Common Control System Vulnerability

November 2005

Prepared by Idaho National Laboratory

Homeland Security

Control Systems Security Center

Control Systems Security Center

Common Control System Vulnerability

Trent Nelson

November 2005

Idaho National Laboratory

Idaho Falls, Idaho 83415

Control Systems Security Center

Common Vulnerability

Introduction

The Control Systems Security Center (CSSC) and National SCADA Test Bed (NSTB) programs have discovered a vulnerability common to control systems in all sectors that allows an attacker to penetrate most control systems, spoof the operator, and gain full control of targeted system elements. This vulnerability has been identified on several systems that have been evaluated at INL, and in each case a 100% success rate of completing the attack paths that lead to full system compromise was observed. Since these systems are employed in multiple critical infrastructure sectors, this vulnerability is deemed common to control systems in all sectors.

Modern control systems architectures can be considered analogous to today's information networks, and as such are usually approached by attackers using a common attack methodology to penetrate deeper and deeper into the network. This approach often is composed of several phases, including reconnaissance, traffic analysis, profiling of vulnerabilities, launching attacks, escalating privilege, maintaining access, and covering evidence. With irrefutable proof that external attacks can lead to a compromise of a computing resource on the organizations business local area network (LAN), access to the control network is usually considered the next phase in the attack plan. Once the attacker gains access to the control network through vulnerabilities in the business LAN, the second phases of reconnaissance begins with traffic analysis within the control domain. Thus, the communications between the workstations and the field device controllers can be monitored and evaluated, allowing an attacker to capture, analyze, and evaluate the commands sent among the control equipment. Through manipulation of the communication protocols of control systems (a process generally referred to as "reverse engineering"), an attacker can then map out the control system processes and functions. With the detailed knowledge of how the control data functions, as well as what computers and devices communicate using this data, the attacker can use a well known Man-in-the-Middle attack to perform malicious operations virtually undetected.

The control systems assessment teams have used this method to gather enough information about the system to craft an attack that intercepts and changes the information flow between the end devices (controllers) and the human machine interface (HMI and/or workstation). Using this attack, the cyber assessment team has been able to demonstrate complete manipulation of devices in control systems while simultaneously modifying the data flowing back to the operator's console to give false information of the state of the system (known as "spoofing"). This is a very effective technique for a control system attack because it allows the attacker to manipulate the system and the operator's situational awareness of the perceived system status. The three main elements of this attack technique are: 1) network reconnaissance and data gathering, 2) reverse engineering, and 3) the Man-in-the-Middle attack.

Network Reconnaissance and Data Gathering

Once access has been obtained on the control system network, be it via the business LAN or some other plausible attack vector (vendor channel, wireless, dial-in access, etc), network reconnaissance is used to gather the information required to develop a plan of attack. By passively scanning, listening, and gathering communication traffic (i.e., protocols), the attacker is able to obtain an initial inventory regarding the architecture components in the control network, as well as direct insight into the communications used by the control devices on the network. After enough information has been gathered, the attacker can begin decoding and assessing the system information flow. This process of passively listening to network traffic is often referred to as 'sniffing'.

Each system evaluated suggested that in order to communicate with the end-point field devices, the application always communicated directly with the device-specific controllers. This identified a critical path on the flow of system information between the controllers and/or field devices and the workstation. Decoding the communications within this flow of information is the key to understanding the system and more importantly, verifying targets on the control network. In order to break the communication layer, the control network traffic had to be monitored and dissected to develop a greater understanding of how the components communicate. From an attack perspective, if a resource has been compromised on the control network, and the attacker has escalated his privileges to an authoritative level, then the task of monitoring control-centric traffic and collecting it undetected for back-end analysis becomes trivial.

Reverse Engineering

In the testing the assessment teams performed, the decoding of the communication (protocols) required reverse engineering. To reverse engineer a protocol, communication packets are captured by the attacker using the compromised machine on the control network and dissected to identify the inner working of the communications. Each packet contains all the required components to operate and control the field devices. The critical aspect of each protocol is to understand how the packet is put together and identify which pieces (bits) within the packet are the commands for controlling the equipment. These pieces are identified through reverse engineering of the protocol, which allows the attacker the ability to manipulate each packet as required.

As control system information infrastructures once existed as closed networks, the threat of external attack was minimal, and it could be assumed that any data on the control network was authorized and permitted. Therefore, data sent to and from control devices and to the operator consoles was usually considered valid. Each control system network component could theoretically communicate with any other component without any verification of sender or receiver. Historically, the Internet has seen myriad attacks where an attacker has been able to capture (sniff) data, analyze it, change critical instructions in the payload of the captured data, and reinsert it back into the network. This new data, with possible harmful instructions, would be accepted by the target resource and command would be executed. This category of attacks are know as 'replay attacks', and the impact they can have in a control systems environment is an obvious concern. While it has been shown that an

attacker can see, collect, reverse engineer and change data payloads in control system traffic, the final task of successfully inserting the modified rogue traffic into the data stream requires that the information flow be uninterrupted. In order to use the information and insert the modified packets into the information flow, a Man-in-the-Middle attack must be carried out.

Man-in-the-Middle Attack

A Man-in-the-Middle attack requires the use of the address resolution protocol (ARP) and an in-depth understanding of the protocol to be manipulated. The ARP Man-in-the-Middle attack is a popular method used by an attacker to gain access to the network flow of information on a target system. This is done by attacking the network ARP cache tables of the controller and the workstation machines. Using the compromised computer on the control network, the attacker poisons the ARP tables on each host and informs them that they must route all their traffic through a specific internet protocol (IP) and hardware address (i.e., the attacker's machine). By manipulating the ARP tables, the attacker can insert his machine between the two target machines and/or devices.

The Man-in-the-Middle attack works by initiating gratuitous ARP commands to confuse each host (referred to as ARP poisoning). These ARP commands cause each of the two target hosts to use the Media Access Control (MAC) address of the attacker as the address for the other target host. When a successful Man-in-the-Middle attack is performed, the hosts on each side of the attack are unaware that their network data is taking a different route through the attacker's computer. The attacker's computer then needs to forward all packets to the intended host so the connection stays in sync and does not time out. Figure 1 illustrates a typical Man-in-the-Middle attack.


Figure 1. Generic man-in-the-middle attack

The Man-in-the-Middle attack is effective against any switched network because it effectively puts the attacker's computer between the two hosts. This means the hosts send their data to the attacker's computer, thinking it is the host to which they intended to send the data. After the ARP tables on both target hosts have been successfully poisoned, the program shuttles packets back and forth between the target hosts. This ensures that all of the current applications on the target hosts will continue to work properly. During the shuttling process, every packet destined for either target host is processed through the attacker's machine and can be manipulated (packet crafting) to send specific commands to each host.

Once the attacker has successfully inserted their machine into the information stream, he or she now has full control over the data communications and can carry out several types of attacks. As previously stated, one attack well suited for the compromise of control system operation is a replay attack. In its simplest form, captured data from the control/HMI is modified to instantiate activity when received by the device controller, and the reverse engineering that has occurred previously will empower the attacker to craft any instruction be it benign or malicious. When considering the activities an attacker will perform during a system compromise, one key element is to maintain covert activity and remove evidence of the attack wherever possible. Bearing in mind that cyber-based attacks on control systems are unique in that they are 'digital' attacks that manifest themselves in 'physical' actions, manipulation of the operator's information is vital to the success of the attack. Control of the information that is accessible by the operator is required to hide the attack. During the earlier data capture phase of the attack, data reflecting normal operations in the control systems are harvested and can be played back to the operator as required. This will ensure that the operator's console will appear to be normal and the attack will go unobserved as the information presented to the operator via the HMI. During this replay attack the attacker can continue to send destructive commands to the controller and/or field devices to cause an undesirable event while the operator is blind to the true state of the system.

Another attack that can be carried out with the Man-in-the-Middle attack is false messaging to the operator, and can take the form of a false negative or a false positive. This aspect of 'perception management' is used to indicate operational activities (good or bad) and thus influence the behavior of the operator who would, under normal conditions, see console activity reflecting the physical manifestation of the cyber attack on the actual devices. The attacker would simply send commands to the operator's console indicating a system change, and when the operator follows normal procedures and attempts to correct the problem, the operator's action would cause the undesirable event. There are numerable variations of how the modification and replay of control data can impact the operations of the system.

Mitigation Techniques

Protocol manipulation and the Man-in-the-Middle attack is a common hacking practice that can be used across control systems in all sectors. However, there are mitigation techniques that can be applied to secure the systems through MAC address locking, static tables, and encryption.

Mac Locking - The ARP Man-in-the-Middle attack requires the attacker to be connected to the local network or have control of a local computer on the network. Port security, also called MAC address locking, is one method to secure the physical connection at the end of each port on a network switch. The high-end corporate class network switches usually have some kind of option for MAC address locking. MAC address locking is very effective against a rogue individual looking to physically plug into the internal network. Without port security, any open network jack on the wall could be used as an avenue onto the corporate network. Port security locks a specific MAC address to a specific port on a managed switch. If the MAC address does not match, the communication link is disabled and the intruder will not be able to achieve his goal. Some of the more advanced switches have an auto resetting option, which will reset the security measure if the original MAC is returned to the port.

Although port security is not hacker proof, it does add a layer of added security to the physical network. The implementation of port security is not difficult, but it increases the complexity of normal maintenance functions such as plugging in a legitimate computer onto the network. It also protects the local network from employees plugging un-patched and out-of-date systems onto the protected network. This reduces the number of target computers a remote attacker can access. These security measures not only protect against attacks from external networks but provide added physical protection as well.

Static Tables - A smaller network that stays relatively static can attempt to implement statically coded ARP tables. Most operating systems have the capability to statically code all of the MAC addresses into the ARP table on each computer. Statically coding the ARP tables on each computer prevents the attacker from changing them by sending ARP reply packets to the victim computer. Static ARP entries are hard to maintain in small networks and impossible to maintain when the network is large and/or dynamic.

Encryption - As a longer term solution, systems need to be redesigned to include encryption between devices in order to make it very difficult to reverse engineer protocols and forge packets on control system networks. Encrypting the communications between the controller and workstations would shut the door on this type of attack, but would be expensive and difficult to retrofit to the existing systems. Along with encryption, monitoring for ARP poisoning provides an added layer of defense. There are several programs available, ARPwatch for example, that can monitor for changing MAC addresses through the ARP packets.