# Using a VMware Network Infrastructure to Collect Traffic Traces for Intrusion Detection Evaluation

by
Frederic Massicotte, Mathieu Couture and
Annie De Montigny Leboeuf
**http://www.crc.ca/networksystems_security/**
**{frederic.massicotte, networksystems-security}@crc.ca**

CENTRE DE RECHERCHES SUR LES COMMUNICATIONS RESEARCH CENTRE

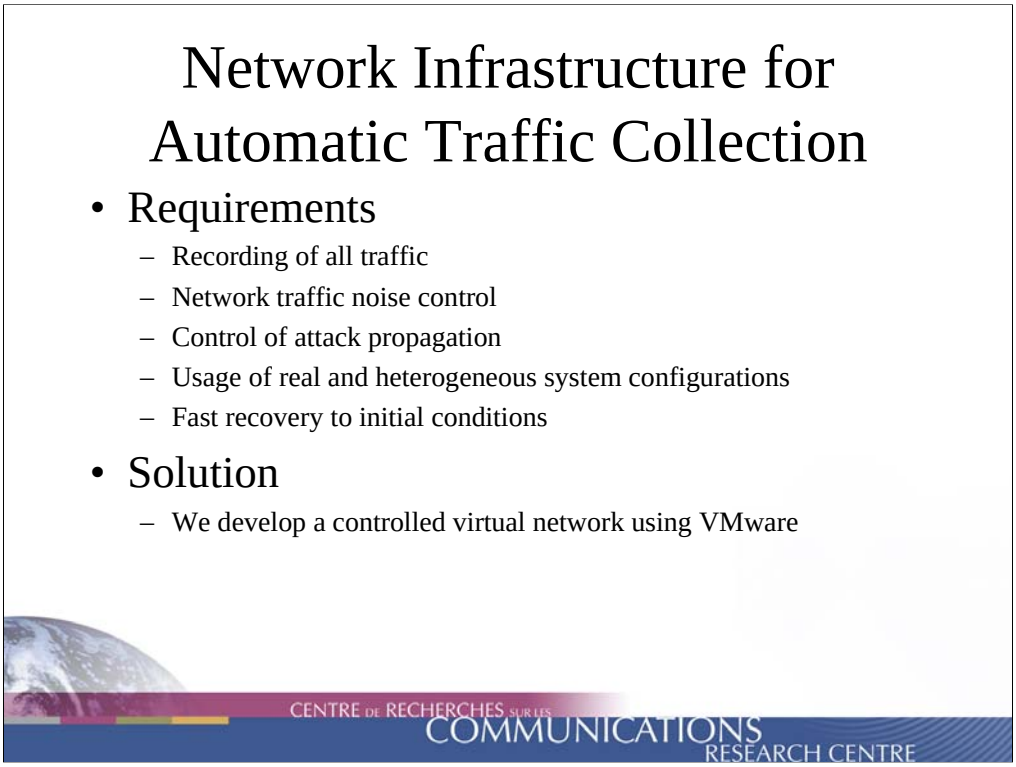


Many research, including those on IDS, do require testing and evaluation.

This work proposes an **automated approach** to develop large data sets of **attack traces**

Our infrastructure had to fulfill the following requirements :

Record traffic, to allow post analysis;

Control noise, everything in the trace is known and relevant to the experiment;

Control of attack propagation, confine attacks to prevent infection propagation;

Realistic targets and a great variety of them;

Fast recovery to initial conditions (prior to attack), to reproduce experiment under the same conditions.

We chose to build a virtualized environment in which a great variety of systems can be tested in an automated fashion.

VMware offers a lot of functionalities “out of the box”: reverting machines to a given state, cloning, and support for many OS families. We have installed over 200 OS versions among the most popular families (FreeBSD, OpenBSD, NetBSD, Linux, Windows).

# Virtual Network Infrastructure

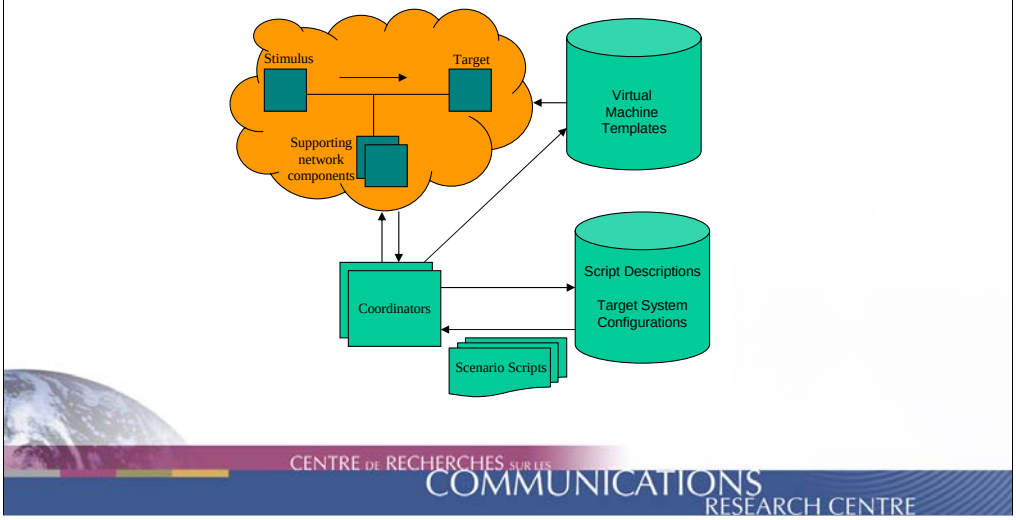


A core component is the coordinator, it has access to a database containing the description of the scenarios, it can pull specific scenario scripts.

From these scripts, the coordinator chose which targets and attackers are required along with other network components, if needed, to support the communications (e.g. DNS, router).

It sets up the virtual network, and give orders to the attackers, it collects the traffic and labels the traces according to the scenario specification. And finally the coordinator tears down the network and reset the virtual machines back to their original states.

# Examples of Application

- Passive Operating System Fingerprinting Data Set
  - Captured over 200 operating system behaviours (with an older version of the virtual network infrastructure)
- Fragmentation Impact Assessment Data Set
  - Captured over 90 packet fragmentation behaviours of operating system using Fragroute (fragmentation overlapping and reassembly timeout).
- **Intrusion Detection Evaluation Data Set**
  - **2343 traffic traces (now over 6000)**
  - **26 operating system versions (now over 85)**
  - **92 vulnerability exploitation programs (now over 95)**

# Intrusion Detection Data Set

- Objectives
  - Automatically execute and test vulnerability exploitation programs
  - Use this data set against IDS
    - Look for false positive and false negative problems
  - **Produce a data set of exploit traffic traces (freely available)**

# Intrusion Detection Data Set Classification and Labelling

| Operating System Family | Operating System Versions | Scenario instances | Vulnerable/ Not Vulnerable | Success/ failure/ unclassified |
|---|---|---|---|---|
| FreeBSD | 7 | 270 | 73/197 | 4/27/239 |
| Linux | 6 | 436 | 79/357 | 10/77/349 |
| Windows | 13 | 1637 | 948/689 | 166/729/740 |

- Automatic classification of attack outcomes
- Attack are launched against vulnerable and non-vulnerable operating systems
  - To identify IDS accuracy for unsuccessful attacks (false positive)

CENTRE DE RECHERCHES SUR LES COMMUNICATIONS RESEARCH CENTRE

Classification and labeling: to be useful, it was felt that the traces had to be properly named (or labeled).

The traffic is separated into multiple traces. Each trace contains the traffic associated to an attack towards one target.

The name of each traffic trace gives the exploit program used, the target OS description, whether the target was vulnerable or not to the attack, and whether the attack was successful or not.

For the data set of traffic trace currently available, it was decided that for each exploit program, all targets running a service on the port targeted by the exploit would be attacked, whether the targets were running a vulnerable version of the service or not.

When determining whether the attack was successful or not, some cases were difficult to classify automatically (without human intervention). Efforts are currently being made to find ways to better discriminate between success and failure automatically.

# Intrusion Detection Data Set
# Exploit Distribution

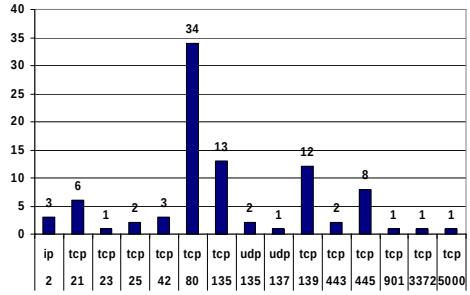


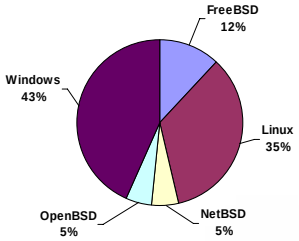


CENTRE DE RECHERCHES SUR LES COMMUNICATIONS RESEARCH CENTRE

On top of having all traffic traces labeled, some basic statistics can be extracted from the database to further document the dataset.

# Questions ???

**Contact Information :**
Frederic Massicotte
http://www.crc.ca/networksystems_security/
**{frederic.massicotte, networksystems-security}@crc.ca**

CENTRE DE RECHERCHES SUR LES COMMUNICATIONS RESEARCH CENTRE