

The Myths and Facts behind Cyber Security Risks for Industrial Control Systems

Eric Byres, P. Eng.
Research Faculty – Critical Infrastructure Security
British Columbia Institute of Technology
Burnaby, BC, Canada

Justin Lowe
Principal Consultant
PA Consulting Group
London, UK

Abstract

Process control and SCADA systems, with their reliance on proprietary networks and hardware, have long been considered immune to the network attacks that have wreaked so much havoc on corporate information systems. Unfortunately, new research indicates this complacency is misplaced – the move to open standards such as Ethernet, TCP/IP and web technologies is letting hackers take advantage of the control industry’s ignorance. This paper summarizes the incident information collected in the BCIT Industrial Security Incident Database (ISID), describes a number of events that directly impacted process control systems and identifies the lessons that can be learned from these security events.

1 A Fine Balance

It is widely accepted in industrial security analysis that the security risk faced by an organization is a function of the both the *Likelihood of Successful Attack* (L_{AS}) against an asset and the *Consequence* (C) of such an attack [1]. The second variable, *Consequence*, while highly site specific, is generally the easiest to get an understanding of. Often it can be estimated in terms of financial loss, acute health effects or environmental impacts; concepts well understood from years of safety analysis of hazardous processes.

Estimating the *Likelihood of Successful Attack* is far more difficult. According to the American Institute of Chemical Engineers’ guidelines it is a function of three additional variables [2]:

Threat (T): Any indication, circumstance, or event with the potential to cause the loss of, or damage to an asset.

Vulnerabilities (V): Any weakness that can be exploited by an adversary to gain access to an asset.

Target Attractiveness (A_T): An estimate of the value of a target to an adversary.

These terms are more difficult to estimate, particularly with respect to cyber security.

This difficulty is largely because we have little reliable historical or statistical data to work with. The details of safety related incidents have been recorded for over a century, while cyber security incidents have less than two decades of occurrence, never mind record keeping. Furthermore, most organizations are

highly reluctant to report security incidents as they are viewed as potential embarrassments. In fact, many organizations have denied that there even is a risk to industrial systems from cyber attack. For example, as recently as March 2002 an article in CIO Magazine entitled “Debunking the Threat to Water Utilities” stated there was no credible risk to SCADA systems from a network-based attack:

“Most public utilities rely on a highly customized SCADA system. No two are the same, so hacking them requires specific knowledge.”[3]

Yet this flies in the face of a number of well documented cyber-related incidents such as the Slammer Worm infiltration of an Ohio Nuclear plant and several power utilities [4] [5] and the wireless attack on a sewage SCADA system in Queensland Australia [6]. Clearly the merging of common information technologies such as Ethernet, Windows and Web Services into industrial controls technology has removed the dubious protective barrier of “*security by obscurity*”.

There is obviously some security risk faced by industrial control systems and, as difficult as it is to estimate, we still need to understand it. We can’t ignore the risk and yet we also can’t afford the infinite cost of perfect security. Sound business practice requires that we balance off the cost of measures to mitigate a risk, with the potential cost of an event occurring. To do so we need to understand the variables at play in defining the cyber security risk for an industrial facility. Furthermore, we need to continuously monitor the

risk variables to determine if they are changing. To be effective from both a technical and cost perspective, our mitigation response must adapt to changes in Threats, Vulnerabilities or Target Attractiveness. As we will show in this paper, the first two of these variables are changing rapidly and demand attention. We will also show that the consequences of successful attacks are not insignificant.

2 The BCIT Industrial Security Incident Database (ISID)

The British Columbia Institute of Technology (BCIT) maintains an industrial cyber security incident database, designed to track incidents of a cyber security nature that directly affect industrial control systems and processes. This includes events such as accidental cyber-related incidents, as well deliberate events such as external hacks, Denial of Service (DoS) attacks, and virus/worm infiltrations.

Data is collected through research into publicly known incidents (such as the Australian sewage spill) and from private reporting by member companies that wish to have access to the database. Each incident is investigated and then rated according to reliability on a scale of 1 to 4 (1=Confirmed, 2=Likely but Unconfirmed, 3=Unlikely or Unknown, 4=Hoax/Urban Legend). Figure 1 shows a typical data entry screen.

Fig. 1 Typical ISID Security Incident Entry Screen

The data collected includes:

- Incident Title
- Date of Incident
- Reliability of Report
- Type of Incident (e.g. Accident, Virus, etc.)
- Industry (e.g. Petroleum, Automotive, etc.)
- Entry Point (Internet, Wireless, Modem, etc.)
- Perpetrator
- Type of System and Hardware Impacted
- Brief Description of Incident
- Impact on Company
- Measures to Prevent Reoccurrence
- References

At the time of this paper 41 incidents had been investigated and logged in the database, with 11 incidents still pending investigation. Of these, 7 were flagged as hoax/urban legend, and removed from the study data, leaving 34 events of sufficient quality for statistical analysis. Figure 2 shows the trend of events between 1995 and 2003. It appears that there is a sharp increase in events occurring around 2001. This may be indicative of an actual increase in attacks or the result of the increased efforts to collect data.




Fig. 2 Security Incidents between 1995 and 2003

Discussions with operators of traditional business crime reporting databases indicate that the typical incident database collects less than one in ten of the actual events. Ten incidents were collected for 2003, so it is likely that industry is experiencing at least 100 incidents per year at the present time. If nothing else, one conclusion we can draw from this statistic is that there is a security problem and it may be more widespread than most engineers believe.

3 The Good Old Days (Threat Sources from 1980 to 2000)

The data was next analyzed for incident type to get an idea of the threat sources. Figure 3 shows the breakdown of 13 incidents between the years 1982 and 2000. Incidents were almost evenly split between accidental, internal and external sources, with only 31% of the events being generated from outside the company. Accidents, inappropriate employee activity and disgruntled employees accounted for most of the problems.

These statistics correlate well with the numbers being expressed by security researchers in the traditional IT world at the time. For example, this statistic was widely quoted in 2001:

“A study by the FBI and the Computer Security Institute on Cybercrime, released in 2000 found that 71% of security breaches were carried out by insiders.” [6]

4 A New Can of Worms (Threat Sources from 2001 to 2003)

The study team then produced the same graph for the period 2001 to 2003, as shown in Figure 4. Externally generated incidents account for 70% of all events, indicating a surprising and significant change in threat source.

Interestingly, the IT world appears to be experiencing the same shift. For example:

“Deloitte & Touche’s 2003 Global Security Survey, examining 80 Fortune 500 financial companies, finds that 90% of security breaches originate from outside the company, rather than from rogue employees.

‘For as many years as I can remember, internal attacks have always been higher than external,’ said Simon Owen, Deloitte & Touche partner responsible for technology risk in financial services.

‘60 to 70 per cent used to be internally sourced. But most attacks are now coming from external forces and that’s a marked change.’”[7]




Fig 3: Security Incidents by Type 1982 -2000




Fig 4: Security Incidents by Type 2001-2003

Why did the threat source change so significantly in such a short period of time? We have no definite answers, but there are a few possibilities to explain the

impact on industrial control systems. First the emergence of automated worm attacks starting with Code Red¹ in July 19, 2001 have meant that many of the intrusions have become non-directed and automated. The control system has become just a target of opportunity rather than a target of choice.

Second, common operating systems (e.g. Windows 2000 or Linux) and applications (e.g. SQL Server) now dominate the Human Machine Interface (HMI), engineering workstation and data historian systems. These often come configured more appropriately to business requirements and are vulnerable to a wide variety of common IT attacks and viruses. Issues with applying patches to these critical systems exacerbate the problem.

Finally the increasing interconnection of critical systems has created interdependencies we haven't been aware of in the past. As the Slammer Worm incident documented by the North American Electric Reliability Council illustrates, Internet incidents can indirectly impact a system that doesn't use the Internet at all. In this case the power utility used frame relay for its SCADA network, believing it to be secure. Unfortunately the frame relay provider utilized a common Asynchronous Transfer Mode (ATM) system throughout its network backbone for a variety of its services, including commercial Internet traffic and the SCADA frame relay traffic. The ATM bandwidth became overwhelmed by the worm, blocking SCADA traffic to substations [8].

Regardless of the reasons, the threat sources are moving from internal to external and this needs to be taken into consideration in the risk assessment process. Determining the actual perpetrators and their probability of attack is currently beyond the ability of the database, but security risk analysts are advised to look at governmental studies of threats to critical infrastructure to obtain some possible threat estimates. A good starting place is the UK National Infrastructure Security Co-ordination Centre's (NISCC) report "The Electronic Attack Threat to Supervisory Control and Data Acquisition (SCADA) Control & Automation Systems" [9].

5. The Backdoor into the Plant

If the threats are becoming increasingly external, then this begs the question, "How are they getting in?" While Internet connections maybe the obvious source, it isn't the only one. For example, database records show that the Slammer Worm had at least four differ-

¹ While Code Red was not the first non-email based worm, it appears to be the first to have had significant penetration into industrial systems.

ent infiltration paths in the control systems it impacted:

- The Davis-Besse nuclear power plant process computer and safety parameter display systems via a contractor's T1 line;
- A power SCADA system via a VPN;
- A petroleum control system via a laptop;
- A paper machine HMI via a dial-up modem.

To answer this question, the study team analyzed the "Point of Entry" data for each of the incidents in the database. The incidents were divided into two groups, namely internal incidents (14) and external incidents (25). Figures 5 and 6 show the statistics for these two groups respectively.




Fig 5: Internal Security Incidents by Entry Point




Fig 6: External Security Incidents by Entry Point

For the internal incident, the business network is the major source. Direct physical access to the equipment was also significant. For the external event, the Internet was a major source, but dial-up connections, VPNs, telco networks, wireless systems and 3rd party connections were all contributors. The obvious conclusion is that there are many routes into a system as complex as a modern SCADA or control system. Focusing on a single intrusion point with a single solution (such as the Internet firewall) is likely to miss many possible attack points.

6. The Consequences of Industrial Cyber Attack

Assessing the consequences of industrial cyber attack is not simply a case of assigning a financial value to an incident. Although there are obvious direct impacts which may be easily quantifiable financially (e.g. loss of production or damage to plant), other consequences may be less obvious. For most companies the impact on reputation is probably far more significant than merely the cost of a production outage. The impacts of health, safety or environmental incidents could be highly detrimental to a company's brand image. Even impacts such as minor regulatory contraventions may in turn affect a company's reputation, and threaten their licence to operate.

For most of the incidents reported in the database, the contributors have been unable (or unwilling) to provide a financial measure of the impact of the industrial cyber attack - in fact only 30% have been able to provide such an estimate. However, although the sample data is not large, it does seem significant that nearly 50% of reported incidents, where a financial impact estimate has been given, have led to sizeable financial losses (>\$1M).

Potentially more significant is the nature of the impacts of the attack. 41% reported loss of production while 29% reported a loss of ability to view or control the plant. Fortunately human impacts have been small with only one unconfirmed (and possibly unreliable) report of loss of life. Overall the reported incidents clearly show that the most likely consequences of industrial cyber attack are loss of view of, or ability to control, the process.

The likely impact of being unable to view or control the process or system is an increased reliance on emergency and safety systems. Traditionally these systems have been totally independent of the main control system and generally considered 'bullet proof'. However, mirroring the trend in the design of the main control systems, these emergency systems are also becoming based on standard IT technologies (such as TCP/IP). They are increasingly being connected to or combined with the main control system, increasing the potential risk of common mode failure of both the main control system and the safety systems. Consequently, in the future, the systemic risks of cyber attack need to be considered in the design of not just the control systems, but also the safety systems.

7 A Brave New World

Looking forward, the study team sees nothing to indicate these trends are likely to reverse in the near fu-

ture. In fact, if anything the situation is likely to get worse. The hacking community is becoming increasingly aware of SCADA and process systems and is beginning to focus their attention on them. For example, the following presentation was reported at the Brum2600 Blackhat Conference, held in Birmingham, UK in October 2003:

“Things started to get a little more interesting... The talk was titled ‘How safe is a glass of water.’ It was a detailed breakdown of the RF systems that are used by water management authorities in the UK and how these systems can be abused, interfered with and generally messed.”[10]

Six months earlier, a presentation was given at the CanSecWest Conference detailing how to attack embedded operating systems used in routers, printers and cell phones [11]. These same embedded operating systems are used in modern SCADA and controls equipment. These presentations indicate that the hacking community is beginning to develop both the interest and the technical expertise to deliberately attack control systems.

8 Some Conclusions for Industrial Cyber Security

The above analysis indicates that there is a clear shift in the source of cyber attacks on industrial control systems (the Threats). Threats originating from outside an organization are likely to have very different attack characteristics to internal threats. Thus companies may need to reassess their security risk model and its assumptions.

In addition, the variation in the infiltration paths indicates a wide variety of vulnerabilities available to the attacker. Considering the difficulty of closing off all of these avenues, it would be wise to assume there will be boundary breaches and harden the equipment and systems on the plant floor to withstand possible attack. In effect, companies need to deploy a “defense in depth” strategy, where there are multiple layers of protection, down to and including the control device.

Achieving a defense in depth solution for industrial systems will require at least four steps. On the system design side, it is recommended that more internal zone defenses and more intrusion detection be deployed. Companies may also need to re-evaluate boundary security in terms of all possible intrusion points and not just focus on the obvious connections such as the business-process link. A single firewall between the business network and control system network is likely to miss many intrusions and will offer little security once the attacker is inside the control system network.

From the control system manufacturers’ side, SCADA and automation devices need to undergo security robustness design and testing prior to deployment in the field. SCADA & control protocols should also be improved to include security features. Currently most devices appear to be highly vulnerable to even minor attacks and have no authentication/authorization mechanisms to prevent rogue control.

Failure to adapt to the changing threats and vulnerabilities will leave the controls world exposed to increasing cyber incidents. The result could easily be loss of reputation, environmental impacts, production and financial loss and even human injury.

9 Literature

- [1] American Institute of Chemical Engineers (AIChE) Center for Process Safety, “Guidelines for Managing and Analyzing the Security Vulnerabilities of Fixed Chemical Sites”, August 2002
- [2] *ibid*, AIChE, pp 10 -12
- [3] Berinato, Scott; “Debunking the Threat to Water Utilities”, CIO Magazine, CXO Media Inc., March 15, 2002
- [4] “NRC Information Notice 2003-14: Potential Vulnerability of Plant Computer Network to Worm Infection”, United States Nuclear Regulatory Commission, Washington DC, August 29, 2003
- [5] “SQL Slammer Worm Lessons Learned For Consideration By The Electricity Sector”, North American Electric Reliability Council, Princeton NJ, June 20, 2003
- [6] Stephanou, Tony; “Assessing and Exploiting the Internal Security of an Organization”, The SANS Institute, March 13, 2001
- [7] Nash, Emma; “Hackers bigger threat than rogue staff”, VNU Publications, May 15, 2003, <http://www.vnunet.com/News/1140907>
- [8] *ibid*, NERC, p.1
- [9] “The Electronic Attack Threat to Supervisory Control and Data Acquisition (SCADA) Control & Automation Systems”, National Infrastructure Security Co-ordination Centre (NISCC), UK, July 12, 2003
- [10] “We have your water supply, and printers’ – Brumcon report”, The Register, October 20, 2003
- [11] FX, “Attacking Networked Embedded Systems” CanSecWest Conference, Vancouver, May 2003