

# Final Take Home Exam: Spring 2012

---

Cpt S 580-03 / Cpt S 483-01 / EE 582-02 / EE 483-01

John Grimes, Riley Huddleston, Jiawei Ning, Lin Zhang

Due May 3, 2012

1. Choose one real distributed application program that you are familiar with (or could become familiar with if you don't know any (!!!)). Write the following:

- A paragraph or two giving an overview of the application. Make sure to include which generic architecture (client-server, publish-subscribe, P2P, . . . ) you believe best categorizes the application, and why it does so. Be sure also to describe what application-specific activities the various pieces in that architecture (clients, servers, whatever) are doing: what they are requesting, processing, etc.
- A paragraph or two describing whether you believe the following runtime issues are important to users of this application, and why you believe it (e.g., what happens in its absence):
  - Low latency network connections
  - High bandwidth available
  - Perfectly “consistent” and “correct” replies/answers/service versus an inconsistent or approximate reply/answer/service

Make sure your answer indicates how you would desire to trade off the above runtime properties, and the “worst” for each it could reasonably tolerate.

OpenVMS is a server operating system that is a highly flexible, general-purpose multiuser system that allows for physical separation of hardware, creating something of a distributed client/ server system. OpenVMS installations are meant for environments where uptime is critical, and can utilize properties of distributed systems and features such as rolling updates to ensure uptime. Among the features of OpenVMS are time sharing, real-time operations, transaction processing, a distributed file system, a GUI, and mixed architecture clustering.

The system itself is run like a server by fulfilling client requests for access to files, printing, application services, communication services, and computing power. These requests may come from users with physical access to a node in the cluster, or from remote users. For example, OpenVMS supports Apache HTTP Server, and thus may be configured with an integrated web server. For other requests, specialized software on the client side may be required.

OpenVMS is an operating system, so it is vital that nodes in the cluster that make up an OpenVMS installation have low latency and moderately high bandwidth. If there is too much latency between nodes in the cluster, then requests for resources such as file access or compute time may go unanswered and result in inefficient resource allocation, or in the worst case result in a node being

considered offline. Bandwidth requirements are moderately high because of the distributed nature of the OpenVMS filesystem; when a client or specific node requests data that is not stored locally, it must be transferred within the network. If the bandwidth is too low, depending on the frequency of these requests, the system may crawl to a halt waiting for important data. Due to the closed nature of an OpenVMS installation, there is no concern about requests returning invalid information.

2. *Explain in a paragraph or two in your own words the analogy "Middleware is to socket programming what high-level language programming is to assembler programming".*

Middleware and high-level language programming both abstract away complexity of their respective lower layers. Middleware simplifies communication between different pieces of software by providing a common interface between machines, meaning that explicit socket programming interface isn't necessary. Middleware usually is tailored for its specific purpose and presents interfaces for that purpose. For example, distributed computing middleware provides relevant paradigms for distributed computing. Similarly, high level programming languages provide abstractions that build upon lower layers to accomplish tasks without requiring a programmer to deal with the lower layers directly. In comparison to assembly programming, C can be viewed as a high level language because it abstracts away registers and the stack while introducing features such as scoped variables, functions, and data types.

3. *Which of the following is false for a NASPInet-like critical infrastructure data delivery service for a power grid:*
  - A. *The scale is such that tracking per-flow state at a router (or similar forwarding device) is feasible.*
  - B. *The admission control perimeter can be complete.*
  - C. **Post-error recovery is sufficient.**
  - D. *The predictability of the delivery latency, even in the presence of a small (bounded) number of failures, is very good or better.*
4. *Which of the following is more burdensome for the application programmer:*
  - A. *Remote procedure call*
  - B. *Publish-subscribe*
  - C. *Remote method invocation*
  - D. **Request-reply protocols**
5. *Which of the following is space-uncoupled:*
  - A. *Remote procedure call*
  - B. **Publish-subscribe**
  - C. *Remote method invocation*
  - D. *Request-reply protocols*

6. Which is not either a motivating observation or a property provided by P2P systems (i.e., what is false):
- A. Compared to exploiting in-network logic, the same level of multicast efficiency and very low delivery latency is achievable (even in the presence of a few faults).
  - B. CPU/GPU and storage resources are often more abundant “at the edges” of a distributed system.
  - C. **A much larger address space can be handled compared to IPv4 or even IPv6.**
  - D. Using only edge resources means that deployment is much simpler than if in-network logic (e.g., GridStat forwarding engines) were required.
7. Which if the following challenge for the bulk power grid can not be mitigated with much better wide-area communications (and sensors feeding it):
- A. The different physics of some kinds of renewable power (e.g., no reactive power) and how that changes the dynamics of a grid at large.
  - B. **Negligible storage of power in the grid.**
  - C. Large numbers of retiring operators who are leaving with a lot of institutional knowledge—intuitive “seat of the pants” understanding—of how their grid operates in many contingencies or unusual situations (and combinations thereof).
  - D. Increasing stress on grid due to insufficient new transmission capacity compared to increases in both generation and load.
8. Tripwire is a software tool intended to assure integrity of system files by detecting unexpected modifications (such modifications are often a sign of rootkit activity). One version of Tripwire reads the names of the directories to be protected from a configuration file. For each file in the specified directories, Tripwire computes its hash value and stores it in a database. What property must this hash function have?
- A. preimage resistance
  - B. **2nd-preimage resistance**
  - C. collision resistance
  - D. all of the above
9. SYN cookies (<http://cr.yip.to/syncookies.html>) are used to mitigate SYN flooding attacks. Which of the following is not used in computing a SYN cookie?
- A. Timestamp
  - B. **TCP flags**
  - C. Maximum segment size
  - D. Port number
10. Suppose you have an intrusion detection system detecting a computer virus on the network with 90% accuracy. Precisely, the IDS detects a connection transferring a virus as an attack with 90% probability and a benign connection as an attack with 10% probability. When 1% of the connections

contain a virus, what is the probability that a connection flagged by the IDS as an attack is actually benign?

- A. 0–20%
- B. 21–50%
- C. 51–90%
- D. 91–100%**

11. The following statements describe Kerberos authentication. Find an incorrect statement.

- A. The KDC is a single point of failure.
- B. KDC and the ticket-granting server can run on a single machine.
- C. Kerberos uses symmetric encryption.
- D. Every message in Kerberos authentication is encrypted.**

12. Which statement is false?

- A. Stuxnet exploited multiple zero-day vulnerabilities.
- B. Stuxnet attacked systems from certain vendors only.
- C. A network physically separated from the Internet can be infiltrated by Stuxnet.
- D. The author of Stuxnet created two fake certificates to sign drivers.**

13. The article “W32.Duqu: The precursor to the next Stuxnet” describes Duqu, a threat similar to Stuxnet. Compare Stuxnet and Duqu from the following aspects.

- (a) Initial infection
- (b) Propagation
- (c) Command and control

### Initial Infection

Stuxnet is initially spread using infected removable drives such as USB flash drives using techniques such as a hidden autorun file. The worm then uses other exploits and techniques, such as peer-to-peer RPC, to infect and update other computers inside private networks that are not directly connected to the Internet. Duqu arrived at the target using a specially crafted Microsoft Word document. The Word document contained zero-day kernel exploit that allows the attackers to install Duqu onto the computer unbeknownst to the user. This shows that both Stuxnet and Duqu used windows kernel vulnerabilities for their initial infection.

### Propagation

Stuxnet used a number of vulnerabilities in order to facilitate mass self-propagation. These vulnerabilities include the infection of WinCC machines through a hardcoded database password, copying itself through network shares, peer to peer, a Print Spooler zero-day vulnerability (MS10-061), and a Windows Server Service vulnerability (MS08-067). On the other hand, Duqu does not self-replicate at all, and the exact method used to replicate within a network has yet to be determined. Duqu has been found to be primarily a remote access Trojan, and the attackers can execute any code they want from the computer, providing an avenue for manual propagation.

## Command and Control

To contact command and control, Stuxnet started by testing the network connectivity of the machine it had infected. If the machine was directly connected to the internet, Stuxnet could then contact the command and control server directly to send information back and forth freely. Duqu used HTTP and HTTPS to contact a command and control server to download executables to be run on the system. Duqu was also primarily meant for data acquisition and sent information about the network it had infiltrated to a command and control server.

14. Read *"21 Steps to Improve Cyber Security of SCADA Networks"* (<http://www.oe.netl.doe.gov/docs/prepare/21stepsbooklet.pdf>) from DoE, and answer the following questions.
- (a) *If properly implemented, how would each step help preventing Stuxnet attacks?*
- (b) *Are those steps sufficient in preventing future attacks similar to Stuxnet or Duqu? If yes, justify it. If no, extend those steps to prevent such attacks.*
- a.1) *"Identify all connections to SCADA networks"*  
Knowing the network topology allows for threat vectors to be identified and, ideally, mitigated.
- a.2) *"Disconnect unnecessary connections to the SCADA network"*  
With less potential attack paths, a worm like Stuxnet may be less likely to infiltrate a system.
- a.3) *"Evaluate and strengthen the security of any remaining connections to the SCADA network"*  
Utilizing best practices in security will likely prevent propagation of any malicious software.
- a.4) *"Harden SCADA networks by removing or disabling unnecessary services"*  
The possibility of zero-day vulnerabilities in unnecessary services increases risk needlessly.
- a.5) *"Do not rely on proprietary protocols to protect your system"*  
Ensuring the security of the proprietary protocols used in a SCADA network will prevent vulnerabilities due to known but hidden shortcomings that a virus like Stuxnet may take advantage of.
- a.6) *"Implement the security features provided by device and system vendors"*  
The more points at which security is implemented within a system, the more difficult attacks will be to execute.
- a.7) *"Establish strong controls over any medium that is used as a backdoor into the SCADA network"*  
A legitimate backdoor is another potential target for attackers.
- a.8) *"Implement internal and external intrusion detection systems and establish 24-hour-a-day incident monitoring"*  
24-hour surveillance by humans may allow operators to stop attacks that are in progress.
- a.9) *"Perform technical audits of SCADA devices and networks, and any other connected networks, to identify security concerns"*

Many attacks use known vulnerabilities, and technical audits by security professionals will likely protect a system from such attacks.

- a.10) *"Conduct physical security surveys and assess all remote sites connected to the SCADA network to evaluate their security"*  
Unrestricted physical access to a SCADA network allows for circumvention of some security measures in the network.
- a.11) *"Establish SCADA "Red Teams" to identify and evaluate possible attack scenarios"*  
Individuals familiar with the systems in place will more easily be able to identify areas of risk.
- a.12) *"Clearly define cyber security roles, responsibilities, and authorities for managers, system administrators, and users"*  
Establishing clear roles and responsibilities of individuals will prevent security oversights such as everyone thinking that someone else secured a certain part of the network.
- a.13) *"Document network architecture and identify systems that serve critical functions or contain sensitive information that require additional levels of protection"*  
Knowing which parts of the system are more important or more vulnerable will allow for easier identification and mitigation of threats.
- a.14) *"Establish a rigorous, ongoing risk management process"*  
Ongoing risk management provides analysis of potential vulnerabilities and the threats that currently exist, allowing for contingencies to be developed in the case of a Stuxnet-like attack.
- a.15) *"Establish a network protection strategy based on the principle of defense-in-depth"*  
Layered defenses will slow infection and better protect the important portions of the network.
- a.16) *"Clearly identify cyber security requirements"*  
Clear requirements will allow prevent unnecessary vulnerabilities and increase the difficulty of malicious insiders to conduct attacks.
- a.17) *"Establish effective configuration management processes"*  
Correctly configuring network components prevents changes from undermining the overall security of the system.
- a.18) *"Conduct routine self-assessments"*  
Self-assessments ideally find and remove vulnerabilities that a worm like Stuxnet could take advantage of.
- a.19) *"Establish system backups and disaster recovery plans"*  
Backup and recovery can restore networks compromised by a virus to a state before infection.
- a.20) *"Senior organizational leadership should establish expectations for cyber security performance and hold individuals accountable for their performance"*  
Holding individuals accountable for their performance may lead to improved attentiveness in securing a network, and thus preventing an attack like Stuxnet.

a.21) *"Establish policies and conduct training to minimize the likelihood that organizational personnel will inadvertently disclose sensitive information regarding SCADA system design, operations, or security controls"*

Targeting a specific network, or specific hardware like Stuxnet did, is more difficult if information about a SCADA network is never revealed to malicious individuals.

b) The twenty-one steps enumerated will go a long way in securing a system, but additional steps are required to further protect a system.

b.1) The development of improved firewalls to better isolate a SCADA system from the internet.

This measure will decrease the risk of remote attackers infiltrating the system.

b.2) Simulate infection scenarios with code built similarly to Stuxnet in order to reveal weak points in the security of the system.

b.3) Update the hardware and software in SCADA systems to ensure the latest and greatest protections, as determined by security experts.