Systems, Networking, and Cybersecurity Ph.D. Qualifier Exam 2019

The following questions relate to the paper published in the reading list. For full citations, please see the reading list.

First, read the following guide: Avoiding Plagiarism: A Guide For Graduate Students at Virginia Tech[1]. In your answers, you should avoid unattributed direct quotations and paraphrases and use proper documentation of all sources you use, which will require you to include a bibliography.

You should answer the following questions.

1. The following questions relate to the paper "TensorFlow: A System for Large-Scale Machine Learning" by Martín Abadi et al.
   a) Express the key ideas of this paper in one paragraph!
   b) Why does TensorFlow assume a static dataflow graph?
   c) Name two separate examples of design decisions which TensorFlow makes that exploit specific assumptions regarding the nature of the machine learning computations for which it is designed! For each decision, discuss how these assumptions influenced the design!
   d) In your opinion, how does this paper advance the state of the art in the area of distributed systems? Justify your opinion!

2. The following questions relate to the paper "My VM is Lighter (and Safer) Than Your Container" by Filipe Manco et al.
   a) Express the key propositions of this paper in one paragraph!
   b) In the paper, the authors state that "The syscall API is fundamentally more difficult to secure than the relatively simple x86 ABI offered by virtual machines where memory isolation (with hardware support) and CPU protection rings are sufficient."
   Discuss the merits of this statement! How strong is the evidence in favor of it that is presented by the authors, in your opinion?
   c) Briefly describe what unikernels are and why they are difficult to use in practice for the purposes at which the paper's authors aim?
   d) How and to what degree does LightVM increase the virtual machine's attack surface, and if so, what influence would that have on the authors' claim of being lighter and "safer"?

[1] https://graduateschool.vt.edu/content/dam/graduateschool_vt_edu/graduate-honor-system/avoiding-plagiarism-short-guide.pdf

3. The following questions relate to the paper “Light-Weight Contexts: An OS Abstraction for Safety and Performance” by James Litton et al.
   a) What is the key assertion made in this paper?
   b) Provide an example of a threat against which lwC cannot defend!
   c) All things considered, would you advocate that support for lwC be added to OS such as Linux? Justify your opinion!

4. The following questions relate to the paper “Application performance and flexibility on exokernel systems” by Kaashoek et al.
   a) What are the tradeoffs of monolithic kernels, microkernels, and exokernels?
   b) Compare and contrast exokernels and containers (e.g., Docker).
   c) What is the Trusted Computing Base for an application running on top of an exokernel?

5. The following questions relate to the paper “Tor: the second-generation onion router” by Dingledine et al.
   a) What security properties does Tor provide?
   b) List five side-channels that exist with Tor and how users can eliminate them?
   c) Describe Tor hidden services (e.g., onion sites) and explain how a search engine finds onion sites.

6. The following questions relate to the paper “LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation” by Lattner & Adve.
   a) What are the advantages and disadvantages of using an intermediate language during compilation?
   b) What is the purpose of the phi instruction? Give a simple example of its use.

7. The following questions relate to the paper “The ZMap: Fast Internet-wide Scanning and Its Security Applications”, by Zakir Durumeric, Eric Wustrow, and J. Alex Halderman.
   a) Describe the key designs in ZMAP that allows it to scan the full IPv4 space within an hour?
   b) If you were to improve the scanning speed even further, what other designs would you like to introduce?
   c) What applications of ZMAP can you think of, in addition to those that are already discussed in the paper?

8. The following questions relate to the paper “Detecting Credential Spearphishing in Enterprise Settings” by Grant Ho, Aashish Sharma, Mobin Javed, Vern Paxson, David Wagner.
   a) Why are traditional machine learning or anomaly detection methods not suitable to detect spearphishing?

b) If you are an attacker that aims to evade the detection of the proposed system, how would you design your phishing campaign? Can you think of any new strategies beyond what has been discussed in the paper?

9. The following questions relate to the paper “Towards Evaluating the Robustness of Neural Networks”, by Nicholas Carlini and David Wagner.
   a) Describe the core idea of the proposed attack. How does the proposed method generate stronger adversarial examples than prior works?
   b) What are the key assumptions made in the threat model? Are the assumptions always valid?
   c) The paper focuses on image classification tasks. If we change the application to malware classification, what are the extra requirements that the attacker needs meet in order to make the attack practical (i.e., generating malware samples that will be classified as benign files)?