

Critical Infrastructure in the Power Grid

Spring 2012 - Final Exam

Anilkumar, Rahul

Edward, Benjamin

Doroshchuk, Ruslan

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.

Answer:

Skype is a peer-to-peer VoIP application that allows users to make voice and video calls and send text messages over an internet connection. The architecture is distributed with two types of nodes: super nodes and Skype client nodes. The Skype client nodes are ordinary nodes that run the Skype service and allow the user to make voice calls and send texts. The super nodes are nodes the Skype client nodes must connect to in order to use the service. Super nodes are like end-points in the ordinary telephone system. Any node is eligible to become a super node as long as it has a public IP address and a sufficient CPU, memory and network bandwidth. There is also a central login server that authenticates users and makes sure each user has a unique username. Each node keeps a list of contacts locally. Searching for a contact is a distributed operation since the local data on each node is searched until the contact is found. Signaling is always done through TCP, while the media is transferred over UDP unless the caller or callee is behind port-restricted NAT and/or a UDP-restricted firewall. In this case, the media flows through an intermediate super node, which forwards the media to the respective node. For three-way calls, the most powerful node (in terms of CPU, memory and bandwidth) is chosen to be the host and is in charge of sending packets to the other nodes. In conference calls with more than three participants, mesh conferencing is used, where multiple nodes do this.

Bandwidth usage is typically 3-16 KBps, where at least 2 KBps is required for reasonable call quality, so high bandwidth is not a strict requirement. Getting perfectly consistent and correct service is not critical either since this is a voice/video application and UDP is used. Low latency on the other hand, can be more of a problem, especially if there is a lot of jitter. In these kinds of applications, users will tolerate low quality sound/video over high latencies and jitter. For sending of texts, these requirements are reversed. Consistent and correct service becomes more important, while latency does not make much of a difference when sending texts.

Sources:

<http://arxiv.org/ftp/cs/papers/0412/0412017.pdf>

<http://mnet.cs.nthu.edu.tw/paper/Chance/041125.pdf>

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”.

Answer:

High-level programming languages provide several constructs that abstract away many of the tedious, low-level aspects of assembly languages that are prone to cause errors. High-level languages also help hide many of the heterogeneities associated with assembly languages and provide a uniform interface for the programmer. In the same way, middleware provides higher level building blocks by abstracting away many of the tedious, low-level aspects associated with distributed systems inherent in socket programming, and hides away varying levels of heterogeneity that make the design of distributed systems so difficult. For example, many programming languages provide memory management, an area that is typically the source of many bugs and memory leaks if done manually. In the same way, middleware provides simplified interfaces for connection management. High-level languages provide libraries of functions that hide the differences in instruction sets of various processor architectures, differences that are visible and need to be accounted for when programming in assembly. Similarly, middleware hides the differences between operating systems and even programming languages present in distributed systems that are all too visible in socket programming.

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.

Answer: C

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

Answer: D

5. Which of the following is space-uncoupled:

- A. Remote procedure call
- B. Publish-subscribe
- C. Remote method invocation
- D. Request-reply protocols

Answer: B

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.

Answer: C

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.

Answer: B

8. Tripwire is a software tool intended to assure integrity of system files by detecting un- expected 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

Answer: B

9. SYN cookies (<http://cr.yp.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

Answer: B

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%

Answer: D

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.

Answer: D

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.

Answer: D

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

Answer:

Primarily Stuxnet was the first discovered malware that spies on and subverts industrial control system or a set of similar systems. Industrial control systems are used in power plants. Stuxnet reprograms industrial control systems by modifying code on programmable logic controllers to make them work in the way the attacker desires and also to hide those changes from the equipment operator. This computer worm spreads via Microsoft Windows, and targets Siemens supervisory control and data acquisition (SCADA) systems.

W 32.duqu is defined as a malware which is nearly identical to Stuxnet, but with a completely different purpose. The same authors created Duqu as Stuxnet, as the source code is similar to Stuxnet. Like Stuxnet, this worm also has a valid, but abused digital signature that is used to collect information and keystrokes to prepare for future attacks. The malware was so named as its keylogger creates temporary files with names starting with “-DQ”.

As stated by FERC, “it is not the size of an entity that is critical, but rather the potential to become a vector of vulnerability to the security posture of interconnected control systems.” After the Internet, the smart grid represents the largest cyber attack surface in North America. Smart meters, substations, and intelligent monitors and sensors on transmission and distribution lines represent millions of physically remote and insecure access points to critical utility networks.

The following three aspects are used to compare Stuxnet and W 32.Duqu and give a better insight to both the worms:

a. Initial Infection:

Stuxnet - The security company VirusBlokAda first identified Stuxnet in mid June 2010. It is estimated to have spread around March or April 2010, but the first variant of the worm appeared in June 2009. The second variant appeared in March 2010, a third with minor improvements appeared in April 2010. Stuxnet attacks computers and networks that meet specific configuration requirements. It makes itself inert if Siemens software is not found on infected computers. Prevention is taken such that the infected computer does not spread the worm to more than three others. Stuxnet has layered attack against three different systems.

Windows operating system – Stuxnet attacked Windows using four zero-day attacks. It spreads initially using infected removable drives such as USB flash drives, and then exploits techniques such as peer-to-peer RPC to infect and update other computers present inside the private networks, which might not be directly connected to the Internet. The malware has both user-mode and kernel-mode rootkit capabilities. Its device drivers are digitally signed with the private keys of two stolen certificates, JMicron and Realtek. This driver signing helps in installing the kernel-mode rootkit drivers, which keeps Stuxnet undetected for a long period of time.

Siemens PCS 7, WinCC and STEP7 software infection – Once installed on Windows system, Stuxnet starts infecting project files belonging to Siemens’ WinCC/PCS 7 SCADA control software. It subverts a key communication

library of WinCC called s7otbxdx.dll. This intercepts communications between the WinCC software running under windows and the target Siemens PLC devices when the two are connected through a data cable. This enables the malware to install itself on the PLCs unnoticed and mask its presence from WinCC.

PLC infection – Stuxnet requires the attacked PLCs to have variable-frequency drives from two specific vendors, Vacon and Fararo Paya. It attacks motors whose spin is between 807 Hz to 1064 Hz. Stuxnet installs the malware into memory block DB890 of the PLC that monitors Profibus messaging bus of the system. It affects the motors by changing the rotating speed drastically. Rootkit is installed for hiding itself.

W32.Duqu – Two variants were initially recovered, the first occurred in April 2011. However it may have been conducted as early as November 2010. An attack was launched via e-mail two times, but only the second one was successful. Both times the e-mail was sent by IP-address based in Seoul, South Korea. Once the e-mail was opened, the worm installed itself and stayed inactive till user activity stopped.

One of the variant's driver files was signed with a valid digital certificate. Duqu consists of a driver file, a DLL and a configuration file. These files are installed by an executable – the installer. The installer registers the driver file as a service, such that it starts at the system initialization. The main DLL is then injected into services.exe by the driver file. The main DLL begins to extract other components, which are further injected into other processes. The main DLL component has eight exports. The installation is handled by exports 4 and 5. Export 4 finds a suitable process for injection. Export 5 is the actual installation routine; it drops the load point driver into the %System%\Drivers\folder with a name defined by the installation configuration file. Next a service is created so the driver is loaded every time windows start. The final step involves the encryption of the main DLL and also reading configuration file, encrypting it, and placing it in the %Windir%\inf\folder. After the complete installation, three files are left on the disk, the driver, the encrypted DLL and the encrypted main DLL configuration file.

b. Propagation:

Stuxnet can propagate using a number of methods. It can propagate by infecting removable drives and also by copying itself over the network using different means. It can also propagate by copying itself to Step 7 projects where it auto executes every time the project is opened.

Network propagation routines – Export 22 is responsible for the majority of the network propagation routines. This export builds a 'Network Action' class that consists of 5 subclasses, each of which are responsible for separate methods of infecting a remote host. The functions of the 5 subclasses are:

- Peer-to-peer communication and updates – this works by installing an RPC server and client. When a computer is infected, it starts listening for connections. Other compromised computers can connect to the RPC server and ask for the version installed on the remote machine. If

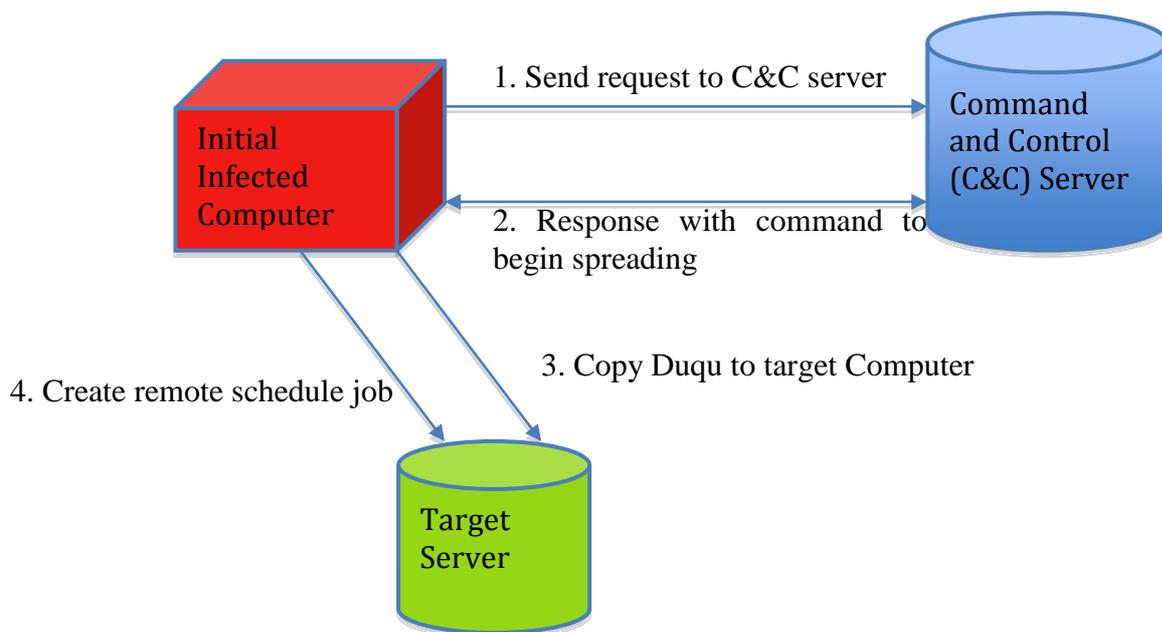
the remote version is newer then the local computer makes a request for the same and updates itself. If the version is older, then the local computer will prepare a copy of it and send it over to update the remote machine. This makes updating the compromised computers on a network fairly easy and spreads throughout the network. Stuxnet is actually a .dll file so in order to send an executable version of itself, it needs to build an executable version of itself by reading in a template .exe from resource and populating it with all the details of the latest configuration data and information.

- Infecting WinCC computers – this class is responsible for connecting to a remote server running the WinCC database software. When it finds a system running this software it connects to the database server using a password that is hardcoded within the WinCC software. After connection it first, sends malicious SQL code to the database that allows a version of Stuxnet to be transferred to the computer running WinCC software and executes it. This infects the computer that is running the WinCC database. Second Stuxnet modifies the existing view adding code that is executed each time the view is accessed. Stuxnet sends an SQL statement that generates a table with an inserted binary value. This binary value is a hex string representing the main Stuxnet dll as an executable file and an updated configuration block. The dll can be extracted from the .cab file when ran locally on a computer with WinCC installed. This allows Stuxnet to execute itself and ensure that it remains resident.
- Propagation through network shares – Stuxnet can also be spread to available network shared through either a schedule job or using Windows Management Instrumentation. Stuxnet will enumerate all user accounts of the computer and the domain, and try available resources either using the user's credential token or using WMI operations with the explorer.exe token to copy itself and execute on the remote share.
- MS10-061 Print Spooler zero-day vulnerability – this is a zero day print spooler vulnerability patched by Microsoft in MS10-061. This allows writing of files to the %System% folder of vulnerable systems. The actual code to carry out the attack is stored in resource 222. This loads the DLL and an IP address along with a copy of the worm. The loaded DLL in turn calls the export 1. Stuxnet can execute itself after copying itself to remote computers as %System%\winsta.exe through the print spooler.
- MS08-067 Windows Server Service vulnerability – this can be exploited by connecting over SMB and sending malformed path string that will allow arbitrary execution. Stuxnet uses this to copy itself to unpatched remote systems.
- Removable drive propagation – Stuxnet uses inserted removable drives for propagation. Operators often exchange data with other computers using removable drives. Stuxnet uses two methods to spread to and from removable drives- one method, LNK Vulnerability that allows auto execution when viewing the removable drive and the other using the autorun.inf file, which instructs Windows to automatically execute a file on the removable drive when the drive is inserted.

- Step 7 Project File infections – Export 16, calls the Export 2, which is used to hook specific APIs that are used to manage a WinCC/Step7 project. Listed files inside the projects with extensions .tmp, .s7p, or .mcp receive special processing. When the infected project is loaded Stuxnet can execute itself.

W 32.Duqu – methods used for propagation can differ as the behavior is not hard-coded into the threat, but actively conducted by the attackers. When Duqu compromises a target network, the threat contacts a Command and Control (C&C) server. One of the files downloaded by Duqu from the C&C server is a keylogger. This helps the attacker to intercept passwords for the local network and any other services accessed by the victim. Other files downloaded from the C&C server allow the attacker to survey the local network and find additional network servers and clients. After locating various computers on the local network and acquiring the passwords, the attacker can begin the process of spreading Duqu across the network.

Duqu will first have to copy itself onto the target computer over a shared folder. By using the credentials intercepted by the keylogger, the infected computer can authenticate itself to the target. Next it can trigger execution of the copied sample on the target computer by creating a schedule task on the target computer, which in turn executes copied version of the Duqu. This indicates that Duqu is running on the target computer. This infected computer does not connect to the C&C server to receive any further commands, instead checks its configuration file as it loads. In order to make the newly infected computer to receive commands from the infecting computer, the configuration file instructs it to connect back to the infecting computer. The diagram below depicts Duqu propagation.



c. Command and Control:

Stuxnet - After Stuxnet is installed and information of the system is gathered, it contacts the command and control server to send some basic information of the compromised system to the attacker via HTTP. Stuxnet used two command and control servers primarily. The addresses of which are `www[.]mypremierfutbol[.]com` and `www[.]todaysfutbol[.]com`. The servers were present in Malaysia and Denmark respectively; however they have been redirected to prevent the attacks. Stuxnet can update itself with new command and control domains. Once system data is gathered the payload contains the machine and the domain name along with the OS information. Attackers can realize if either ICS programming software Siemens step7 or WinCC is running on the machine. The payload data is XOR-ed with a byte value of 0XFF and sent to the target server. The target process can be an existing Internet explorer process (`iexplore.exe`). The information is sent to one of the malicious Stuxnet server specified in the Configuration data block. Two legitimate web-servers referenced in the configuration data block are queried to test the network connectivity: `www.windowsupdate.com` and `www.msn.com`. Once the test is passed the network packet is built. The payload is then XOR-ed with a static 31-byte long byte string that is found inside Stuxnet. The result is hexified in order to transform binary data to ascii string. The payload is then sent to the aforementioned urls, as the 'Data' parameter. The malicious Stuxnet server processes the query and sends a response to the client. The response payload is present in the 'Context' section of the HTTP. This response payload is binary data, encrypted with 31-byte long XOR key. The server response is decrypted and depending upon the command byte, the payload module is either loaded in the current process, or in a separate process via RPC. This allows Stuxnet backdoor functionality, to upload and run any code on an infected machine.

Duqu – Duqu uses HTTP and HTTPS to communicate with the command and control (C&C) server. Currently known C&C servers include 206.183.111.97 hosted in India, 77.241.93.160 hosted in Belgium, and 123.30.137.117 hosted in Vietnam. All of which are inactive. The C&C servers contacted by Duqu are proxies redirecting connections to either the true C&C server or yet another proxy. They are virtual machines running the Linux OS. The main DLL extracts an embedded file, Resource 302. The payload DLL is found in the `.zdata` section of Resource 302, along with its configuration data. The purpose of the `.zdata` DLL is command and control functionality, which allows downloading and executing updates and additional payload modules. The command and control protocol is custom protocol using any of the five following methods:

- Encapsulated in HTTP over port 80
- Encapsulate in HTTP over port 80 using a proxy
- Directly over port 443
- Encapsulated in HTTPS over port 443
- Encapsulated in SMB primarily for peer-to-peer command and control.

Through the command and control servers the attackers can download an infostealer that can perform actions such as enumerating the network, recording keystrokes, and gathering system information. The information is logged in an encrypted and compressed local file, and then exfiltrated out. The command and control functionality can download new executables and either execute them directly in memory or write them to disk. When written to disk, they are saved encrypted using a file name defined in the configuration data. Typical filenames are %Temp%\~[VARIABLE].tmp. The command and control protocol primarily downloads or uploads .JPG files. Additional encrypted data is appended to the dummy .JPG files for exfiltration. Using .JPG files is to obfuscate network transmissions. When using HTTP, the client sends repeated GET requests to the server. The server replies with modules to execute. To return data, Duqu uses a POST and sends a small blank JPG file appended with the data to send to the server. When using HTTPS, the same happens, except within an encrypted HTTPS session.

Peer-to-peer command and control protocol is not used by default but is configured for use in cases where a computer cannot reach the external C&C server. The attackers set a byte in the configuration file to one, and instead of specifying an external IP address, provides an IP address or string representing a remote resource. The peer-to-peer command and control protocol can use HTTP or Inter Process Communication over Server Message Block (SMB), also known as named Pipes. When sending traffic directly to port 443 or named pipes, no encapsulation is used and the Duqu protocol traffic is sent directly with the addition of eight initial bytes, which is a validation key. The command and control protocol is a reliable protocol, which implements fragmentation, reordering, and handles duplicate and missing packets via sequence and ACK numbers.

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.

Answer:

The 21 steps to improve cyber security in SCADA systems clearly identify key aspects of the system that need protection, and why cyber security of SCADA systems is of prime importance.

- a. The 21 steps if properly implemented could prevent Stuxnet attacks in the following ways :

Step 1: By identifying all possible connections to the SCADA network, the paths and situations under which the virus can propagate can be curtailed and tighter security can be levied on them. A comprehensive review of connected

equipment and their security can be the essence of preventing another virus attack.

Step 2: The isolation of SCADA network from any other network reduces the risk of creating a pathway to or from the internet. Thereby leaving the system contained. Data transfer mechanisms to business networks should be more secure by Utilizing Demilitarized Zones and Data Warehousing.

Step 3: If tight security is maintained on any remaining connections to the SCADA network by conducting penetration testing and vulnerability analysis and used in conjunction with risk management processes it could provide robust protection strength. The use of Firewalls and IDS could help keep the system free of containment.

Step 4: Once disabled or unnecessary services are removed they reduce the risks of virus attacks. The virus could be logged in such a service, and anybody executing them could release it from containment. Examples of services to remove include Email Services, and those to disable include remote maintenance.

Step 5: Sometimes relying solely on proprietary protocols or factory settings, could lead to an open back door going unnoticed.

Step 6: If all the security measures as required by the device and provided by the system vendor, are implemented, it could help in securing the system in a better way. For this all the security parameters should be set to maximum.

Step 7: Once devices that could act as backdoors like Modems, wireless connections are controlled strongly, their activities could be monitored for discrepancies.

Step 8: Implementation of a 24 hours a day incident monitoring mechanism software would be capable of immediately notifying any difference in behavioral patterns.

Step 9: Performing technical Audits of SCADA devices and Networks would help in identifying vulnerabilities to determine their risks, Track corrective actions and study trends, Retest the systems after corrective action to see if vulnerabilities are eliminated and also scan non production environments to identify any potential problems.

Step 10: Conducting Physical security surveys help in identifying inventory access points at facilities connected to the SCADA system and also identify and access any source of information.

Step 11: Creation a SCADA red team would be an advantage in having experience people provide information about identified potential attacks and evaluate risks of malicious insiders. Their information would be useful in risk assessment processes.

Step 12: Making the roles and responsibilities of all personnel clear to them is essential in ensuring the proper authority carries out the assigned responsibility.

Step 13: Developing and documenting robust information security architecture is essential to establish an effective strategy to protect information on the system that is critical and needs extra protection. It is essential to deploy these procedures at a very early stage to prevent any virus vulnerabilities that include packet sniffing.

Step 14: Performing a rigorous risk management process including a stage of assessment and a subsequent protection strategy is essential to make routine changes due to rapidly changing technology and new forms of Stuxnet that may come to effect like Duqu.

Step 15: Utilizing technical and Administrative controls to mitigate threats from identified risks help in preventing single points of failure and containing the impact of security incidents. Protecting systems from other systems in the same layer is essential to prevent spread of the virus

Step 16: When background checks are conducted and network privileges are limited, the virus spread from insiders can be prevented. User Agreements , Notifications and Warning Banners are key components in this alert system.

Step 17: When the hardware and software configurations are covered by processes that evaluate and control any changes , the network can be made more secure.

Step 18: Conducting self-Assessment routines increasing ability to identify issues conduct root cause analysis and implement corrective actions immediately. This also provides information of how effective the cyber security policies and technical implementations are.

Step 19: When system backups and disaster recovery plans are deployed, they allow for information backup and rapid recovery from any emergency. When the personnel are familiar with these plans, they are more in place to take quicker action.

Step 20: When the expectations of the cyber security programs are clearly identified by senior organizational leadership, the objectives of system security from viruses become much clearer. Also holding individuals accountable for their performance, adds an advantage to have everyone work their best.

Step 21: When personnel are trained not to release data or information to unauthorized people, they become more aware of the risks of their actions. Training on this aspect is essential in preventing any insider attack or outsider for that matter. For example : Divulging information about physical and logical locations of computers and networked systems.

- b. The steps identified are efficient in preventing future attacks, but are missing a few points. There is a need to implement key reminders that there is a chain that needs to be followed. It is essential to be on the defensive, watch for cyber-attacks, have a plan, know how to recover, document and report the attack and again be on the defensive. Considering that Stuxnet initially propagated through a pen drive,

Media Protection should be defined. This should define the use of both digital media and non digital Media. Also steps should be provided at the design level of SCADA systems, since these viruses can compromise the control actions of PLC.