

NIST

**National Institute of
Standards and Technology**

U.S. Department of Commerce

U.S.
Resilience
Project

BEST PRACTICES IN CYBER SUPPLY CHAIN RISK MANAGEMENT

Intel Corporation

Managing Risk End-to-End in Intel's Supply Chain

INTERVIEWS

David A. Brown

Senior Principal Engineer, Data Center Group

The Next New Things in Cyber Security Supply Chain Risk Management

- **Providing near real-time transparency** into a part's provenance reduces the risk of counterfeiting. This includes tools and methods to audit provenance claims of a part at any location in the supply chain, just before a part is installed into a platform, and in-situ.
- **Improving continuity of supply** reduces the risk of counterfeiting. Because obsolete parts are at much higher counterfeiting risk, long-term authorized sources for critical parts are an essential part of cyber supply chain risk management.
- **Temporarily increasing controls immediately after natural disasters also reduces counterfeiting risks.** Part shortages that occur when supply is interrupted by disasters such as earthquakes and floods attract counterfeiting attacks. Additional checks and controls are implemented until authorized supply sources are restored.

Company Overview

Intel Corporation (commonly referred to as Intel) is an American multinational technology company headquartered in Santa Clara, California. Intel is one of the world's largest and highest valued semiconductor chip makers, based on revenue. It is the inventor of the x86 series of microprocessors, the processors found in most personal computers. Intel has a large and diverse product portfolio, including motherboard chipsets, server boards, chassis and systems, network interface controllers and integrated circuits, flash memory, graphics chips, embedded processors and other devices related to communications and computing.

Supply Chain Organization and Management

Intel is one of many Original Component Manufacturers (OCMs) that operate in a worldwide market that offers both opportunity and risk. All OCMs face these in-common market risks, and each OCM implements risk management according to the OCM's business needs and policies. Likewise, OCM supply chains face similar in-common risks, and Intel's supply chain faces the same risks and threats as does every other OCM. The risks and threats discussed in this paper are common to all OCMs, so the risk management techniques offered here can be shared.

Supply chain risk management at Intel spans many business units and functional groups, but is primarily coordinated through a centralized Technology and Manufacturing Group (TMG). This corporate-wide organization has oversight over all wafer fabrication factories, all assembly and test plants that convert the wafer into finished integrated circuits, all the warehousing and shipping of finished goods, and commodity management of all incoming materials used by these operations.

Intel supply chain risk management encompasses both the inbound and outbound supply chains. The four distinct operations include:

- **Inbound materials:** Sourcing physical ingredients used to make electronic parts.
- **Function development:** Designing the electronic part functions and logical processes.
- **Enterprise & manufacturing processes:** Actually making the parts.
- **Outbound finished goods and spares:** Getting the parts to end users and supporting them.



Each operational area faces distinctly different technical challenges and potential risks, therefore each operational area implements different risk mitigation controls. Representative examples of potential issues and risks by operational area are:

Inbound Materials

- Non-conforming parts and materials
- Conflict minerals

Function Development

- Faulty, inadequate, or misused design tools
- Architectural/design vulnerability
- Unintended consequences of intentional design changes
- Compromised or stolen secrets

Enterprise & Manufacturing Processes

- Network/system vulnerability
- Unauthorized facility access
- Business continuity
- Infrastructure availability
- Unauthorized changes to machine settings
- Improper configuration of factory options
- Incomplete testing

Outbound Finished Goods and Spares

- Freight theft
- Tampering
- False description
- Product substitution
- Counterfeiting

Inbound Materials

The practices that Intel follows to manage incoming materials are well-established and time-tested. These practices include developing long-term relationships with top-tier vendors who have proven track records for consistently delivering high quality ingredients. Intel periodically performs joint audits with these vendors to identify potential quality and functional issues, and if found, jointly evaluates corrective actions to address the audit findings.

Developing leading-edge products frequently requires incorporation of new technologies and materials. This, in turn, requires continuous efforts to on-board new vendors to source these new technologies. Intel has a supplier selection process that considers many factors such as quality, availability, and security to develop mitigation plans that compensate for the absence of a long-term relationship and proven track record.

Intel also leads development of industry-wide supply chain initiatives, such as [conflict-free minerals](#). Risks and concerns here include health and safety of workers and communities, living wages and other labor rights, displacement and resettlement, environmental impacts, and other social and sustainability issues. In this example, mitigating supply chain risk required many years of consistent leadership with never-ending focus on goals established by a broad international community.

Function Development

Intel has created a set of policies, procedures, tools, indicators, and consulting practices called the Security Development Lifecycle (SDL). The SDL provides an evaluation framework designed to help the company determine whether the product meets technical specifications, delivers to security objectives, supports the protection of privacy and personal information, and does not contain malicious software or hardware when shipped.

Intel is an active member in security and privacy industry consortia and monitors and benchmarks its own SDL, and adherence to its SDL, against industry peers, global standards, and specific regulations. Some of the international standards that provide guidance to Intel's SDL include ISO/IEC 27001, 27002, 27034-1 and 27036-3.

Enterprise and Manufacturing Processes

For more than 40 years, Intel has been designing and delivering high-quality reliable products across multiple manufacturing and assembly locations. Intel follows a “Copy Exact!” methodology to match manufacturing environments with the development environment at all levels for physical inputs and statistically matched outputs. Statistical monitoring and cross-site audits can match or detect deviations throughout the production process.

Intel uses the ISO 9001:2008 International Standard as the baseline for its quality system. A third party registrar maintains information relative to mature wafer fabrication sites, assembly and test sites, and logistics centers. The ISO 9001:2008 International Standard describes a system of standards and procedures that help ensure product uniformity across manufacturing sites. Certificates are available for viewing at the [Intel® ISO Registrations](#).

Wafer manufacturing processes used at Intel mitigate many risks of counterfeit ingredient infiltration. During wafer processing, it is also very difficult to modify the IC design to insert malware. These risk mitigation benefits are created by two deeply rooted processes involved with wafer manufacturing: mask sets and common mode yield analysis.

Mask Set: Once the function development is completed, the resulting IC design is converted to a wafer mask set. This conversion process is very compute-intensive and utilizes proprietary algorithms and other trade secrets that are explicitly tuned for Intel wafer processing.

Once made, a mask sets is extremely difficult to edit. While it is possible to use an ion mill to change a mask, the details of that change must be calculated by the proprietary mask conversion process. Access to mask sets is very carefully controlled. Mask sets are stored in ultra-clean containers within an inventory management robot. People are not allowed near the masks because they would carry in dust or other particles that could contaminate them. The inventory management systems log all mask movements. The bottom line is, by protecting the mask sets, Intel reduces the risk of malicious function [malware] being introduced during wafer manufacturing.

Common Mode Analysis: The yield of every wafer is measured through testing. In order to understand what causes changes in yield, Intel keeps track of the batch identity for all the ingredients used to fabricate each wafer. Computerized records are maintained that detail which machine processed each individual wafer, including which batch of chemical or other material was used. Changes in

yield are correlated to changes in material, machine settings, or environment. Intel investigates any change in yield. Counterfeit materials that cause a change in yield are investigated and correlated to the batch of counterfeit material.

Outbound Finished Goods and Spares

Since the mid-2000s, concern over the risk that counterfeit electronic parts might cause failure of a business critical application has grown. For many, the [October 1, 2008 BusinessWeek magazine cover story “Dangerous Fakes”](#) made the counterfeit threat a priority. In 2011, U.S. Congress mandated in the National Defense Authorization Act for fiscal year 2012 that the Department of Defense must mitigate the threat of counterfeit electronic parts. This mandate was signed into law on December 31, 2011. One of the most important goals of supply chain risk management is to mitigate the threat of “fakes” infiltrating and contaminating business critical systems.

The two most broadly used database services for reporting instances of counterfeit electronic parts [GIDEP and ERAI] contain only a small percentage of what are believed to be the total actual incidents. According to David Brown:

“Adversaries and counterfeiters do not publicly report the scale of their businesses, so we don’t know how many counterfeits are in the open market. Lacking statistically significant data, we are forced to rely on a few public, anecdotal stories and examples where counterfeiting clearly creates a threat to public safety and significant potential for harm. These examples include counterfeit parts in [helicopter engine and hydraulic controls](#), high speed train brake controls, hostile radar tracking systems in fighter jets, Automated External Defibrillators (AED), airport signal light controls, and so on.

“We also see evidence that adversaries are moving beyond basic re-marking counterfeiting schemes to actually cloning functional integrated circuits. Some of these clones perform better at incoming inspection electrical tests than the original OCM version. These clones are currently targeting older, less complex components, but their learning curve is fast.”

Intel has been researching and implementing methods to cost effectively mitigate counterfeiting risk for decades, for example, the [CPUID](#) instruction was added in 1993.¹ Most of Intel’s anti-counterfeiting research is shared through industry consortiums like the Semiconductor Industry Association (SIA), and standards development organizations like SEMI Standards International, ANSI, JEDEC, SAE, and ISO.

1 A software program can use the CPUID instruction to determine the factory tested speed and performance levels. This allows end uses to verify they received the CPU performance levels they paid for.

As active as Intel is in security and privacy industry consortia which include many Standards Development Organizations at industry, national, and international levels, Intel cannot be directly involved in every standards action worldwide. Studying standards developed by others is an important opportunity to discover and learn about new ideas.

Supply Chain Risk Management Constantly Evolves

One standard that Intel has studied is NIST Special Publication 800-161: *Supply Chain Risk Management Practices for Federal Information Systems and Organizations*, with a particular focus on the cyber supply chain sections [SA-10, SA-11, & SA-12]. Taken together, these sections identify 25 unique controls. [See Table 1].

Table 1. Controls listed in SA-10, SA-11, & SA-12 of NIST SP800-161

SA-10: Developer Configuration Management	SA-12: Supply Chain Protection
1. Software/firmware integrity verification	1. Acquisition strategies/tools/methods
2. Alternative configuration management processes	2. Supplier reviews
3. Hardware integrity verification	5. Limitation of harm
4. Trusted generation	7. Assessments prior to selection/acceptance /update
5. Mapping integrity for version control	8. Use of all-source intelligence
6. Trusted distribution	10. Validate as genuine and not altered
SA-11: Developer Security Testing and Evaluation	11. Penetration testing/analysis of elements, processes, and actors
1. Static code analysis	12. Inter-organizational agreements
2. Threat and vulnerability analysis	13. Critical information system components
3. Independent verification of assessment plans/evidence	14. Identity and traceability
4. Manual code reviews	15. Processes to address weaknesses or deficiencies
5. Penetration testing/analysis	
6. Attack surface reviews	
7. Verify scope of testing/evaluation	
8. Dynamic code analysis	

Intel analyzed these 25 controls against two key criteria:

1. What is the dominant risk or problem corresponding to the control?
2. How well do Intel's existing controls and processes mitigate these risks?

Dominant Risk Addressed by Controls: Risk experts at Intel found that there is not a one-to-one mapping of risks and controls. Many controls mitigate more than one risk and many risks are mitigated by several controls. The analysis identified 37 distinct risks or threats that relate to the supply chain controls recommended in the NIST guidance. [See Table 2.]

Table 2. Risks derived from Control Mapping

- | | |
|---|--|
| <ul style="list-style-type: none"> ▪ Malicious Firmware ▪ Malicious OpSys ▪ Refusal to support system ▪ Inability to resolve major bug ▪ Incorporation of unfit IC's ▪ Incorporation of unfit assembly ▪ Implementation allows known vulnerability ▪ Implementation lacks a consistent threat mitigation strategy or policy ▪ No testing for common vulnerabilities ▪ Log files get edited, removed or corrupted ▪ Undocumented configuration changes occur ▪ Maintenance spare parts get substituted, tainted, or corrupted ▪ Debug capability gets added or unlocked after final acceptance testing ▪ Programmers are allowed to use insecure coding practices ▪ Obvious problem-code goes undetected because no one is looking for it ▪ Development teams lack experience in how adversaries typically attack ▪ Customer cannot review what was done to evaluate security robustness of platform ▪ Customer requires evidence that attacks [risks] of concern are mitigated ▪ Customer requires evidence that supply chain attacks [risks] of concern are mitigated | <ul style="list-style-type: none"> ▪ Development teams have blind spots and are unaware of their own blind spots ▪ Automated checkers will not have tests for all attacks. ▪ Automated checkers also have blind spots ▪ Developers may not understand where and how adversaries penetrated similar systems ▪ Supplier's supply chain may have a needlessly large attack surface ▪ Vulnerabilities caused by careless suppliers are not detected and mitigated ▪ Supply chain risk assessments are not performed ▪ Suppliers fail to mitigate identified supply chain risks because there is no supplier benefit to mitigating the risk ▪ Supplier is unaware of a risk known to others and is not informed about that risk so the risk remains unmitigated ▪ Loss of shipment ▪ Harm to delivery personnel ▪ Tampering during shipment ▪ Insider attacks supply chain ▪ Regulations or contracts prohibit exchange of information that would be useful to mitigate a risk ▪ Regulations or contracts allow access that creates a risk ▪ Denial of service because a critical part is not available ▪ It takes too long [or cost is high] to find the one [or a few] bad part among a large majority of good parts ▪ The ripple effects of one bad part cannot be contained and trash the entire platform |
|---|--|

Effectiveness of Existing Risk Mitigation: To answer the second question (How well do Intel’s existing controls and processes mitigate these risks?), content experts at Intel were tasked to perform a risk assessment against the 37 risks and threats shown in Table 2. Risks were considered across several different disciplines such as: silicon hardware design, platform hardware design, firmware development, software drivers & utilities, manufacturing, and shipping.

Outbound Supply Risk Assessment Results

Intel believes that threats and risks are constantly changing and evolving. In particular, experts at Intel anticipate three risks listed below will increase in the future. Intel is increasing risk mitigation controls in these three areas:

- Infiltration of malicious firmware
- Infiltration of counterfeit sub-assemblies
- Infiltration of counterfeit Integrated Circuits (ICs)

Risk of malicious firmware or other firmware corruption

Some of the controls in place for Intel products that help to mitigate firmware corruption risk include:

- Hardware features that limit unauthorized modification of firmware.
- Hardware that dynamically checks that firmware is digitally signed before it gets loaded.

On server platforms where Intel has permanently attached a [Trusted Platform Module \(TPM\)](#)² on board, it is adding control steps that capture the TPM Public Endorsement Key. A Platform Certificate signed by Intel will provide credible provenance information about the hardware and firmware that shipped from the Intel factory.

More hardware protections are being developed that further mitigate threats to firmware, including remote/automated recovery in case of an attack. Intel is consistently innovating new features that will be included with future products.

² A device used as a hardware root of trust.

Risk of using counterfeit or tainted Sub-assembly

In the late 1990s, Intel found that increasing transparency reduces the risk of counterfeit product infiltrating a supply chain. This was first observed on Intel® Pentium® III processors, with the introduction of the [Intel® Processor Frequency ID Utility](#). The tool made false claims about the intended CPU frequency obvious to customers, and Intel observed a 97 percent reduction in CPU counterfeiting as the market moved from Intel® Pentium® II processors to Intel Pentium III processors. The Intel Pentium II and early Intel Pentium III processors were multi-chip assemblies build onto printed circuit boards, which properly fits the definition of a sub-assembly.

In the 2000s, Intel introduced another method that brings transparency to network interface cards (which are also sub-assemblies) where any customer can use a unique [QR code](#)³ found on each card to access provenance information at any time. The introduction of this technology was also followed by a dramatic decrease in Network Interface Card (NIC) counterfeiting.

The new control Intel is introducing to mitigate risk of counterfeit board and sub-assemblies is called Transparent Supply Chain (TSC). Transparent Supply Chain (TSC) provides the following services

- Capturing and archiving “As Built” component identity data for each assembly
- Linking component identity data to sourcing information
- [Platform Certificate](#)⁴ [only when TPM or equivalent root of trust is present]

TSC provides credible data that the parts and ingredients used to build an assembly were procured from an authorized channel. Analysis of counterfeit product incident reports in GIDEP and ERAI shows that a “Buy Only Authorized” procurement policy mitigates a majority of the counterfeit product infiltration risk. Our customers benefit because they have proof the systems hosting their critical business process and data were built using components traced back to OCM authorized sources.

3 A two-dimensional barcode that works well with smartphones apps.

4 Section 3.3 of TCG Credential Profiles standard.

Transparency is an effective deterrent because it decreases the time to detect counterfeiting. It only takes one end user that notices a product as delivered does not match with data reported by the manufacturer to launch an investigation. Once one end user reports the fraud, the investigation that may follow will push counterfeiters to exit the business. Making the data available to customers brings transparency which makes it very difficult for counterfeiters to sustain their fraudulent business practice.

Risk of using counterfeit or tainted IC

Counterfeiting attacks on Intel branded products fall into two major categories:

- Re-marking of an IC that Intel actually manufactured.
- Printing an Intel trademark on an IC that Intel did not make.

Historically, when a counterfeit IC was functional, it was a re-marked device. This is because manufacturing a functional IC required generation of a wafer mask set and the cost and effort to generate a mask set has shown to be an effective obstacle to creating modern high complexity ICs such as Intel® Core® processors and Intel® Xeon® processors. To mitigate the risk of re-marking, Intel developed the [Intel® Processor ID utility](#). When an Intel CPU passes this utility, it is unlikely that the CPU is counterfeit. The utility is a free download available from Intel support website.

Historically, if the counterfeit IC was not manufactured by the brand owner, that IC was non-functional. Again, this was because the cost and effort to generate a mask set was an effective obstacle. Unfortunately, the cost of a mask set is no longer an effective mitigation for low complexity ICs. Platforms that cannot upgrade to “modern” electronics have become exposed to risk of cloned ICs. Lower complexity devices, such as embedded controllers designed 20 years ago, are at risk of being cloned. With lower complexity devices, it is extremely important that devices are sourced from authorized distributors. To mitigate the risk of procuring counterfeit IC’s, Intel recommends working with manufacturers that are transparent about their sourcing practices and who provide credible evidence that they actually follow their sourcing practices.

The risk of procuring a cloned and functional IC increases when ICs reach the end-of-life phase of their lifecycle. Any deviation in the end-of-life safety stock ordering process can create future demand for the obsolete IC that exceeds the supply of safety stock. It is well documented that machines, platforms, and missions sometimes outlive the lifecycle availability of the ICs used to maintain them. There is a risk that the IC or component manufacturer will stop making

spare parts for a machine or platform long before that machine or platform is retired and decommissioned. When authorized sources for an IC are depleted, the IC cloning operations quickly move in.

The end-of-life counterfeiting threat can be mitigated by using an authorized legacy supplier. An authorized legacy supplier is licensed by the original component manufacturer [OCM] to continue production of obsolete ICs. In many cases, the OCM transfers proprietary design information about the ICs, enabling the authorized legacy supplier to manufacture additional trustworthy parts as needed in the future. In other cases, trusted legacy part providers can be licensed to reproduce trustworthy legacy spares when unexpected disaster causes unexpected demand. Authorized legacy suppliers are also a potential source of securely and properly-stored unused parts.

If, however, a platform is designed using parts that do not have a legacy support plan, the prospects for finding trustworthy spare parts in the future are not good. The existing controls to build or pre-purchase a lifetime supply of trustworthy spares are sometimes overwhelmed by unexpected events. Choosing ingredients and parts that do not consider legacy support may appear lower cost during the design phase of a platform lifecycle, but these choices can put those platforms at higher risk of using spare parts procured from the open market.

The National Defense Authorization Act for fiscal year 2012 also places the cost burden of replacing counterfeit electronic parts onto the contractor that supplied the platform. This burden remains on the contractor for the entire lifetime of the platform. Reducing the risk of a counterfeit part infiltrating a platform clearly benefits these contractors by reducing the risk of this burden. The concept of Transparent Supply Chain enhanced by improved legacy product support is a potential method to ensure a safer long term future for the platforms that Intel is designing and building today.

Finding an authorized source for any obsolete part is challenging the entire industry.

A Final Word about TSC

Component security controls at Intel are often based on quality controls, and it is difficult to distinguish between the two. Intel's focus on quality also helps to mitigate many security threats and risks. As an example, both quality and security controls call for event logs so security uses the log files originally developed as quality controls. The two controls are the same in that both call for creation of the log file. Where the controls differ is how the two groups act upon the same log file data. A quality control will use the log to initiate more frequent training or modify the period between maintenance cycles. A security control will use that same log file to guide a forensic audit to narrow the timeframes for examining security videos, or access logs into sensitive IP archives. Security wants adversaries to know the company is watching, and by crosslinking log files, it becomes very difficult to make unauthorized modifications without being detected.

A Final Word about Controls

Controls are not free, they add cost.

The counterfeiting threat increases dramatically when procurement policy is overly dominated by lowest price and the value of fitness-for-use is discounted. A counterfeiter can offer a more attractive price simply because they have little or no R&D overhead. A counterfeiter creates product on demand and has no inventory overhead. The material counterfeiters use to create fake ICs has not been properly stored and reclaim techniques often introduce damage. The authorized legacy suppliers hold inventory for many years and must recover those costs to stay in business. While a counterfeit IC may initially appear to cost less, the hidden costs that manifest later will soon exceed the perceived price savings.

Shortage of supply caused by natural disasters interrupting availability of [raw] materials can also increase the threat of counterfeiting. Whenever a part or ingredient is hard to find, counterfeiters move in. When a natural disaster interrupts supply, it is prudent to temporarily invoke additional controls during the time when normal supply chains are disrupted. The additional controls can be removed after sources recover.

In both cases, increasing transparency into the sourcing process increases the probability of detecting tainted and counterfeit materials as well as exposing any false claims of sourcing policy.