SAND2015-0813C

Framework for Identifying Cybersecurity Risks in Manufacturing

Margot J. Hutchins[1], Raunak Bhinge[2], Maxwell K. Micali[2], Stefanie L. Robinson[2], John W. Sutherland[3], and David Dornfeld[2]
[1]*Sandia National Laboratories, Livermore, CA, USA*
[2]*University of California, Berkeley, CA, USA*
[3]*Purdue University, West Lafayette, IN, USA*
mjhutch@sandia.gov, raunakbh@berkeley.edu

Abstract
Increasing connectivity, use of digital computation, and off-site data storage provide potential for dramatic improvements in manufacturing productivity, quality, and cost. However, there are also risks associated with the increased volume and pervasiveness of data that are generated and potentially accessible to competitors or adversaries. Enterprises have experienced cyber attacks that exfiltrate confidential and/or proprietary data, alter information to cause an unexpected or unwanted effect, and destroy capital assets. Manufacturers need tools to incorporate these risks into their existing risk management processes. This paper establishes a framework that considers the data flows within a manufacturing enterprise and throughout its supply chain. The framework provides several mechanisms for identifying generic and manufacturing-specific vulnerabilities and is illustrated with details pertinent to an automotive manufacturer. In addition to providing manufacturers with insights into their potential data risks, this framework addresses an outcome identified by the NIST Cybersecurity Framework.

Keywords: Cybersecurity, Risk, Data Analytics, Advanced Manufacturing

1 Introduction

A globally-interconnected digital information and communications infrastructure, which may be referred to as “cyberspace” supports the functionality of almost every system in the modern world [Office of the White House, 2009; Clark et al., 2014]. Economic, transportation, communication, energy and security systems, among other systems, are highly reliant on information and communications technology (ICT). Small businesses would be unable to continue their day-to-day operations without access to the current cyber infrastructure, and such a situation would be even more profound on Wall Street. ICT provides mechanisms for more efficient and convenient transfer of information than ever before; however, it also brings substantial risks. The retailer Target believes as many as 70 million customers were affected by a data breach in late 2013; Home Depot estimated that up to 56 million credit cards may have been compromised in a 5-month-long breach in 2014; and approximately 76 million households and 7 million small businesses were affected by the cybersecurity breach at JP Morgan in the summer of 2014 [Forbes, 2014; Wall Street Journal, 2014a; Wall Street Journal, 2014b]. More recently, an attack on Sony Pictures Entertainment, in which an alleged tens of terabytes of data were leaked including executive salaries and pirated films, resulted in economic sanctions by the U.S. government against senior North Korean officials [New York Times, 2014; New York Times, 2015].

As early as 1994 [Fortune, 1994], proponents of a digital factory touted the benefits of software (and data) driven manufacturing. This has become a reality in the ensuing 20 years with the prevalence of high-speed Internet, interoperability, advanced computation, and computers. The President's Council of Advisors on Science and Technology (PCAST) define advanced manufacturing as "a family of activities that depend on the use and coordination of information, automation, computation, software, sensing, and networking" [PCAST, 2011]. As the conventional manufacturing sector transitions into the advanced manufacturing sector, it becomes progressively more connected and data-driven. As a result, it benefits greatly from the enabling power of ICT. The ability to transfer data quickly from one physical location to another has allowed manufacturers to decrease the time required to bring a product to market. Demand for a product can be more easily identified, and the product itself can be more readily conceived, engineered, manufactured, delivered, and tracked. Additionally, with the increased ability to collect and share information, it has become easier to characterize the product, and the materials and processes used to create it. This facilitates improved product quality, increased process efficiency, and reduced resource consumption. However, manufacturers assume some risks associated with the connectivity of the systems within their enterprises and to the outside world. In fact, there is evidence that manufacturing is at high risk for spear phishing attacks, directed attacks that embed malware in target computers, which may be delivered via email or other targeted means [NDIA, 2014; Symantec, 2013]. Spear phishing is a cause for concern because these attacks may seek out personal information in a highly sophisticated manner. For perspective, stolen or weak credentials were exploited in 76% of all network intrusions that occurred in 2013 [Verizon, 2014].

Manufacturing is not exempt from the many cybersecurity challenges faced by other critical infrastructure sectors, which vary from nuisance vandalism to theft of intellectual property to destruction of capital assets [FBI, 2014; Cardenas et al., 2008]. Additionally, the manufacturing sector faces unique challenges due to its objectives and context. Maintaining safe environments, safeguarding process and machine confidentiality, and ensuring compliance with specifications and requirements are concerns for manufacturers. A framework is needed with which to identify the vulnerabilities faced by a specific manufacturing organization before assessing their cybersecurity risks. This work proposes such a framework, which addresses the cybersecurity vulnerabilities throughout a manufacturing enterprise and its supply chains.

2 Background

To understand the basis for constructing a framework for identifying and addressing cybersecurity risks in manufacturing, it is necessary to first understand the environment in which manufacturing and security reside and define essential terms. In addition, some lexicon for risk must be given, some background on the consequences of cybersecurity failures presented, and the risk management, as it relates to manufacturing and cybersecurity, needs to be introduced.

2.1 Language of Risk

Engineers are accustomed to addressing risk, which Merriam-Webster defines generally as “the possibility that something bad or unpleasant (such as an injury or a loss) will happen” [Merriam-Webster, 2014]. Engineers construct models and experiments to quantify the possibility of something happening and then make adjustments to ensure or increase the likelihood that, with a specified level of confidence, undesirable outcomes will not occur. For example, when selecting bolts that will support a bridge, an engineer may calculate the expected magnitude and direction of forces on the bolt, establish the amount of uncertainty in those measurements, and then select a bolt with appropriate dimensions and material properties to ensure it will not fail within a given factor of safety. Two important elements of risk emerge from this discussion: outcome and probability.

The U.S. Department of Homeland Security (DHS) has developed a set of definitions that are a useful basis for a conversation about risk. The DHS has been charged with ensuring the United States is “safe, secure, and resilient against terrorism and other hazards” [DHS, 2012]. This includes protecting 18 sectors of critical infrastructure and key resources, which includes critical manufacturing resources, as well as safeguarding and securing cyberspace. In the DHS Risk Lexicon, risk is defined as the “potential for an unwanted outcome resulting from an incident, event, or occurrence, as determined by its likelihood and the associated consequences” [DHS, 2010]. The two elements identified previously, “outcome” and “probability,” are included here and referred to as “consequences” and “potential,” respectively. Furthermore, DHS often adopts a representation of risk in which risk is a function of threats, vulnerabilities, and consequences. These terms are defined as follows:

- threat – natural or man-made occurrence, individual, entity, or action that has or indicates the potential to harm life, information, operations, the environment, and/or property
- vulnerability – physical feature or operational attribute that renders an entity, asset, system, network, or geographic area open to exploitation or susceptible to a given hazard
- consequence – effect of an event, incident, or occurrence

In short, a threat is something that has the potential to do harm; a vulnerability enables a threat to do harm; and a consequence is the harm that is done.

For clarity in this discussion, it is useful to differentiate between hazard, threat, and risk. A hazard i) is a source or cause of an issue, ii) can be actual or potential, and iii) is **not** specifically directed by an individual or organization at an entity or geographic area [DHS, 2010]. The third element of hazard, its lack of direction toward a target, differentiates it from threat. Finally, the notion of risk includes consideration of the probability of occurrence. There is a qualitative or quantitative likelihood that a threat or hazard will result in harm or difficulty. The likelihood is not inherent in the threat or hazard, but is an inherent part of risk.

2.2 Cybersecurity

Information security professionals often categorize consequences based on a loss of confidentiality, integrity, or availability (CIA) [SANS, 2013]. Data is confidential when access is limited to authorized individuals or organizations, and its integrity is maintained when it is accurate, authentic, and complete. Availability indicates that data and cyber-physical systems are available to authorized users on demand. Health care systems provide illustrative examples of CIA consequences. Data confidentiality is compromised if unauthorized individuals access a patient’s medical records. If the patient’s medical records were altered, data integrity is degraded. Data availability is affected if the patient’s records were removed from the system entirely; availability of a cyber-physical system is affected if the functionality of medical equipment is altered. This demonstrates that confidentiality, integrity, and availability are interrelated. Furthermore, when developing security measures, reducing potential consequences in one area may impact another. For example, incinerating hard drives containing patient data will ensure confidentiality, but severely diminish availability.

Cyber threats have the potential to affect confidentiality, integrity, and availability in a manufacturing setting as well. Manufacturers possess a substantial amount of information that they want to remain confidential. This is most often the case because the information provides some economic or technological advantage or because release of the information would result in negative competitiveness repercussions. Examples include product designs, process plans, business plans, financial information, and employees' personal information.

A cyber espionage group, which the computer security firm Mandiant labeled APT1 (where "APT" stands for advanced persistent threat) was particularly prolific in terms of its ability to exfiltrate data from companies [Mandiant, 2013]. Based on investigations of cybersecurity breaches for a number of organizations, Mandiant believes it characterized elements of the APT1 attacks and ultimately "uncovered a substantial amount of APT1's attack infrastructure, command and control, and modus operandi." According to Mandiant's report, APT1 compromised 141 companies spanning 20 major industries, including aerospace, energy, construction, manufacturing, engineering services, and metals and mining. APT1 is reported to have stolen hundreds of terabytes of data, including broad categories of IP, such as blueprints, proprietary manufacturing processes, test results, pricing documents, and partnership agreements. Accessing and distributing information of this nature could negatively impact production by the company with legal rights to the information while enabling other organizations to quickly replicate the technology and potentially take over market share.

The importance of data integrity for manufacturing can be seen in relation to part production. Altering product and process specifications could be detrimental to product quality and reliability. Wells et al. presented a case study of a cyber attack on a milling system [Wells et al., 2014]. Groups of 3-4 sophomore-level engineers were asked to design and manufacture a tensile test specimen. The groups generated tool paths using Computer Aided Manufacturing (CAM) software and the file was transferred onto a PC-controlled mill. The PC with the CAM software was infected with a virus that altered the tool path file, ultimately altering the resulting specimen's dimensions. Wells et al. demonstrated that i) part quality could be degraded through such an attack, ii) if quality control systems are not in place, changes in quality may go unnoticed by engineers, and iii) even when a systematic analysis of the process reveals the step in which the file is altered (i.e., incorrect transfer of the file to the PC), a "cyber" event may not be identified as a possible root cause. While the purpose of this study was, in part, to clarify the need to educate future engineers about the cyber threats and provide them with tools to identify cyber attacks on manufacturing processes, the study also solidifies the need for data integrity within manufacturing systems.

Data and cyber-physical system availability is also critical to manufacturing productivity. The information that describes what is going to be manufactured and how it will be manufactured must be available as required throughout the manufacturing process. At an even more fundamental level, the equipment necessary for manufacturing can be rendered unavailable through cyber attacks, as Stuxnet has demonstrated. Stuxnet, a computer worm that has been found in industrial control systems worldwide, is an often-cited example of the extreme end of availability concerns. Stuxnet is believed to have targeted centrifuges used for enriching uranium at a plan in Natanz, Iran, ultimately causing the destruction of about 1,000 of the 9,000 centrifuges on site [Albright et al., 2011]. It eliminated the *availability* of physical systems by reducing the *integrity* of the software supporting the systems. Stuxnet worked by initially infecting machines running Microsoft Windows via USB sticks; Stuxnet would check for Siemens Step 7 industrial control system (ICS) software; if this software was found, Stuxnet would attempt to access the Internet to update to the most recent version of itself [Kushner, 2013]. Then, Stuxnet would seek to compromise the system's programmable logic controllers (PLC), first observing the targeted system, and then taking control of the centrifuges, causing them to intermittently run at high rates and fail. Stuxnet's design and architecture is not domain-specific, and it could be adapted to attack PLCs in industrial plants, power stations, or other facilities with automated machinery.

As Wells et al. pointed out, “manufacturing systems are more than a collection of control systems; they are highly integrated with the product lifecycle” [Wells et al., 2014]. Therefore, manufacturing enterprises must be cognizant of cyber threats from initial design through final inspection to product use and throughout the enterprise from the corporate headquarters to the tool tips to their ICS. An approach is needed to bring all of these issues into existing risk management practices.

2.3 Risk Management

There are three primary components of risk management: identifying, assessing, and addressing risk [NRC, 2005]. Because any organization, system, project, or product faces a multitude of risks and because resources available for risk mitigation are usually limited, effective risk management often includes prioritizing risks. Decision makers must evaluate the potential opportunity or benefit associated with a given action as well as its risks. This is especially challenging when considering cyber risks because the Internet was designed to facilitate interoperability and efficiency rather than security. Within the Internet, there is an inherent trust among components that provide ample opportunity for those with ill intent to do harm. Therefore, decision makers must carefully consider when and how to connect systems and share data. While there are many opportunities to create new efficacies by harnessing the power of big data in manufacturing, there are also substantial risks.

Numerous tools exist to help organizations, and manufacturers specifically, to manage risk [NRC, 2005; DeVor et al., 2007; Yoe, 2011]. These tools include root cause and essential function analysis, probabilistic risk assessment, Pareto diagrams, process control charts, failure modes and effects analysis, event trees, evidence maps, and waterfall diagrams. Tools that have even broader applicability and can also be leveraged to guide risk assessment and prioritization include multivariate statistical analysis, multi-criteria decision analysis, system dynamics, sensitivity analysis, and stochastic simulation. Cyber attacks are yet another risk that must be evaluated and prioritized among all other risks facing an enterprise.

In Executive Order 13636, the President charged The National Institute of Standards and Technology (NIST) with developing a Framework for Improving Critical Infrastructure Cybersecurity [2013]. The Framework was released in February 2014, as Version 1.0. It is expected that the framework will be updated and improved using feedback from industry as they implement it. The framework builds on existing standards, guidelines, and practices and is intended to complement an organization’s risk management processes, and enhance their cybersecurity. The Framework Core, presents cybersecurity activities identified by industry as helpful in achieving specific cybersecurity outcomes. The Framework Core is organized by function category and subcategory, and provides references to standards related to each subcategory. The function categories are: identify, protect, detect, respond, and recover. The need to identify organizational communication and data flows is called out specifically; however, it is mapped to references such as NIST’s Security and Privacy Controls for Federal Information Systems and Organizations [NIST, 2013] and the information transfer policies and procedures included in ISO/IEC 27001:2013 [ISO, 2013]. These references are useful and broadly applicable to critical infrastructure, but do not address the unique cybersecurity challenges facing manufacturing. In fact, the authors are unaware of a previously established, systematic approach to identifying cybersecurity risks in the manufacturing sector.

One of the challenging aspects of risk management is identifying risks [NRC, 2005; Yoe, 2011]. Once risks have been identified, many risk assessment tools are available. It is often difficult to imagine what could go wrong, especially in the case of a willful adversary, as opposed to risks that can be traced back to the physical or natural hazards. One key contribution of this paper is mapping data flows as suggested by the NIST Framework, functional subcategory 3 of Identify-Asset Management. This is necessary to identify vulnerable systems and enable assessment of their overall risk. This information can then be included in standard risk management practices to guide allocation of resources to decrease an organization's overall risk.

3 Manufacturing Data Security Framework

We propose a systematic framework for analyzing the opportunities to extract or insert data in a manufacturing system. From a risk management perspective, these opportunities are vulnerabilities. Once they are identified, the potential cost associated with the consequences of an exploit may be evaluated against the cost of mitigation. In the sections that follow, we define all the terms we will use in the framework, describe the data flows and interactions between the different levels of the manufacturing system, and provide a cybersecurity case study specific to advanced manufacturing.

3.1 Framework Structure

The framework we propose is based on the flow of data between activity levels in a given manufacturing enterprise, as well as across its supply chains. Dornfeld et al. [2009] describe six levels of manufacturing activity in the context of the need for interoperable systems to realize computer integrated manufacturing: enterprise, factory, line, machine tool, component, and sensor data. Because our focus is also on data flows, we will consider these levels, but incorporate component and sensor data into a single level, process detail, which reflects the tool-workpiece interface. The nomenclature for this work will be enterprise, facility, line, machine, and process detail, as depicted in the general framework view in Figure 1. The nodes are the various levels of the manufacturing enterprise. Data flows to and from nodes along the edges, which exist between levels in the manufacturing enterprise and span the supply chain between analogous levels of connected enterprises.

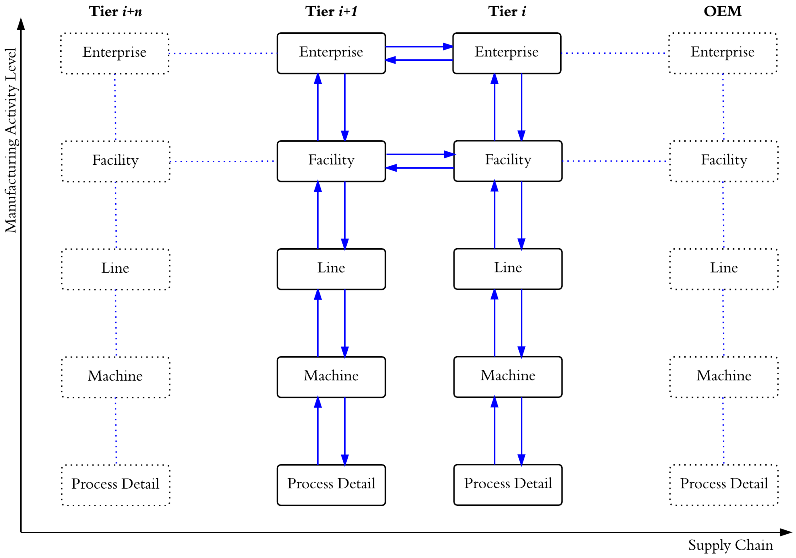

Figure 1: Graphical representation of the activity levels within a manufacturing enterprise and the interactions between tiers of a supply chain (occasional interaction will take place between higher and lower levels of contracting and contracted tiers; however, those arrows have been removed from this figure for simplicity)

There are five categories of data that can be compromised in a manufacturing system:

- high-level digital data,
- low-level digital data,
- financial data,
- physical data, and
- user data.

Financial data refers to all the data that are related to the accounts and finances of the enterprise. Physical data includes actual components and assets in the facility, from which information about the facility and its operations can be inferred. User data refers to user credentials, user information, and information distributed inappropriately (i.e., leaked) by a user in the enterprise. High-level and low-level digital data are distinguished by the level of manufacturing activity in which they flow. High-level digital data flow above the line level, while low-level digital data flow below it.

3.2 Data transfer

We now describe the data that flows in this framework, up and down the enterprise system and throughout the supply chain. The automotive industry will be used as an example to ground this discussion because it i) produces complicated products and components, which are familiar to most manufacturing engineers, and ii) utilizes extensive supply chains and manufacturing processes. We follow the clauses of the ISO/TS 16949:2009 (E) closely while describing the data flows within an automotive manufacturing industry [ISO/TS, 2009]. Appendix A summarizes the clauses in this document and their focus areas. Section 3.2.1 will focus on data that flows along edges in this framework and section 3.2.2 will discuss the data that resides at specific nodes.

3.2.1. Data transfer between enterprise nodes

As shown in Figure 2, at the process detail level in a manufacturing enterprise, low-level digital data like process parameters are dynamically moving and being used while the machine executes a process. For example, process parameters dictate the machine movements, and process feedback from sensors is communicated back up to the machine level. Apart from this, the actual components or materials (i.e., physical data) being processed flow down to the process detail level and back up to the machine level. In the automotive industry, these are referred to as “work in progress” (WIP), “raw materials” (RM) if they are an input to the first operation of a process chain, or “finished goods” (FG) if they are the output of a process chain. All of these pieces of physical data contain clues and insights about the facility and its operations to someone who can intercept or examine them.

At the machine level, WIP and RM are passed on from the line level in a typical manufacturing enterprise. Other examples of physical data are tooling and Quality Control (QC) equipment, which also bring with them insights about the processes in use when they enter into the system. User information like credentials, responsibilities, access levels, and skills are also passed down from the line level. This is an important data type at the machine level, since the machine operator could be a source of valuable information about the machine, its products, its operating schedule, and its setup. Low-level digital data such as machine programs, product and process specifications, and production plans are also transferred from the line level to a machine. The machine gives feedback to the line level in the form of digital QC data and machine-level production information. In the automotive industry, QC feedback in Parts Per Million (PPM) and First-Piece Inspection Reports (FPIR) are sent back to the line level along with production information like target completion ratio and resource consumption. WIP and FG also flow from the machine to the line level.

Flowing from the facility to the line level, high-level digital data manifests itself in the form of design specifications, process chain plans, and production plans. Information about the user is also passed from the facility to the line. This is a key data type, especially because of the type of information they have access to. High-level QC data, like overall PPM levels, process capability (e.g., C_p, C_{pk}), and process/product design feedback flow up from the line level to the facility level. Critical information like inventory data, productivity and production data, and process chain plans flow along the edge from the line to the facility level.

At the highest level of the enterprise, financial information like costs, targets, goals, and objectives flow from the enterprise to the facility. This could be in the form of a balance sheet in an automotive enterprise, for example. These goals and targets are based on the aggregated information on actual costs, inventory consumption, design feedback, productivity, and quality, which flow to the enterprise from the facility, as well as information provided through user analysis at this level. The highest-level form of design specifications originate and flow down from the enterprise node.

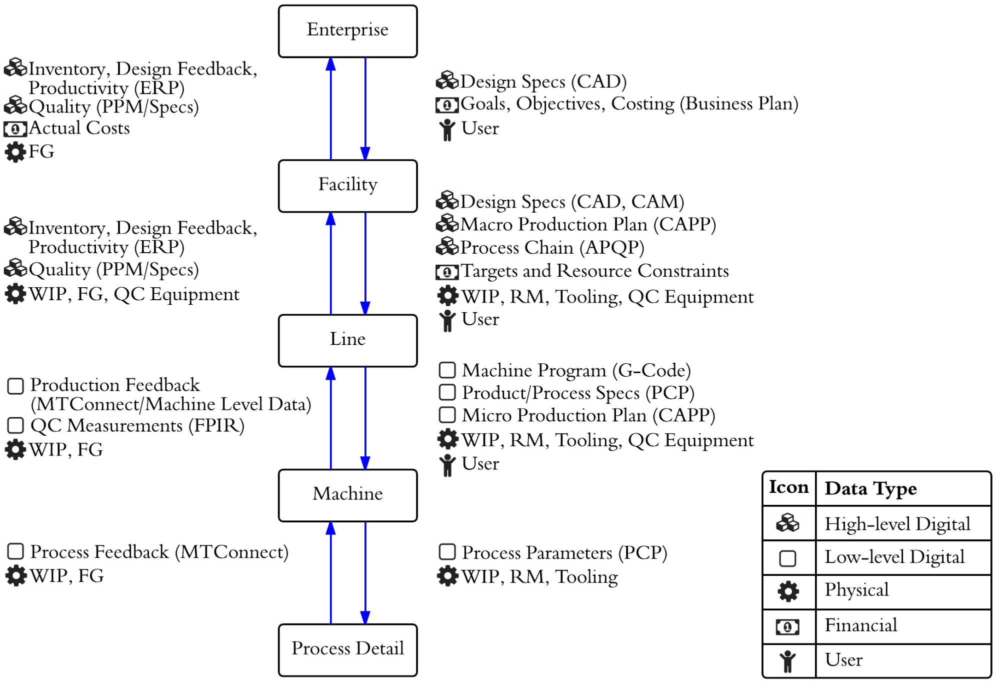
Enterprise
Inventory, Design Feedback, Productivity (ERP)
Quality (PPM/Specs)
Actual Costs
FG
Design Specs (CAD)
Goals, Objectives, Costing (Business Plan)
User
Facility
Inventory, Design Feedback, Productivity (ERP)
Quality (PPM/Specs)
WIP, FG, QC Equipment
Design Specs (CAD, CAM)
Macro Production Plan (CAPP)
Process Chain (APQP)
Targets and Resource Constraints
WIP, RM, Tooling, QC Equipment
User
Line
Production Feedback (MTConnect/Machine Level Data)
QC Measurements (FPIR)
WIP, FG
Machine Program (G-Code)
Product/Process Specs (PCP)
Micro Production Plan (CAPP)
WIP, RM, Tooling, QC Equipment
User
Machine
Process Feedback (MTConnect)
WIP, FG
Process Parameters (PCP)
WIP, RM, Tooling
Process Detail
Icon
Data Type
High-level Digital
Low-level Digital
Physical
Financial
User

Figure 2: Schematic showing an isolated manufacturing enterprise and the data types that flow between its manufacturing activity levels

3.2.2. Data at the nodes

In the previous section, we discussed data flow between enterprise nodes (i.e., the various activity levels in a manufacturing system). This section explores all the data that resides at the nodes. Figure 3 shows the types of data that can be found at the different levels in a manufacturing enterprise.

Most of the data that flows throughout the enterprise system originates and resides at one of the nodes in the enterprise chain. However, apart from the data that flows, there are additional data that are static at these nodes. At the enterprise level, in-depth details of the cost analysis, cost composition for quotations, profitability analyses, business plan calculations and account details, payments, customer and supplier details, and tax information are static and could be critical to the survival of the enterprise. The enterprise's Intellectual Property (IP) and the Unique Selling Point (USP), which provides market advantage, also reside at this top node.

We assume that any servers utilized by the facility reside at the facility level. The server, apart from all the digital information described in section 3.2.1, is likely to maintain other critical information like closed-circuit TV (CCTV) footage, product identification strategies, documentation and standards, correspondence, and important quality and production data. "The cloud" may also store this information; Wang et al. [2014] discuss cloud security and manufacturing. The Enterprise Resource Planning (ERP) system of a typical manufacturing industry also resides at the facility level, which stores in-depth details of all transactions and production data over a large period of time.

Material handling equipment, machines in the process chain, and the interactions within process chains are the sources of information that can be directly compromised at the line level. Machine set-up and operation sequences are low-level digital data that can be derived from the machine at the process detail levels, apart from information about the process parameters, tooling, machines used, and numerical control (NC) codes, which have been discussed in Section 3.2.1 as flowing along edges.

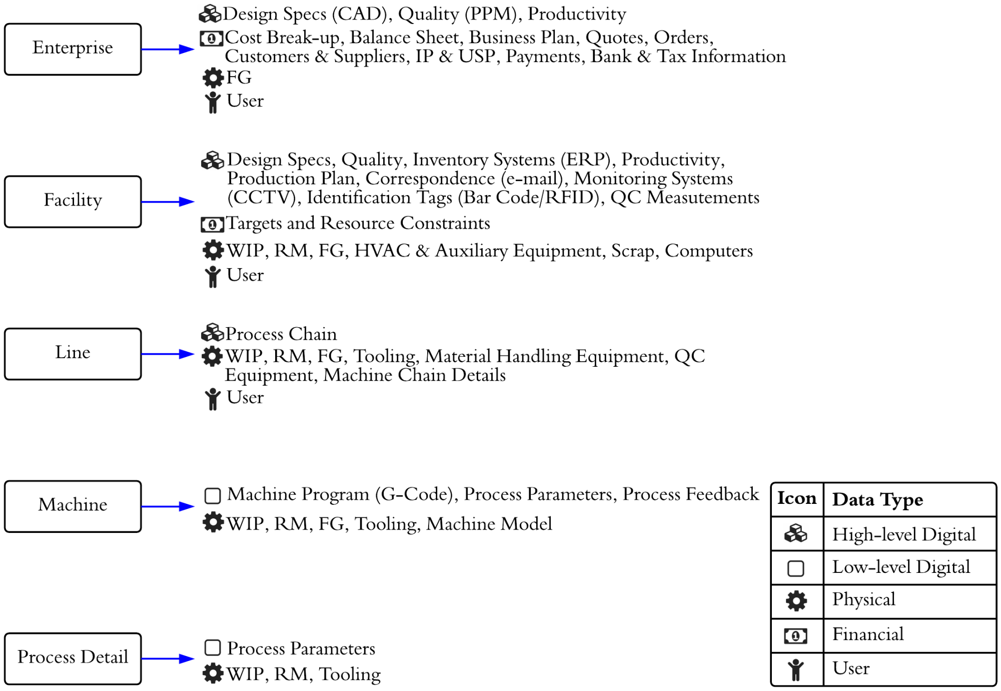
Enterprise
Design Specs (CAD), Quality (PPM), Productivity
Cost Break-up, Balance Sheet, Business Plan, Quotes, Orders, Customers & Suppliers, IP & USP, Payments, Bank & Tax Information
FG
User
Facility
Design Specs, Quality, Inventory Systems (ERP), Productivity, Production Plan, Correspondence (e-mail), Monitoring Systems (CCTV), Identification Tags (Bar Code/RFID), QC Measutements
Targets and Resource Constraints
WIP, RM, FG, HVAC & Auxiliary Equipment, Scrap, Computers
User
Line
Process Chain
WIP, RM, FG, Tooling, Material Handling Equipment, QC Equipment, Machine Chain Details
User
Machine
□ Machine Program (G-Code), Process Parameters, Process Feedback
WIP, RM, FG, Tooling, Machine Model
Process Detail
□ Process Parameters
WIP, RM, Tooling
Icon
Data Type
High-level Digital
□ Low-level Digital
Physical
Financial
User

Figure 3: Types of data that can be found at different activity levels within a manufacturing enterprise

3.2.3. Data transfer across the supply chain

We use the concept of tiers while describing the supply chain of a manufacturing enterprise. The data flows between two adjacent tiers in a manufacturing supply chain are shown in Figure 4.

For the purposes of this framework, we assume that data only flows between the analogous enterprise levels and facility levels. High-level digital data and financial data flow between the enterprise levels, whereas physical data and high-level digital data flow between the facility levels. Supplier audits, a mandatory clause within the ISO/TS standard used by the automotive industry, are audits conducted by an enterprise on its suppliers. These audits extend down to the process detail level of the suppliers. However, we consider this a data flow between enterprises since it is facilitated by the enterprises directly.

Financial data, such as quotes, schedules, and payments, and high-level digital data, such as process details and design specifications, flow directly between the enterprises. RM, WIP, FG, tooling, and QC equipment are the physical data that flow between facilities. This framework incorporates laboratories and calibration agencies as well as lower-tier suppliers. Quality information related to the physical products, like Pre-Dispatch Inspection Reports (PDIR), also flows with these products between facilities. Users are critical data types involved during these inter-enterprise and inter-facility data flows because the users handle all the data transfers. Enterprises also exchange critical design specifications, drawings, and process details, especially during the Part Production Approval Process (PPAP) stage in the automotive industry.

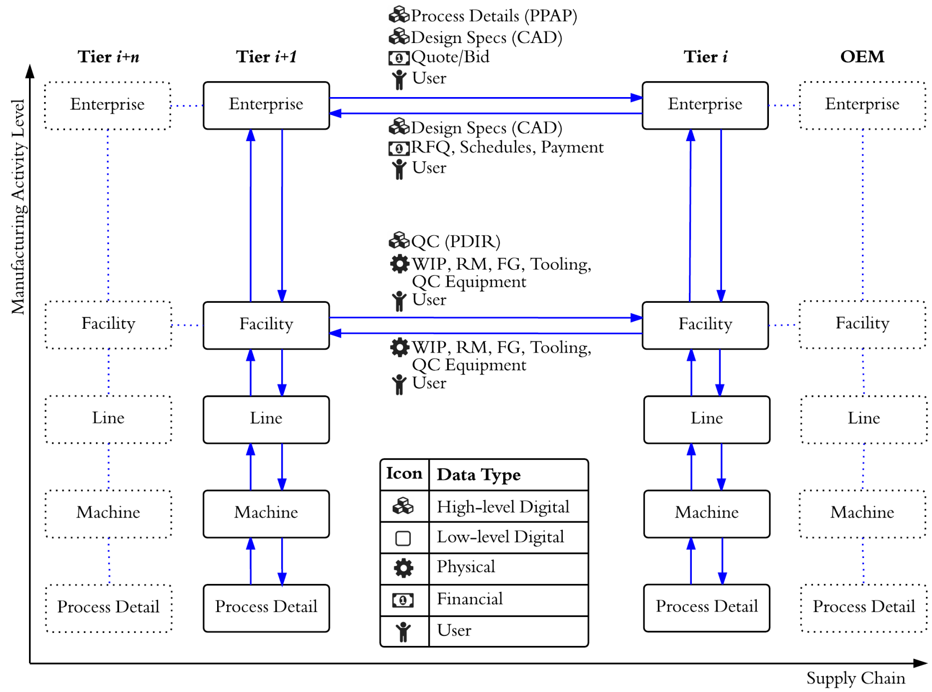

Figure 4: Data flow between two adjacent tiers of a manufacturing supply chain

3.3 Machine Level Example: MTConnect in a Component Manufacturing Enterprise

Organizations are growing increasingly aware of the potential consequences associated with breaches of their IT infrastructure. The number of vulnerabilities is large and increasing in number, the attack surface is changing as technology evolves, and the problem combines technical and non-technical factors. However, some risks are shared among different industrial sectors. Many organizations support development of mitigations for these shared risks, such as securing industrial control systems. The framework presented in this paper enables identification of general cybersecurity risks as well as manufacturing-specific risks. In the following example, we will consider manufacturing process feedback flowing on the edge from the process detail level to machine level.

Manufacturing firms' interest in big data analytics and machine monitoring is increasing as they become more energy and resource conscious. Enterprises of all sizes are striving to minimize energy consumption in their plants and improve their overall equipment efficiencies. There are several data sources in machine tools (i.e., data flows on the edge between process detail and machine levels) that allow firms to obtain large amounts of information from their machines for monitoring and data analysis.

MTConnect, an interoperability standard, aggregates data from multiple sources for systematic data analysis and knowledge extraction [Vijayaraghavan et al., 2008]. Operators and process engineers can obtain large amounts of process data from MTConnect, which is useful for data-driven process refinement [Helu et al., 2014; Bhinge et al., 2014]. As shown in Figure 5, process information such as machining process parameters and strategies can be extracted through the MTConnect data stream, along the edge between the process detail and machine levels, as shown in the framework in Fig 2.

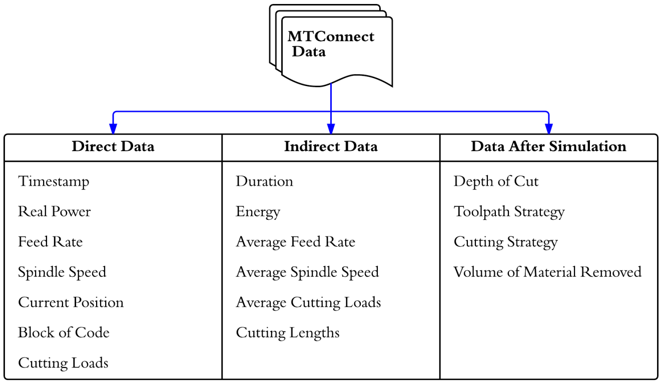

Direct Data	Indirect Data	Data After Simulation
Timestamp	Duration	Depth of Cut
Real Power	Energy	Toolpath Strategy
Feed Rate	Average Feed Rate	Cutting Strategy
Spindle Speed	Average Spindle Speed	Volume of Material Removed
Current Position	Average Cutting Loads	
Block of Code	Cutting Lengths	
Cutting Loads		

Figure 5: Data that can be collected from MTConnect [Bhinge et al., 2014]

Though this data is extremely useful for process improvement, its flow represents a vulnerability in the system; this sensitive information could be accessed by unauthorized individuals. As shown by Bhinge et al. [2014], the entire cutting process of a machine can be simulated and reproduced with just this stream of information. In metal cutting, for example, process parameters such as feed rates, spindle speeds, depths of cut, entire NC codes, and tool details can be extracted. This information is particularly valuable to component and part manufacturers because their process parameters are often what provide their competitive advantage, especially for the machining or processing of specialty materials and alloys.

As MTConnect is currently configured in many applications, the only piece of information required to gain access to a machine tool at a remote site and abstract data from it is the IP address of the MTConnect agent. Although this configuration is relatively vulnerable, the system can be secured by IP masking, authorization, or data encryption. Since the process data is available through the Internet, the threat of compromised information is truly global in reach. The implementation of such a data acquisition scheme is enabling for manufacturing capabilities, but introduces vulnerabilities that must be mitigated.

This scenario is becoming increasingly common in manufacturing firms, and a cybersecurity framework is needed to provide such data-driven manufacturing systems with a structured approach to identifying the vulnerabilities in their systems and enable them to flourish [NIST, 2014; NIDA, 2014]. We have proposed a framework in this paper that provides such an approach. The example at the machine level presented here is limited to low-level digital data on an edge between two nodes in a manufacturing enterprise. As the framework suggests, there are likely many other nodes and edges that should be examined within the enterprise and across the supply chain.

4 Discussion

In light of the framework described in Section 3, we will discuss the cybersecurity risks to manufacturing more holistically. We will provide further insight into the vulnerabilities that the framework helps to identify; discuss the two other elements of risk: threats and consequences, from a manufacturing perspective; and explore the interrelationships among risk elements to create risk. Finally, we outline risk mitigation strategies.

4.1 Threats

The above framework describes all the typical data flows and information exchange within a manufacturing enterprise and between its supply chains. This provides an understanding of the types of data that can be compromised and points to areas in the system that are vulnerable to cyber threats. In order to evaluate the risks associated with vulnerabilities in a system, a manufacturer must first consider the threats to that system. As discussed previously, a threat is anything, human or otherwise, that has the potential to do harm. Cyber threats to manufacturers span issues from a hot day and an HVAC (heating, ventilation, and air conditioning) system that could be modified to cause temperatures to rise beyond acceptable levels in a server room, to a bad actor seeking to make a political statement by causing substantial damage to a manufacturer's capital assets.

When an organization seeks to identify its risks, it must consider the types of threats that could exploit different vulnerabilities and what the consequences would be to the organization. Cyber threats can be characterized by their capabilities and resources, which may be linked to the types of consequences these threats can achieve [DoD, 2013]. For example, virtually all organizations have competitors that would be interested in their IP, but the data that contains the most valuable IP may exist in different edges and nodes, depending on the organization and product of interest. A brand manufacturer may value its unique product design and marketing strategy far more than the sequence of its assembly processes. However, for a commodity component manufacturer, it may be its sequence of manufacturing processes that provides its competitive advantage; therefore, data that provides insights into this sequence may be of most interest to a competitor. The commodity component manufacturer in this example must then consider the vulnerabilities in the system that could provide access to this data and the type of threat (e.g., highly capable) required to exploit those vulnerabilities. Once these assessments are complete, the organization can assess the relative likelihood of such an event (i.e., likelihood a threat exists with interest in the data and the capability to exploit the vulnerability) and what mitigation strategy it should take.

Manufacturers should consider a range of potential threats when conducting their cyber risk analyses. Low-level threats, those individuals and groups with limited capability and resources, may be problematic if a manufacturer's systems are extremely vulnerable, regardless of the consequences that could be achieved. These threats may simply be interested in testing their skills and causing nuisance problems. High-level threats, those with considerable capabilities and resources, may be more interested in consequences that provide them with good return on investment (e.g., valuable IP) or are highly visible and cause public fear. After considering a range of threats, the manufacturing enterprise can determine which are most likely to contribute to a risk that it should mitigate.

4.2 Vulnerabilities

We distinguish vulnerabilities in the context of manufacturing cyber security as intrusive and interceptive. Intrusion refers to physical or remote (cyber) entrance into a system. This could include hacking into a server and compromising information or physical intrusion into an enterprise to access information regarding the manufacturing strengths of that enterprise. Users are key data types in the case of intrusion. Users can either leak information willingly, as in the case of privileged user abuse, or may give out information to an intruder unwillingly (and perhaps unknowingly). Another form of intrusion is hacking into a central server and leaking information or altering information to cause a temporary shutdown. In intrusion, the data existing at the nodes are the most vulnerable. This relates back to the data at the nodes described in Section 3.2.2 and shown in Figure 3.

Interception, as the term suggests, refers to data being compromised while flowing from one node to another along an edge. Interception may occur as information and physical data move through the supply chain. Data and physical products in transport from one organization to another are extremely vulnerable to these forms of threats. Tapping phone lines, leaking information during data transfer and interception of payments and shipments are all forms of interception. The key aspects to address are the requirement of data identification, tracking, and encryption during any of the data exchanges discussed in Sections 3.2.1 and 3.2.3, and shown in Figures 2 & 4.

4.3 Consequences

The consequences of cyber exploits in manufacturing vary widely between the different levels of a manufacturing enterprise and across the different levels of a supply chain. Typically, Original Equipment Manufacturers (OEMs) are the most concerned about protecting their information and designs, and the strict confidentiality requirements have to be complied with by the lower tiers of the supply chain. The consequences of compromising data at the higher levels of an OEM enterprise could be disastrous, since they hold the IP or USP of the product, which the entire supply chain depends upon. Data at the lower levels of a manufacturing enterprise are more specific, of higher frequency, and provide little abstraction. However, replication of lower level processes at an external site could be a consequence of compromising lower level data.

The IP or USP of a manufacturing enterprise could reside within any level of activity in the manufacturing enterprise. The marketing strategy and business model, for example, could be the core strength of an enterprise, and would likely reside at the enterprise level. A radical design or a patented invention are examples of IP that reside at the facility level. Special processes and process chain novelty such as hard turning and unconventional machining process chains are examples of IP that reside at the line level. At the machine level and process detail level, the material properties, process parameters, tooling choice, or the surface quality maintained could provide a competitive advantage and be considered IP. In short, IP could reside at any level and the consequences of compromising this information could be catastrophic ranging from loss of critical proprietary data or knowhow to serious disruptions to production or delivery schedules.

A serious consequence of a cyber exploit could be the loss of integrity. Quality systems in manufacturing are becoming increasingly digital and connected, and at the same time vulnerable to cyber exploits. As Wells et al. have described, hacking into and altering the measurement systems for product quality can result in non-conforming products being passed up the supply chain (or conforming items being scrapped), despite the outcome of the product's evaluation by the compromised quality system's tests [Wells et al., 2014]. This can result in lower product quality, expensive recalls, decreased brand value, and product failures in the field.

Loss of availability due to a cyber exploit directly impacts quality and/or productivity. Such a consequence could be in the form of power disruption, a compromised machine tool, or a malfunctioning temperature controller in a calibration facility, for example. This is a direct result of a bad actor not just intruding into the system, but also modifying it, as in the previous Stuxnet worm example. The probabilities of such consequences must be determined by manufacturing enterprises and factored into their contingency plans.

4.4 Risk Mitigation

Ultimately, the information a manufacturer gathers about its cyber risk (i.e., threats, vulnerabilities, and consequences) needs to be included in its existing risk management processes and used to guide decisions about allocation of resources to mitigate its risks. The strategies an organization chooses to employ or actions it chooses to take may include development of a cybersecurity policy, establishing access controls (or enforcing existing access controls), utilizing multi-factor authentication, patching and updating regularly, judicious use of firewalls, and encrypting data. As manufacturers implement risk mitigation strategies, it may become necessary to identify and characterize the relationships between information technology and operation technology architecture. This has been challenging in other sectors and presents an interesting opportunity for interdisciplinary research focused on manufacturing. A detailed discussion of mitigation strategies is beyond the scope of this paper; however, there are numerous resources that could be leveraged to identify useful mitigation strategies [CCS, 2014; ISO, 2013; ISA, 2009]. Many of these referenced approaches are well aligned with existing quality management practices. An example is the Failure Modes and Effects Analysis (FMEA) conducted in many industries with regard to risks in product quality, which involves assignment of a Risk Priority Number (RPN) to each risk based on its threat, vulnerability, and consequence. The risks are then ordered by their RPNs and appropriate mitigation strategies are implemented for the risks with high RPNs [Dyadem Press, 2003].

5 Conclusion

Increasing connectivity and the use of digital computation provides potential for dramatic improvement in manufacturing productivity, quality, and cost. However, there are also risks associated with increasing connectivity, generating larger amounts of data, and the potential access competitors or adversaries may be able to gain. Enterprises have experienced cyber attacks that steal confidential and/or proprietary data, alter information to have an unexpected or unwanted effect, and destroy capital assets. Manufacturers need tools to consider these risks in their existing risk management processes.

In this paper, a framework was established that considers the data flows within a manufacturing enterprise and throughout the supply chain. The framework provides several mechanisms for identifying generic and manufacturing-specific vulnerabilities in these data flows and is illustrated with details pertinent to an automotive manufacturer. In addition to providing manufacturers with insight into their potential data risks, this framework addresses an outcome identified by the NIST Cybersecurity Framework.

However, much work remains to fully integrate cybersecurity concerns into existing risk management practices. Vulnerabilities in the system must be linked to potential consequences. In turn, the costs associated with consequences must be better defined so that they may be prioritized. Additionally, mitigation strategies must be identified that effectively address cybersecurity risks across all levels of manufacturing enterprises and supply chains. Finally, enterprises face growing pressure for transparency throughout their supply chains to support sustainability analyses of their products. There is a desire for data to understand and address issues such as life-cycle resource use and inappropriate working conditions. This will contribute to increased data flow, making the need to address these vulnerabilities even more urgent.

Acknowledgements

This work was supported in part by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-AC04-94AL85000. This work was also supported by the Laboratory for Manufacturing and Sustainability (LMAS) sponsors as part of the Sustainable Manufacturing Partnership (SMP) and Machine Tool Technologies Research Foundation (MTTRF) For more information, please visit http://smp.berkeley.edu and http://lma.berkeley.edu.

References

Albright, D., P. Brannan, and C. Walrond, 2011, *Stuxnet Malware and Natanz: Update of ISIS December 22, 2010 Report*, Institute for Science and International Security Report. Accessed from http://isis-online.org on 10/27/14

Bhinge, R., Park, J., Biswas, N., Helu, M., and Dornfeld, D., Law, K., and Rachuri, S., 2014, *An Intelligent Machine Monitoring System Using Gaussian Process Regression for Energy Prediction*. IEEE International Conference on Big Data (IEEE BigData 2014), Washington, DC

Cárdenas, A.A., S. Amin, S. Sastry, 2008, *Research Challenges for the Security of Control Systems*, HOTSEC'08 Proceedings of the 3rd conference on Hot topics in security, Article No. 6

Cárdenas, A.A., S. Amin, B. Sinopoli, A. Giani, A. Perrig, and S. Sastry, 2009, Challenges for securing cyber physical systems. In: Workshop on Future Directions in Cyber-Physical Systems Security. DHS; 2009.

Clark, D., T. Berson, and H.S. Lin, Editors, 2014, At the Nexus of Cybersecurity and Public Policy: Some Basic Concepts and Issues, National Academies Press, Washington DC.

CCS (Council on CyberSecurity), 2014, Critical Security Controls for Effective Cyber Defense. Accessed from http://www.counciloncybersecurity.org/ on 11/04/14.

Dyadem Press, 2003, Guidelines for Failure Mode and Effects Analysis (FMEA), for Automotive, Aerospace, and General Manufacturing Industries. CRC Press, Boca Raton, Fl.

Department of Homeland Security, 2010, *DHS Risk Lexicon: 2010 Edition*.

Department of Homeland Security, 2012, *Our Mission*. Accessed from http://www.dhs.gov/our-mission on 10/22/14

DeVor, R.E., T. Change, and J.W. Sutherland, 2007, Statistical Quality Design and Control. Prentice Hall.

DSB (Defense Science Board), 2013, *Resilient Military Systems & Cyber Threat.* Accessed from http://www.acq.osd.mil/dsb/reports/ResilientMilitarySystems.CyberThreat.pdf on 11/4/14.

Dornfeld, D., P. Wright, A. Vijayaraghavan, and M. Helu, 2009, *Enabling Manufacturing Research through Interoperability*, Transactions of NAMRI/SME, 37, 443-450.

Federal Bureau of Investigation, 2014, *Intellectual Property Theft*. Accessed from www.fbi.gov on 10/30/14.

Forbes, 2014, *Target Data Breach Spilled Info on as Many as 70 Million Customers*. Accessed from forbes.com on 10/14/14

Fortune, 2014, *The Digital Factory*. 130(10), 92-110

Helu, M., Robinson, S., Bhinge, R., Bänziger, T., and Dornfeld, D., 2014, *Development of a Machine Tool Platform to Support Data Mining and Statistical Modeling of Machining Processes*. Proc MTTRF 2014 Annual Meeting, San Francisco, CA, 2014.

ISA (International Society of Automation), 2009, *ANSI/ISA–62443-2-1: Security for Industrial Automation and Control Systems: Establishing an Industrial Automation and Control Systems Security Program*.

ISO (International Organization for Standardization), 2009, *Technical Specification ISO/TS 16949, Quality Management Systems; Automotive Suppliers*.

ISO (International Organization for Standardization), 2013, *Standard: Information technology — Security techniques — Information security management systems — Requirements*.

Kushner, D., 2013, The Real Story of Stuxnet, IEEE Spectrum. Accessed from http://spectrum.ieee.org on 10/27/14

Lindsay, J.R., 2013, *Stuxnet and the Limits of Cyber Warfare*, Security Studies, 22(3), 365-404

Mandiant, 2013, *APT1: Exposing One of China's Cyber Espionage Units*. Accessed from http://intelreport.mandiant.com on 10/26/14.

Merriam-Webster Dictionary, 2014, *Definition of Risk*. Accessed from http://www.merriam-webster.com/dictionary/risk on 10/20/14

New York Times, 2014, *Sony Films Are Pirated, and Hackers Leak Studio Salaries*. In press December 2, 2014, accessed from www.nytimes.com on 1/26/15.

New York Times, 2015, *More Sanctions on North Korea After Sony Case*. In press January 2, 2015, accessed from www.nytimes.com on 1/26/15.

National Research Council, 2005, The Owner's Role in Project Risk Management. The National Academies Press, Washington D.C.

NDIA (National Defense Industrial Association), 2014, *Cybersecurity for Advanced Manufacturing*. Accessed from www.ndia.org on 09/22/14.

NIST, 2013, *Security and Privacy Controls for Federal Information Systems and Organizations*. Accessed from http://dx.doi.org/10.6028/NIST.SP.800-53r4 on 1/26/15

NIST, 2014, *Framework for Improving Critical Infrastructure Cybersecurity*. Accessed from http://www.nist.gov/cyberframework/ on 2/28/14

Office of the White House, May 2009, *Cyberspace Policy Review: Assuring a Trusted and Resilient Information and Communications Infrastructure*. Accessed from http://www.whitehouse.gov/assets/documents/Cyberspace_Policy_Review_final.pdf on 07/30/14

The President, 2013, *Executive Order 13636: Improving Critical Infrastructure Cybersecurity*. Federal Register, Vol. 78, No. 33.

SANS, 2013, *Security 401: SANS Security Essentials Bootcampstyle*. Course material.

Symantec, 2014, *Internet Security Threat Report: 2014*. Accessed from www.symantec.com on 10/26/14.

PCAST, 2011, *Report to the President on Ensuring American Leadership in Advanced Manufacturing*. Accessed from http://www.whitehouse.gov on 11/04/14.

Web Wall Street Journal, 2014a, *Home Depot's 56 Million Card Breach Bigger Than Target's*. Accessed from online.wsj.com on 10/15/14.

Wall Street Journal, 2014b, *J.P. Morgan Says About 76 Million Households Affected by Cyber Breach*. Accessed from online.wsj.com on 11/04/14.

Verizon, 2014, *2013 Data Breach Investigations Report*. Accessed from https://ics-cert.us-cert.gov on 10/26/14.

Vijayaraghavan, A., Sobel, W., Fox, A., Dornfeld, D., and Paul, W. (2008). "Improving machine tool interoperability using standard interface protocols: MTConnect," Proceeding of 2008 International Symposium on Flexible Automation, 2008, Atlanta, USA.

Wang, Y., S. Ma, and L. Ren, 2014, *A security framework for cloud manufacturing*, Proceedings of the ASME 2014 International Manufacturing Science and Engineering Conference, June 9-13, Detroit, Michigan, USA.

Wells, L.J., J.A. Camelio, C.B. Williams, and J. White, 2014, *Cyber-physical security challenges in manufacturing systems*, Manufacturing Letters, 2, 74-77.

Yoe, C., 2011, Principles of Risk Analysis: Decision Making Under Uncertainty. CRC Press, Boca Raton, Fl.

Appendix A

Areas of focus in an automotive manufacturing industry described as examples and their corresponding clauses in the ISO/TS 16949:2009 (E) Standard [ISO/TS, 2009]

No.	Description	ISO/TS 16949:2009 (E) Clause
1.	Supplier Audit	7.4.3.2
2.	Contingency Plan	6.3.2
3.	Process Control Plan	7.5.1.1
4.	Part Production Approval Process	7.3.6.3
5.	Maintenance	7.5.1.4
6.	Measurement Systems Analysis	7.6.1
7.	Calibration	7.6.2
8.	Internal Audit	8.2.2
9.	Corrective Action	8.5.2
10.	Quality Control	8.2.4
11.	Advanced Product Quality Planning	7.1
12.	Design Failure Modes and Effects Analysis	7.3.3.1