September/October 2010

Process Automation

Balancing security and safety with risk

From simplicity to complexity

Fast Forward

- Process control systems have almost universally moved from being a bespoke development to utilizing a Microsoft Windows platform.
- Malware writers have now set their sights on control system vulnerabilities and started to develop viruses.
- Patching vulnerabilities as they are found is often difficult, and operators often need to look at the risks and consequences.

By Graham Speake

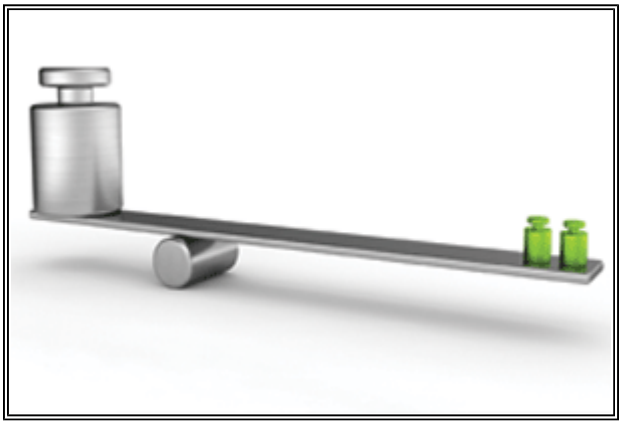

Initially when control and safety systems moved away from being hardwired and relay-based to computerized systems, vendors and asset owners were more interested in functionality than security. Typically, especially in high-risk environments in refineries and off-shore oil installations, the systems were standalone with a dedicated Safety Instrumented System. The advances in computer technology during the 1980s and 1990s caused a rapid shift from these proprietary systems to a typically Intel hardware with Microsoft-based operating systems. This was primarily driven by the end user to reduce costs and for standardization with the rest of the IT infrastructure. At that time, patches and updates to the base operating system came out sporadically from Microsoft, and the security aspects were rarely considered.

The new millennium saw this situation change rapidly, with Code Red and Nimda malware following quickly on from the events of 9/11 and a rapid re-evaluation of control system security was undertaken. Many companies assumed the control systems were still being operated as islands of automation—completely separate from the business network. In the majority of cases, this proved to be a misconception, and these “islands” were firmly attached to IT business mainland. What was thought of as a low security risk quickly escalated to being something definitely on the radar of risk managers.

To compound the risk problem, when these control systems were analyzed, it was quickly discovered the individual components would be at home in a standard IT environment—file servers, SQL (database) servers, web servers. The major differences were these were often unpatched, installed without anti-virus software, and configured with weak or no passwords—amongst numerous other security risks that would not normally be tolerated on servers located on the business network. The risks attached to these systems were rising by the minute.

Control system complexity

The last 10 years has seen a continuation in the rapid rise of the complexity of these control systems, with asset owners demanding more functionality and vendors competing to add new bells and whistles to differentiate their system from the competition. As these systems are typically marketed and sold to control engineers, these advancements are made to the parts of the system that will appeal to them—newer and better HMIs, faster control loops, wireless I/O, etc. These engineers often have little experience with IT infrastructures or cyber security, but are also unwilling to enlist the IT security experts who have scant knowledge of process systems. A few evangelists have emerged, from the vendor community as well as from users, but often security and complexity take second place to system operations.

As stated, the modern control system is likely to contain many standard IT server-type components. Complexity starts to grow when the location of these devices needs to be considered. Should they be installed on the process control network or are they better suited to the business network? It is likely that both scenarios will be deployed—sometimes by choice, but often by accident as no one has looked at the architecture as a whole. This is combined with the fact that often the process control engineers at a plant will approve system architecture designs put forward by third parties (consultants, Main Automation Contractors, etc.) but do not have in-depth knowledge of current IT practices to judge if these designs are suitable. The systems are often deployed by third parties that challenge the understanding of the actual users who make poor choices in terms of security.

The larger vendors have stepped up their efforts to ensure control systems manufactured today have far more security built into them. Anti-virus is considered to be a standard, for instance, with comprehensive advice on how to integrate this into the control system available from vendors. Patch updates for Windows are tested and approved for deployment a few days after Microsoft's "Patch Tuesday"—no small feat considering the number and diversity of control systems most vendors support. However, control systems consisting of units from multiple vendors, multiple versions of the software from Microsoft, and the system vendors can overwhelm the control engineer.

Measuring risk

There are numerous definitions and equations for "risk," and they change depending on what industry you are in and what discipline. Ask the question to someone in the Health, Safety, and Environment department of a company, and the answer will likely be different had the question been asked to a control systems engineer. A common risk equation can be defined as:

Risk = threat × vulnerability × target attractiveness × consequence

The problem with this is it is difficult to assign actual numbers to the equation. In addition, it would likely need to be recalculated on a regular basis due to external changes. For instance, Microsoft releases patches to their operating systems every month; these may directly affect the vulnerability of one or more of your control systems and hence drive the risk up. Asset owners have to decide if the increased risk is worth accepting, eliminating (by installing the patch), or mitigating by other means. The difficulties of testing and installing an operating system patch to a system located on a drilling platform miles from shore cannot be underestimated. This may be exacerbated by the fact that these control systems may be utilized for critical functions and therefore the availability of these systems is essential. The CIA triad (confidentiality – integrity – availability) is used widely within IT security circles, but the importance is often reversed in control systems as availability is usually far more important than protecting actual information.

Considering the "threat" variable, the recent Stuxnet incident (see sidebar "Virus targets Siemens control systems") has obviously made this increase. The vulnerability of control systems has been touted at

security conferences for many years, as a Google search will quickly testify to. However, these were generally vague presentations and alluded to unpatched system, poor/non-existent passwords. Stuxnet was specifically targeted at Siemens equipment, and there are likely to be more attacks of this kind in the future. The increase in news channels around hacking critical infrastructure (CBS “60 Minutes,” amongst others) obviously raised the target attractiveness index as well.

The final factor, consequences, needs to be very carefully considered. The disaster aboard the Deepwater Horizon drilling vessel in April 2010 showed how easily it is to underestimate the consequences of any incident. Apart from the human and environmental cost, the industry as a whole will likely be impacted by increased regulation.

Assessing security

ISA has life-cycle models for security and safety—defined in ISA99 and ISA84, respectively. As seen in the ISA99 sidebar, there is some momentum in defining a process to assess the security of control systems (and IT systems) using a scale very similar to that used in the safety industry. Asset owners can start to look at the security of their control systems today using another process taken from the safety industry—the HAZOP (Hazards and Operability analysis). This can be used as a basis for a control hazards and operability analysis tool (CHAZOP), and it is being used by a number of enterprises in the control systems space. The time taken to fully complete a CHAZOP cannot be underestimated, especially on a large interconnected system.

Within the security world, the phrase “defense in depth” is used widely, and it basically is a means to deploy numerous defensive mechanisms throughout the control systems to block (or at least delay) hackers trying to break into a system. Given enough deterrents, potential adversaries will hopefully move on to try another target. The number and sophistication of these deterrents will decrease the likelihood of an attack succeeding. The safety world can again be plagiarized as they use a layer-of-protection analysis extensively to estimate the likelihood of an event. While this does need to be modified for the security field, there is a good benefit in using it.

The terms safety and security are becoming synonymous, and indeed many European languages use the same word
for safety and security. Even in English, we will have a security system installed in our house, which is making us feel safer. When we assess security, we are also going to be indirectly assessing safety as well due to their similarities and
interconnections.

Balancing act

Assessing the hazards or risks in a system is only part of the process as a decision on what action to take on each risk needs to be made, and this is where an asset owner will need to get out the scales and start the balancing act. There will be a big difference to assessing and fixing the risks at the early stages of a project, ideally at the Factory or System Acceptance Test phase than if this is done to a fully deployed system. At the test phase, patches can be deployed easily, software added or removed, and problems in the overall architecture can usually be fixed without too much trouble. When you have deployed this system to a refinery or on an oil platform, the fix becomes more problematic.

Look at the problems an asset owner is facing every month when Microsoft releases the security updates for its current products. The vendors will approve and release the patches shortly afterwards and will comment on the criticality of the updates if these are not deemed to be relevant to their products. If they do fix a flaw in Windows that is being used, Microsoft’s assessment is usually accepted. However, these

systems will generally operate in a different world than Microsoft normally operates in. Does the security assessment (critical, high, medium, etc.) assigned by Microsoft apply here, or should it be higher or lower? Can other mitigations be put in place to reduce this risk and allow the patch deployment to be delayed?

As an example, assume Microsoft announces a critical vulnerability in Internet Explorer that could allow an attacker to gain information from your system if you connect to a malicious web site. For the enterprise with numerous laptops, this may be a serious issue and require the rollout of the patch. A control system will not, or should not, have systems that connect to the Internet or the enterprise Intranet, and therefore the risk from this vulnerability may not be rated as high. The business may see the deployment of this patch in the control system space as only needed to be undertaken as a regular change.

Not all vulnerabilities are as clear cut as the one outlined above. Suppose there is a vulnerability in Microsoft Server 200x that could allow an attacker to take control of the system, and you have a number of these deployed in your process control network. If there is no virus or worm currently “in the wild” (released and active on the Internet), you may take the decision to hold off on deploying the patch, as it requires a reboot and this will disrupt production. If a worm is released and is active prior to your scheduled deployment of the patch, have you a mechanism to reassess your initial evaluation of the risk?

The systems you normally see and can touch are usually on the list for patches and software updates and, hopefully, assessed regularly to determine the level of risk they have to the control system. There are often “hidden” systems that can be overlooked—those in laboratory instruments such as analyzers or control systems deployed in a subsea environment. The operating system deployed on them may not always be as easy to spot as they can boot up directly into a custom application. These systems can have glaring vulnerabilities that if exploited could cause major disruption to the rest of the control network unless they are properly segregated from the main systems.

Although we have mainly mentioned Microsoft Windows throughout this article, it should not be presumed that other operating systems are immune from vulnerabilities. While flavors of UNIX or Linux may not have viruses or worms targeted against them, there are vulnerabilities within these operating systems that can lead to exposure given the right circumstances. Assessment of these systems and the risks they pose to the infrastructure must also be undertaken.

Adding in security controls may also have an added benefit, as often these tools offer multiple benefits. They may reduce risk, but as we have seen, this is hard to quantify, but they also may reduce uncertainty in other ways, such as improving network or system reliability. If you can stop a worm entering or propagating on your control network, for instance, there will be no degradation in performance. If these security controls can be shown to correlate to process trend events, then the introduction of these security measures can improve security and operations. An increase in the operational performance of the plant may be easier to justify than a decrease of the security risk, which is hard to quantify.

Getting the balance right

Asset owners will have a battle for many years in ensuring control systems are deployed in a secure manner and are kept secure. The longevity of control systems in the past has left many outdated systems still in production, often not supported by the vendor. Regulation such as CFATS in the chemical sector is focusing asset owners’ minds onto security, but it is only a start. ISA is acting as one of the front runners in defining how security should be assessed and measured in this space with its work on the ISA99 standard. The next couple of years should see a number of additional documents in the ISA99 series being released, aiding asset owners and vendors. Control system security has to be more widely

practiced around all parts of the industry to ensure that all vendors produce secure products that can be installed and operated securely by asset owners. Ensuring the knowledge encompassed in ISA99 is disseminated to process control and IT engineers will start the process of everyone understanding the risks and ensuring that these risks can be mitigated.

ABOUT THE AUTHOR

Graham Speake is a principal systems architect at Yokogawa Electric Corporation, a major industrial automation supplier, where he provides security advice and solutions to internal developers and customers in many countries. His specialties include industrial automation and process control security, penetration testing, network security, and network design.

RESOURCES

- Safety first
  www.isa.org/intech/200906web
- Defense in cyberspace
  www.isa.org/intech/20080901
- *Protecting Industrial Control Systems From Electronic Threats*
  www.isa.org/link/Elec_Weiss_bk

Virus targets Siemens control systems

In 2010 Stuxnet, a virus specifically targeted against a control system was released into the wild; luckily it only infected a small number of systems worldwide. This malware spread via infected USB thumb drives and used a zero-day (unknown or undisclosed) vulnerability in Microsoft Windows. The malware had the ability to steal proprietary data and transmit it back to its command center and possibly modify or disrupt the actual control system itself.

The sophistication of Stuxnet is troubling, as its creators must have a detailed knowledge of Siemens WinCC and PCS7 products, as well as having the expertise to combine this with the previously unknown Windows vulnerability. The consequences could have been far more disastrous if the payload had been more malicious or a more widely deployed control system had been targeted. Luckily, the number of actual control systems infected was small (around six), and apparently non-production systems. Asset owners are usually reticent to discuss actual attacks that have succeeded in their plant, so the real number and consequence may never be known.

ISA99

ISA99 is a committee set up under the auspices of ISA, and the purpose of which is to develop and establish standards, recommended practices, and technical reports that will define and improve the cyber security of industrial automation control systems. The committee is made up of thought leaders and individuals from the vendor community, asset owners, and academia, all of whom have a desire to improve security. The documents that have emerged from the committee have all been accepted by

industry and are being implemented by companies around the globe.

One of the current tasks of the committee is to define a Security Assurance Level (SAL), which will be utilized in a similar fashion to how Safety Integrity Levels are used to look at process hazards. Briefly, the four SALs will be:

Level 1 – Basic. Minimal system protection. Little or no damage.

Level 2 – Moderate. Basic authentication, configuration management, etc. Damage limited to process interruptions or stoppages.

Level 3 – Significant. Some redundancy in networks and control, availability, etc. Injury up to death and dismemberment, loss of public confidence.

Level 4 – Extreme. Extensive redundancy in network and controls, strong authentication. Multiple fatalities, including members of the public.

Further information can be found at http://isa99.isa.org.

Related Files

process auto septoct2010