EDITOR: Michael Mattioli, michael.mattioli@gs.com

DEPARTMENT: SECURITY

The Next Security Frontier: Taking the Mystery Out of the Supply Chain

Michael Mattioli, *Goldman Sachs & Co., New York, NY, 10282, USA*

Tom Garrison and Baiju V. Patel, *Intel Corporation, Mountain View, CA, 94040, USA*

The modern technology supply chain is highly complex. Throughout the various stages of design, manufacture, assembly, transport, and operation, a compute device is subject to tampering (malicious or otherwise). This exposes end users and consumers of compute devices to a variety of risks at varying levels of impact. Two potential technology solutions, which aim to provide transparency and insight into the systems in use, are proposed. The success and adoption of these choices, or any other solutions developed, is highly dependent on the participation of ecosystem partners such as foundries, original device manufacturers (ODMs), and original equipment manufacturers (OEMs).

By the time a computing system (e.g., laptop, desktop, server, smart watch, tablet, etc.) is delivered to its intended end user, the sum of its parts has traveled through a highly complex supply chain. This supply chain, illustrated in Figure 1, includes diverse component suppliers, subsystem manufacturers, integrators, and original equipment manufacturers (OEMs).

The final product may go through several warehouses and may be transported/handled via several shipping companies before finally being received, unboxed, and deployed/used.

Considering ever-increasing threats to this supply chain; end users have a growing need to know that the final product they received is indeed the product they ordered. Unintentional mistakes/errors, poor handling, or intentional fraud are key risks at play. Additional risks may come from malicious actors, including nation states and well-funded criminal organizations who are motivated to tamper with systems in the supply chain. The consequences of these risks may include financial or reputational harm.

Frankly, many treat an arbitrary computing system as a "black box" and blindly trust the supplier(s) and transport methods. However, a growing portion of the population, such as financial services or government institutions and even consumers are keen to understand and gain insight into the systems they use to store and process sensitive information as well as conduct high-stakes transactions (e.g., financial, health care, education, etc.). They are actively taking steps to ensure that the computing system, as delivered, meets their risk profile and can fulfill their compliance, security, and performance requirements.

It is important for the modern technology supply chain ecosystem to take measures to ensure a growing portion of the population can trust the supply chain with increased accuracy and decreased cost. Improving transparency in the supply chain will help meet the need for security and quality assurance among broader portions of the population as both awareness and risks continue to grow.

KEY RISKS TO SECURITY

Any typical component or system changes hands dozens of times from inception to deployment and, ultimately, retirement. The supply chain is a continuous, evolving process and, as illustrated in Figure 2, may not end even when it leaves the end user's hands (e.g., recycle or donate).

Participants in the supply chain, such as the OEMs, also treat the role they play as intellectual property

0272-1732 © 2021 IEEE
Digital Object Identifier 10.1109/MM.2021.3072174
Date of current version 25 May 2021.

FIGURE 1. High-level, simplified depiction of the modern technology supply chain.

(IP) or some sort of "business secret," which makes it even more challenging to evaluate and manage risk.

Design

The first window of opportunity to insert risk into a computing system is by attacking the design and manufacturing of the individual components that will eventually comprise the system. Schematics can be altered, design tools can be compromised, and collaboration solutions can be manipulated to alter a component from its intended composition. Subsequent risks are introduced when end users work with the supplier, such as an OEM, to design and/or configure a platform (e.g., build-to-order systems). During such engagements, there is room for unintentional human error, such as misinterpretation of phone calls and emails. Risks might also include intentional malicious action that can cause harm in the process. Behind the scenes, the vendor also works with their respective suppliers, such as original design manufacturers (ODMs), to do similar functions such as source subcomponents and/or assemble components; this further exacerbates the potential for human error.

Assembly and Manufacturing

Further down the line, during assembly and manufacturing, there are more opportunities for malicious actions

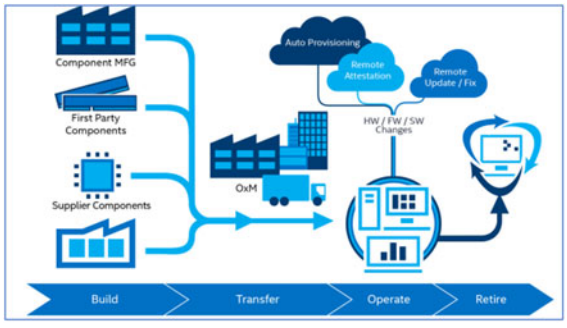

FIGURE 2. High-level depiction of the supply chain lifecycle.

such as circuit design modification (e.g., hardware trojans, scan chain attacks, etc.), firmware modifications, and even counterfeiting.[1] A disruption to the supply chain, such as a factory fire (local) or the 2020 COVID-19 pandemic (global), may also force OEMs to substitute parts to ensure timely delivery of computing systems to end users. The many levels and layers of the supply chain make it nearly impossible to keep a watchful eye on every single party involved and ensure that no harm, either intentional or unintentional, was done.

Shipment

After a system has left the factory, there are many opportunities for tampering, modifications, or changes within the hardware, firmware, and software. Once a device is received and deployed, it may be serviced (e.g., components replaced) in different locations or by different parties throughout its lifecycle. It is practically impossible to have complete confidence that a system has not been tampered with in some manner. The potential for tampering applies to each of the functional components of a computing system. For example, a solid-state drive (SSD) shipped to an ODM for integration into a computing system could be tempered with by having the firmware within the drive replaced with a malicious version.

Operation

Transparency in the supply chain helps the end user ensure that they received exactly what was ordered. This is a significant step in maintaining continuous trust in the computing ecosystem. Since most computing systems are frequently upgraded with the latest functional and security updates, it is extremely important to maintain the current state of the system throughout its entire lifecycle. Ideally, one should have the complete and latest information (e.g., firmware version updates) about each critical component of the computing system, and should be able to validate the version updates and configurations against the

expected state in real-time with reasonable assurance that nothing malicious has modified the system to an unacceptable configuration/state. Granted, this might be more feasible for large-scale commercial entities but is rather complex and dizzying for consumers; in either case, it is difficult to achieve.

Impact

Just as causes of unintended changes to the platform vary, the impact to the end user can vary from benign to more serious consequences rooted in malicious intent, such as the following.

- Financial—Stealing sensitive information can lead to financial gain through access to nonpublic information.
- Reputational—Exposing and fabricating information, or interrupting operations, can cause severe reputational damage to a person/company.
- Revenge—Personal attacks against a company or organization by a disgruntled employee or end user can cause potential damages ranging from physical alterations, compromise of IP, to destruction of reputation.
- Chaotic—Some individuals and criminal organizations simply want to "watch the world burn."

Different entities, private and public, are exposed to these and other threats at varying levels. It is imperative that supply chain transparency is designed to scale to not just different size businesses but to consumers as well over time.

POWER OF TRANSPARENCY

As might be expected, transparency has a cost. Beyond the obvious, managing the supply chain means keeping track of each component through manufacturing, warehousing, and delivery. This, in turn, limits the flexibility of the manufacturer.

For example, a manufacturer may wish to have multiple sources for a keyboard and would want to be flexible in selecting the keyboard to be integrated into a laptop based on the availability of parts and the cost at any given time. However, one may require or specify a keyboard from a specific supplier, which may impact a timely delivery and also introduce challenges and additional cost due to inventory management.

Inventory may also be impacted due to unexpected natural calamities such as the 2020 COVID-19 pandemic. In order to deal with such disruptions, as well as those caused by smaller unforeseen events, it is typical to include multiple sources for every significant component to support business continuity.

At the same time, suppliers also need to protect their confidential business information including complex business relationships, procurement processes, and inventory management. This type of information in the wrong hands can have significant impact to the business and can lead to compliance issues. Therefore, the ecosystem needs to align on the right technology as well as processes to minimize misuse of transparency while meeting end user requirements. Additionally, transparency needs to maintain a balance between the needs of supply chain and the needs of the end user and consider a baseline that may be available to everyone with additional levels of detail requiring different agreements.

TRUST REQUIRES INDUSTRY-WIDE PARTICIPATION

A computing system is commonly composed of various components. Many of these components can be classified as "smart"; these are components that typically contain their own firmware. Firmware has significant functional capabilities including the ability to access, and potentially manipulate, critical or sensitive (e.g., personal, mission-critical, etc.) data.

The smart components are typically manufactured by different suppliers and, in turn, are composed of both active and passive subcomponents. It is important that one has full transparency of active components due to their significant capabilities to access and modify data. Passive components, such as resistors, capacitors, physical packaging, printed circuit boards (PCBs), or displays can also pose a risk to one's data if not implemented correctly. For example, if PCB routing had certain wires close to the surface or implemented sockets for expansion or options, an adversary might be able to easily substitute or even add a component without the knowledge of the end user.

While it may not be practical to have full transparency at each passive component level, it is important, at a minimum, to have transparency at the active component level and to establish trust in the entire supply chain; this requires an industry-wide initiative and agreement on key technical underpinnings in order for the initiative to be successful. With the ultimate goal of achieving full transparency, we recommend starting small yet critical components of the computing system first and then develop technology and processes to progress ahead.

While we do not attempt to specify a comprehensive list of smart components for initial implementation,

some examples that the industry should consider as starting points are the CPU, SSD, wireless (e.g., Bluetooth, WiFi, etc.) interfaces, motherboard, and any/all embedded controllers (ECs). It is not enough to know the manufacturer of these components or subsystems, one needs to know details such as the firmware running on these components, the date and location of manufacture, and even the design revision.

PROPOSED SOLUTIONS

Technology choices we make to establish transparency will have long-term impact on the industry both from longevity and cost perspective. We propose two models which are 1) use of a ledger or database and 2) a method of self-reporting.

Ledger or Database

As a computing system goes through assembly and delivery in the supply chain, each entity makes an entry in a ledger, thus, creating a record that can be retrieved later by the end user to meet their need for transparency. Once a computing system is delivered to the end user, the same process can be used to record changes, such as a change of ownership, to maintain a continuous record for the entire lifespan of the computing system. Two fundamental approaches exist here.

Centralized Approach

A trusted third-party maintains the ledger or database of all the transactions and manages updates to and retrieval of information per agreement between trusted third party and members of the supply chain. The fundamental trust model is based on the idea that without collaboration between multiple member companies in the supply chain, any inconsistency (intentional or otherwise) might not be detected. For example, if a system manufacturer entry shows they shipped 1M computing systems with a specific SSD from a specific supplier, a corresponding entry by the supplier in their database is necessary to corroborate.

Blockchain Approach

Application of blockchain technology to ingredient supply chains has been widely discussed. A 2019 Fortune article outlined Bumble Bee Tuna's use of blockchain to create an electronic and distributed ledger chronicling a fish's capture, processing, and travel history to ensure freshness and product quality[2]. Beyond this ability to track ingredients, the distributed ledger can be expanded to keep a running and growing record over the operational lifecycle for each system. Tightly controlled and permissioned ledger entries can continue for updates, changes, and ownership transitions that occur as the system travels through distribution and integration and then, ultimately, provisioning, operation, and modification by the end user or owner. The private and public cryptographic controls inherent in existing blockchain infrastructures can be applied to balance the simultaneous needs of all ecosystem participants. Component suppliers can contribute product information into the blockchain without the worry of sharing sensitive product details to competitors. Similarly, system manufacturers can create system-level ledgers from component and subsystem vendors in private transactions shared only with those whose access permissions have been cryptographically proven. Ownership can be transferred to operators who can manage ledger updates to track key information regarding location, application, upgrades, updates, and usage statistics. The usefulness of this ledger can extend from the resale of the device into the secondary market; the ledger can provide sellers with better insights into possible IP loss based on the usage and application while providing buyers a better understanding of the provenance of the device.

Self-Reporting

Tracking the components in the supply chain while the device is manufactured and readied to be delivered to the end user helps ensure that intended components are used in the device. Once the device is in the hands of the end user, the device can provide a cryptographic report (i.e., attestation) of the current configuration of all of the smart components in the device. The configuration can include current firmware version, security version, hardware ID, etc. Each report may further include information about nonactive or passive components that are part of the smart component or even include a report-out of smart subcomponents as well as revision or change log for additional transparency.

In order to establish trust into the report, the smart component needs to include a cryptographic key and associated certificate, as depicted in Figure 3, and, when queried, produces a signed report along with the certificate so that the end user can validate the authenticity and trust in certificate and subsequently trust the report.

Finally, the entire computing system can either collect report-out from each of the smart components and may produce a comprehensive report for the entire system or may leave it up to the end user to query each smart component of interest. The benefit of a system-level report-out is that the end user gets the full picture without necessarily having to understand how to query each of the components (as the

FIGURE 3. Simplified illustration of computer components with a certificate of authenticity.

manufacturer of the system has the best understanding). This would require the system to also have a trusted smart component that cannot only collect all the information but is able to produce a cryptographically signed comprehensive report.

The shortcoming of this type of approach is that each component has an additional cost of obtaining a certificate and complexity of providing a report-out. Additionally, the end user does not have the benefit of transparency from any component that does not implement this capability.

GOVERNANCE

There are two potential paths toward solution adoption. The first is a self-regulated or market-driven approach where each participant in the supply chain may choose their own degree of participation in transparency. End users, based on their purchasing preference, will drive supply chain to adopt the "right" level of transparency to meet market needs. The other is by means of an industry group or consortium in which computing system supply chain participants form a consortium to establish technical direction and a governance model to provide consistent and sufficient transparency to meet broad industry needs.

RECOMMENDATION

With increasing reliance on information technology for business-critical and personal data, having transparency in the supply chain, as well as information on the current state of the computing systems in use, is vital to the economic health of the information technology ecosystem. Transparency by itself is of the limited value without the information needed to assess risk or compliance. Beyond the ability to verify that what the end user received is indeed what they ordered, there is an additional value of transparency post deployment as well. For example, if a security risk is identified (e.g., a publication in the common vulnerabilities and exposures system), the end user can use transparency to precisely pinpoint impacted systems and take appropriate remediation steps. Due to the diverse and complex nature of the ecosystem, no one company or institution can single-handedly solve this problem. The choice of technologies and collaboration will have long-term implication on evolution of capabilities to meet current and future needs as well as the ability to scale and manage costs. We strongly recommended that industry leaders come together and take a comprehensive long-term approach to address this important problem.

REFERENCES

1. M. Mattioli, "Consumer exposure to counterfeit hardware," *IEEE Consum. Electron. Mag.*, to be published, doi: 10.1109/MCE.2020.3023873.
2. Fortune, "Bumble bee foods aims to put all its fish on a blockchain. It's starting with 'fair trade' tuna," Mar. 2010. Accessed: Feb. 17, 2021. [Online]. Available: https://fortune.com/2019/03/08/tuna-blockchain-bumble-bee-sap/

MICHAEL MATTIOLI currently leads the Hardware Engineering team within Goldman Sachs, New York, NY, USA. He is responsible for the design and engineering of the firm's digital experiences and technologies. He is also responsible for the overall strategy and execution of hardware innovation both within the firm and within the broader technology industry. Contact him at michael.mattioli@gs.com.

TOM GARRISON is currently the Vice President and General Manager of Security Strategy and Initiatives within the Client Computing Group, Intel Corporation, Mountain View, CA, USA. He leads efforts to help customers and manufacturers deploy tooling and processes for greater security assurance, supply chain transparency, and cybersecurity innovation. He also launches industry-wide initiatives and research with ecosystem partners and academia to accelerate cybersecurity product assurance. Contact him at tom.garrison@intel.com.

BAIJU V. PATEL is currently an Intel Fellow within the Client Computing Group, Intel Corporation, Mountain View, CA, USA. He is responsible for setting the technical direction for the Client Computing Group and Intel's security technologies. Contact him at baiju.v.patel@intel.com.