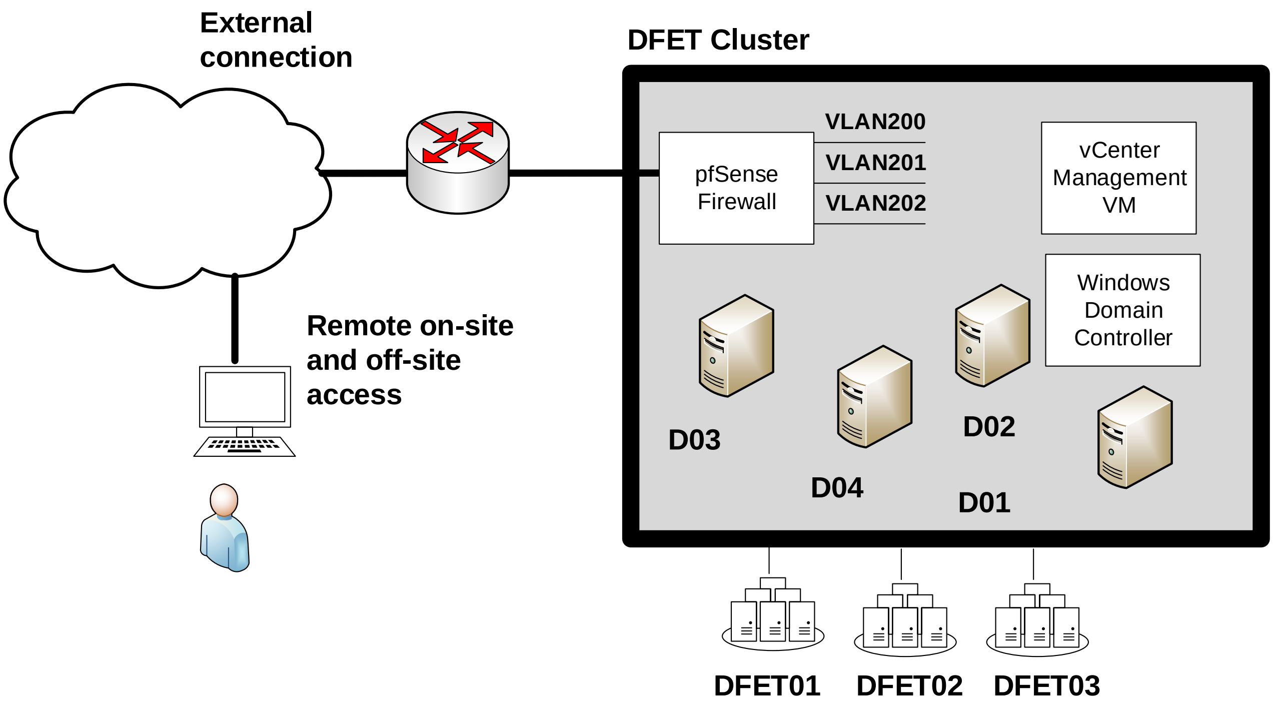

Figure 2: Overview of DFET cluster

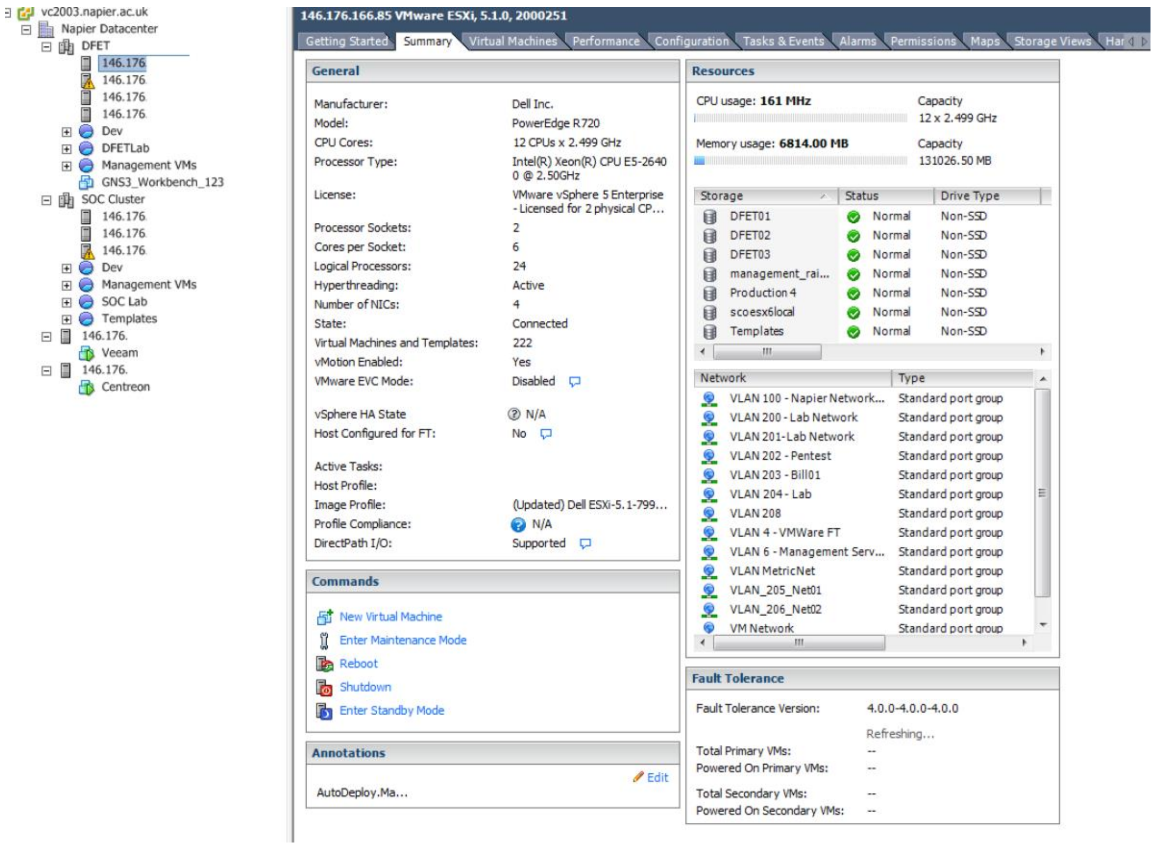

- vc2003.napier.ac.uk
  - Napier Datacenter
    - DFET
      - 146.176
      - 146.176
      - 146.176
      - 146.176
      - Dev
      - DFETLab
      - Management VMs
      - GNS3_Workbench_123
    - SOC Cluster
      - 146.176
      - 146.176
      - 146.176
      - Dev
      - Management VMs
      - SOC Lab
      - Templates
    - 146.176.
      - Veeam
    - 146.176.
      - Centreon

146.176.166.85 VMware ESXi, 5.1.0, 2000251

**General**

| | |
|---|---|
| Manufacturer: | Dell Inc. |
| Model: | PowerEdge R720 |
| CPU Cores: | 12 CPUs x 2.499 GHz |
| Processor Type: | Intel(R) Xeon(R) CPU E5-2640 0 @ 2.50GHz |
| License: | VMware vSphere 5 Enterprise - Licensed for 2 physical CP... |
| Processor Sockets: | 2 |
| Cores per Socket: | 6 |
| Logical Processors: | 24 |
| Hyperthreading: | Active |
| Number of NICs: | 4 |
| State: | Connected |
| Virtual Machines and Templates: | 222 |
| vMotion Enabled: | Yes |
| VMware EVC Mode: | Disabled |
| vSphere HA State | N/A |
| Host Configured for FT: | No |
| Active Tasks: | |
| Host Profile: | |
| Image Profile: | (Updated) Dell ESXi-5.1-799... |
| Profile Compliance: | N/A |
| DirectPath I/O: | Supported |

**Commands**

**Annotations**

AutoDeploy.Ma...

**Resources**

CPU usage: **161 MHz** Capacity 12 x 2.499 GHz

Memory usage: **6814.00 MB** Capacity 131026.50 MB

| Storage | Status | Drive Type |
|---|---|---|
| DFET01 | Normal | Non-SSD |
| DFET02 | Normal | Non-SSD |
| DFET03 | Normal | Non-SSD |
| management_rai... | Normal | Non-SSD |
| Production 4 | Normal | Non-SSD |
| scoesx6local | Normal | Non-SSD |
| Templates | Normal | Non-SSD |

| Network | Type |
|---|---|
| VLAN 100 - Napier Network... | Standard port group |
| VLAN 200 - Lab Network | Standard port group |
| VLAN 201-Lab Network | Standard port group |
| VLAN 202 - Pentest | Standard port group |
| VLAN 203 - Bill01 | Standard port group |
| VLAN 204 - Lab | Standard port group |
| VLAN 208 | Standard port group |
| VLAN 4 - VMWare FT | Standard port group |
| VLAN 6 - Management Serv... | Standard port group |
| VLAN MetricNet | Standard port group |
| VLAN_205_Net01 | Standard port group |
| VLAN_206_Net02 | Standard port group |
| VM Network | Standard port group |

**Fault Tolerance**

| | |
|---|---|
| Fault Tolerance Version: | 4.0.0-4.0.0-4.0.0 |
| | Refreshing... |
| Total Primary VMs: | -- |
| Powered On Primary VMs: | -- |
| Total Secondary VMs: | -- |
| Powered On Secondary VMs: | -- |

Figure 3: Cluster outline

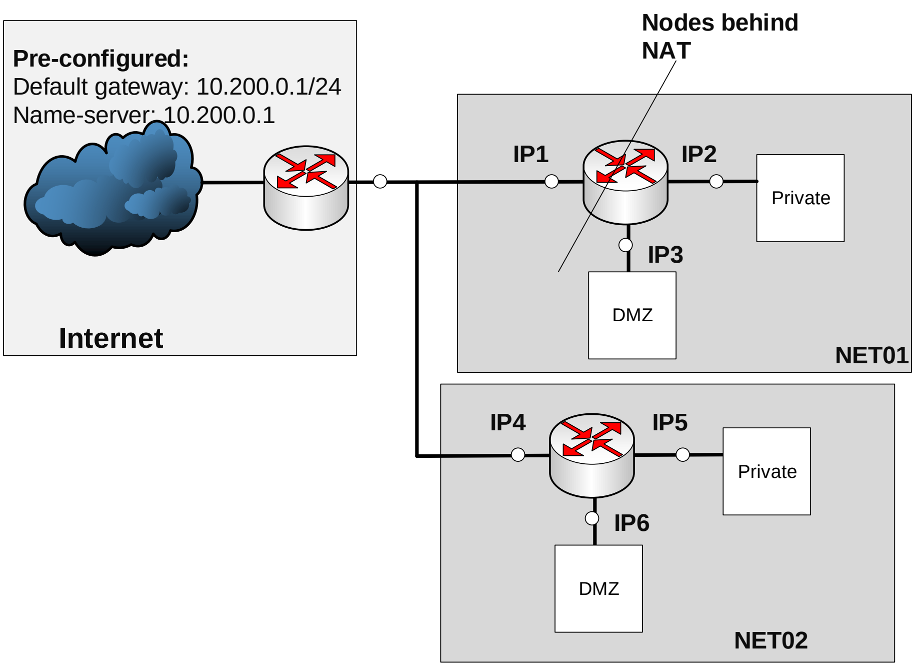

Figure 4: Firewall lab setup

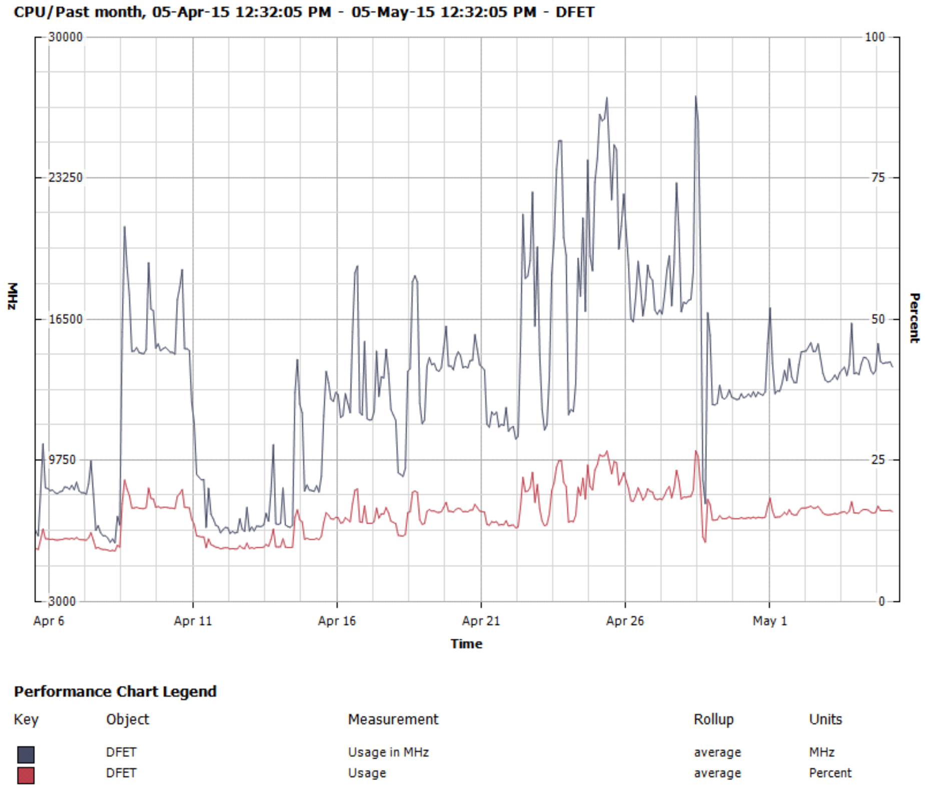

Performance Chart Legend

| Key | Object | Measurement | Rollup | Units |
|---|---|---|---|---|
| 🟦 | DFET | Usage in MHz | average | MHz |
| 🟥 | DFET | Usage | average | Percent |

Figure 5: DFET CPU April 2015

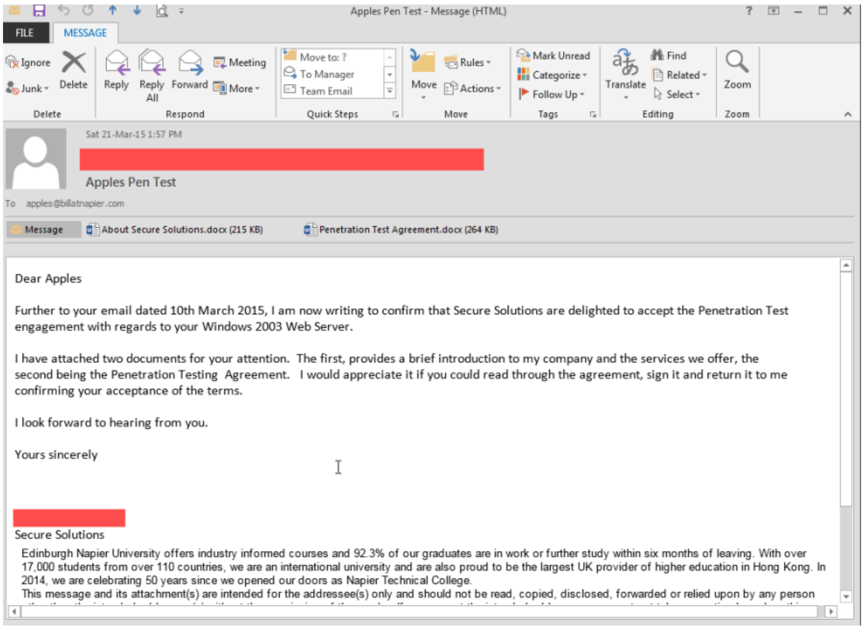

Sat 21-Mar-15 1:57 PM

[redacted]

Apples Pen Test

To apples@billatnapier.com

Message About Secure Solutions.docx (215 KB) Penetration Test Agreement.docx (264 KB)

Dear Apples

Further to your email dated 10th March 2015, I am now writing to confirm that Secure Solutions are delighted to accept the Penetration Test engagement with regards to your Windows 2003 Web Server.

I have attached two documents for your attention. The first, provides a brief introduction to my company and the services we offer, the second being the Penetration Testing Agreement. I would appreciate it if you could read through the agreement, sign it and return it to me confirming your acceptance of the terms.

I look forward to hearing from you.

Yours sincerely

[redacted]

Secure Solutions

Edinburgh Napier University offers industry informed courses and 92.3% of our graduates are in work or further study within six months of leaving. With over 17,000 students from over 110 countries, we are an international university and are also proud to be the largest UK provider of higher education in Hong Kong. In 2014, we are celebrating 50 years since we opened our doors as Napier Technical College.
This message and its attachment(s) are intended for the addressee(s) only and should not be read, copied, disclosed, forwarded or relied upon by any person

Figure 6: Sample engagement email

## VII. Conclusions and Future Work

The DFET cloud infrastructure has supported over 200 students running with over 800 virtual machines. Overall the range of tools and environments that can be used within the Cloud enhances computer security education, while making sure that the labs work in a dependable way. A key factor in the setup of the DFET is the usage of VM pools which are created once from templates, and then created with a snap shot. When students take the lab, the VM is deployed into their folder, and allocated to the correct network. Once finished it can be returned back to the snap shot and put back in the pool. Students can also keep with the same VM throughout the module (and reset it for each lab).

Plans are underway for another scale-up with new servers with 192GB of physical memory. This would allow for thousands of VMs, running full penetration testing tools, and where each student gets their own environment, which links to other student networks, allowing for group work and inter-group collaboration.

The Penetration Testing coursework has been a particular success, especially in terms of supporting not only the development of practical skills, but in developing business communication skills (Figure 6). Often students who had the best evaluation for the coursework, also were the best at keeping in touch with the company involved in the Penetration Test. Contacts with local companies have identified that they were keen to recruit students who had strong technical skills, but also had a good background in customer engagement skills.

One challenge is that, although there were several target servers setup for the penetration test, students typically focused on one of them, and if it was compromised, other students would not get the correct target. For the forthcoming year, each student will get their own isolated environment with their own target. They will also have full rights to take snap shots, and also to return them to any of the snap-shots, along with full rights to reboot the target server. The team are working on the creation of the target Web servers, so that every server will be different in setup of the most advanced information that can be gained from the target.

An important element of the module, and the usage of the Cloud, is the mix of Microsoft Windows and Linux, especially where students get used to setting up both these systems for the given scenario (which differed each week, with different IP ranges), along with them using command line tools, rather than GUIs – which local industry have highlighted that there is an increasing required for command line tools and scripting. Overall every student used the Linux Kali instance for their Penetration Test coursework, with only two students also requesting Windows 2003 for the usage of Nessus. This highlights the preference of students to use command line tools, as they were shown in the lab.

## VIII. References

[1] CSN09112 (BEng Module), https://asecuritysite.com/csn09105

[2] CSN10107 (BEng (Hons) module), https://asecuritysite.com/csn10107

[3] CSN11123 (MSc module), https://asecuritysite.com/csn11123

[4] pfSense Lab, https://dl.dropboxusercontent.com/u/40355863/csn09112_lab03.pdf

[5] http://thecyberacademy.org/about/dfet/

[6] http://asecuritysite.com/subjects/chapter50

[7] http://asecuritysite.com/subjects/chapter14