3

HIGH PERFORMANCE COMPUTING & CYBER SECURTY

CSIR-4PI has been providing state-of-the-art High Performance Computing (HPC) facility to the computational scientists and researchers of CSIR to address significant problems in their frontier areas of science and engineering. The centralized HPC facility located at CSIR-4PI is operational on a 24x7 basis with high uptime efficiency. The facility at CSIR-4PI is one of the top supercomputers in the country and provides a computing platform suitable for multiple domain specific applications. All the CSIR laboratories access the facility through the high speed, National Knowledge Network with redundant path. In addition to providing the HPC facility, the group is also actively involved in research on various aspects of cyber security and cryptography. Extending the research carried out under the 12th Five Year Plan of CSIR, the team is participating in a Mission Mode Project of CSIR on Intelligent Systems in which the team is concentrating on the security and privacy issues in connected vehicles and biometric based transactions.

Inside

3.1 Cryptojacking in HTTP/2 Framework

Cryptojacking is a malicious activity in which compute resources are used for mining crypto currency without the consent of the owner of the resources. As the popularity and usage of crypto currency has been increasing rapidly, the resource requirement for mining them is also increasing proportionately. Consequently, Cryptojacking has emerged as a new security threat on the cyberspace. Identifying novel means of exploitations and countermeasures have become a topic of interest to cyber security community.

Typically, an adversary performs Cryptojacking either by hacking into the victim computer or by injecting malware with Cryptojacking code into the victim computer for auto-execution. Once injected, the code runs on the victim computer in the background and mine the currency silently.

In this work, we focus on an alternate way of performing Cryptojacking. In particular, we explore the feasibility of exploiting Internet middle-boxes (like open proxies, ToR exit points, etc.) for injecting Cryptojacking code. Note that the usage of open proxies for accessing the Internet has been increasing because of anonymity reasons. Many users, who prefer to use such services to hide his/her identity like Internet Protocol (IP) Address and physical location may become a victim of Cryptojacking.

We used a relatively simple experimental setup consisting of three entity: A client machine that acts as a base user; an anonymous open proxy server through which all the client communications passes; and an HTTP/2 web server that provides HTTP/2 service to the client machine. The proxy component was modified to inject the Cryptojacking code into the client in a seamless manner. For this, the proxy intercepts the response from the HTTP/2 webserver, modify the html page and insert the java script containing the malicious code. Neither the client nor the server is aware of the code injection process. The client browser just displays the original html page served by the HTTP/2 webserver, and the Cryptojacking code execution is not visible on the browser.

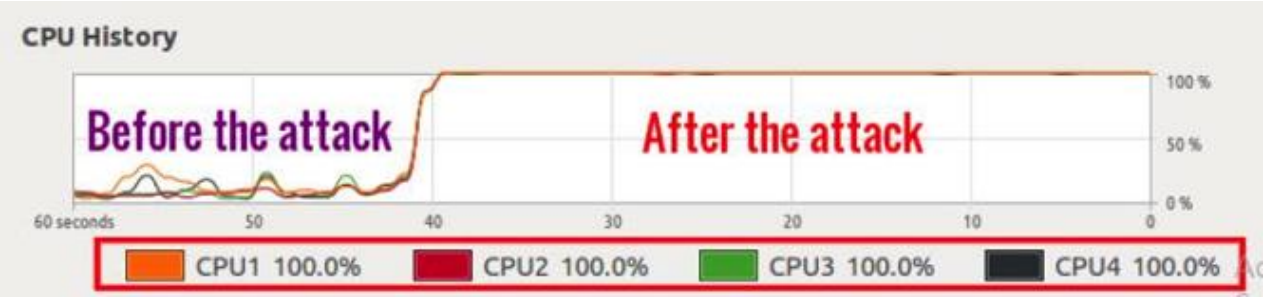

Figure 3.1 CPU Utilization of the client before and after the Cryptojacking Attack. The axes X and Y, respectively, show the running time and percentage of CPU utilization

In Figure 3.1, we capture a running screen-shot of the CPU utilization of the client. As seen in the Figure, before the code execution, the CPU utilization was small and mainly used by the Operating System and some default daemons. Once the Cryptomining is started, the CPU utilization of all the four CPUs has gone to 100%. The only indication that a careful user will have is that the system response is relatively slow.

3.2 Tree parity machine based group key agreement protocol

There are large number of network applications that involve multiple users and demand security. Hence, there is a need to develop multi-party key exchange that can provide secure communication among parties. In this work we have extended Tree Parity Machine (TPM) based key exchange to group based key exchange. Here, each participant is treated as a node of a binary tree, holding his or her own TPM. Each participant starts a synchronization/learning process, with his or her sibling partner at the lowest layer. At the end of synchronization, the siblings act as the parent and again participate in synchronization with their sibling in the next upper layer. The synchronization process moves upwards in the binary tree, till the root is reached. This indicates that all the participants have the same synchronized key to participate in the multi party application. Figure 3.2 shows the mean synchronization step for different number of participants. The figure shows the difference between, the traditional ring structures with our proposed binary tree structure. It is clear that the tree based mechanism takes less learning steps to generate a group key.

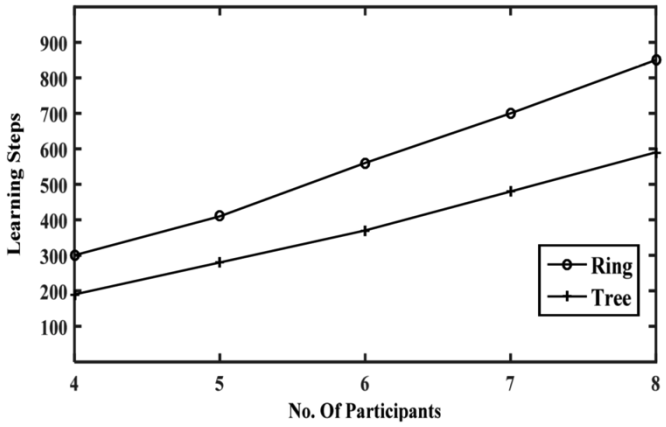

Figure 3.2 The mean synchronization steps with different number of participants between ring and tree based group key agreement for same TPM structure

3.3 High Performance Computing

The “Ananta” Supercomputer (Figure 3.3) continues to serve the CSIR computational science community tirelessly for the 6th year in a row as the centralized facility. The computational workforce got a major boost with the addition of 127 TF of additional computational capability powered by 48 numbers of Intel Skylake processor based nodes.

Figure 3.3 CSIR centralized 489 TF High Performance Computing Facility

With this upgrade, the system currently has the capability to carry out about 489 Tera floating-point operations in a second. The sustained performance using a High Performance LINPACK (HPL) is about 334 TF for the original system and 84 TF for the additional nodes and the system continues to be the largest Supercomputer of CSIR. In addition, the centre also hosts an Altix-ICE medium range HPC along with a hierarchical storage infrastructure.

Currently listed as the 9th fastest system in the country, the supercomputer Ananta, is a cluster consisting of 1088 computing nodes (two numbers of Intel Xeon E5 2670, 8 core processors each) and 48 numbers of upgraded nodes

(two numbers of Intel Xeon Gold 6140, 18 core processors per node), distributed over 18 racks. While the original system has 64 GB per node, the upgraded nodes have 192GB memory per node. This amounts to about 77TB of distributed memory for the total system. The inter-node communication is powered by high speed FDR (for original nodes) and EDR (in the upgraded nodes) infiniband providing a dedicated 56 and 100 Gbps interconnect bandwidth respectively. However, both the set of nodes access the common LUSTRE parallel file system of about 2.1 Peta Byte of usable capacity, which is capable of providing a minimum of 20 GB/s simultaneous read and write performance. PBS Pro workload manager ensures efficient usage of the system.

One of the reasons of “Ananta” Supercomputer providing un-interrupted service for more than 6 years is due the Tier-3 equivalent state-of-the-art data center along with the state-of-the-art energy farm. The highlight of the datacenter is the water based cooling mechanism called Rear Door Heat Exchangers (RDHx) that has resulted in providing one of the best power efficient and high-density datacenter (Power Usage Efficiency (PUE) of less than 1.5) in the country. The energy farm consists of two numbers of redundant compact substations of 1.25MVA, three numbers of 1010 KVA diesel generators, an underground diesel yard (more than 15000 liters) and three numbers of UPS with battery backup for ensuring 24x7 power supply to the datacenter.