Industrial Agents Cybersecurity

Stamatis Karnouskos
SAP, Germany
Email: stamatis.karnouskos@sap.com

Contents

1 Introduction

Industrial Agents (IA) are considered as a key enabler for industrial applications (Leitão et al., 2013) and therefore sophisticated approaches have been developed over the past years. Probably the most well known example of such system is the one deployed in the factory of DaimlerChrysler as analyzed by Schild and Bussmann (2007). However, it is noticeable that most of the existing approaches focus on the provision of core functionalities relevant to the application, while other aspects, such as security and privacy, that are not immediately visible are considered as second class priorities and are often neglected or realized only at very basic level. With the emergence of Cyber-Physical Systems (ACATECH, 2011; Colombo et al., 2014), and especially their application in industrial domain, the business landscape is changing, as they offer sophisticated capabilities that may be transformed to competitive business advantages. However, as

Porter and Heppelmann (2014) point out, in such environments, underestimating the security and privacy pose one of the greatest strategic risks.

Due to increasingly sophisticated security threats (Cheminod et al., 2013), it has been repetitively shown that industrial systems are largely becoming vulnerable and so is the critical infrastructure they control. However, although awareness is raising, dealing effectively with these is still not adequately addressed. Security, trust and privacy are such aspects that also in the Industrial Agent domain are not given the appropriate importance and considered usually as a future add-ons, once the Industrial Agents achieve their breakthrough. Although this may have been somehow acceptable some years ago, where their utilization e.g. in factories was done in highly-controlled and isolated environments with low probability of misuse, today we are far away from such "safe-haven" systems. The recent Stuxnet worm (Karnouskos, 2011) exposed the vulnerability of modern industrial systems even in the most controlled environment of a nuclear facility. In addition, the penetration of Internet technologies and concepts, the amalgamation of industrial networks and IT systems, as well as the need for tackling increasingly complex industrial systems with common means, has increased the risks introduced to and by Industrial Agent systems.

2 Technology Trends and Industrial Agents

To better understand the transformation on industrial systems and how this affects Industrial Agent approaches, we have to consider the vision of future industrial systems (Colombo and Karnouskos, 2009; Kagermann et al., 2013) as well as the trends in technologies to realize it. Today we see an increased penetration of Internet technologies in industrial settings and an amalgamation of the different concepts and technologies on enterprise and shop-floor (Colombo et al., 2014). Some key trends we witness include:

- Information Driven Interaction: Future integration will not be based overwhelmingly on the data that can be collected and delivered, but rather on the services and intelligence that each device/system can deliver to an infrastructure. These information points will be distributed and provide local intelligence (including monitor and control capabilities) via well defined interfaces while their interworkings are hidden. The interactions that happen among them, will give emergence to system wide characteristics and capabilities. Industrial Agents fit well in this role due to their characteristics. Security though will be critical as the task to empower modern scenarios that rely on such interactions without revealing key competitive advantages to other parties is challenging.

- Distributed Business Processes: In large scale sophisticated infrastructures, business processes can be distributed in-network e.g. in the Cloud and on the device. Thus processing of information and local decisions can be done where it makes sense and close at the point of action, while only necessary info is propagated to higher levels for system view. Industrial Agents provided with the right capabilities and resources, can host the logic to execute business processes and become part of complex orchestrations. However, how to securely and efficiently outsource such functionali-

ties to Industrial Agents, especially in enterprise-wide and cross-enterprise scenarios still needs to be properly addressed.

- Cooperation: Highly sophisticated networked devices are able to carry out a variety of tasks not in a standalone mode as usually done today, but taking into full account dynamic and context specific information. As such we see the emergence of a highly distributed intelligent infrastructure that is able to cooperate, share information, act as part of communities and generally be an active element of a more complex system (Marrón et al., 2012). Industrial Agents can be seen as an add-on to such devices which can take over management of interactions and cooperation. Security aspects relevant here are manifold including modern research in reputation systems and building of collaborative systems and infrastructures.

- Cloud Computing and Virtualization: Virtualization addresses many enterprise needs for scalability, more efficient use of resources, and lower Total Cost of Ownership (TCO) just to name a few. Cloud Computing is emerging powered by the widespread adoption of virtualization, Service-Oriented Architecture and utility computing. For Industrial Agents this is of relevance as now resources can be dynamically adjusted to the needs of an Industrial Agent for execution of a scenario. This means that Industrial Agents can cohabit resource constrained devices and systems while in parallel outsource more demanding (resource consuming) functionalities to the cloud. However this strong dependence and communication between the local and cloud Industrial Agents may raise some security concerns or may not be wished or appropriate e.g. in critical infrastructures.

- Multi-core systems and GPU computing: The last ten years we have seen the rapid prevalence of multi-core systems that nowadays start to dominate not only everyday devices but also traditional embedded industrial systems. The general trend is towards chips with tens or even hundreds of cores, simultaneous multi-threading, memory-on-chip, etc. which promise high performance and a new generation of parallel applications unseen before in embedded systems. Additionally in the last decade we have seen the emergence of GPU computing where computer graphic cards are taking advantage of their massive floating-point computational power to do stream processing. For Industrial Agents and generally multi-agent systems this adds new capabilities especially towards running complex simulation scenarios. For instance specific analytics on an embedded device can now be realized at high performance in the GPU which empowers new Industrial Agent applications at the edges.

- Infrastructure Servicification: Service Oriented Architectures have penetrated modern infrastructures from larger systems down to even simpler networked embedded devices. As the latest have become more powerful with respect to computing power, memory, and communication, they are starting to be built with the goal to offer their functionality as one or more services for consumption by other devices or services. Due to these advances we are slowly witnessing a paradigm shift where devices can offer more advanced access to their functionality and even host and execute business logic, therefore effectively providing the building blocks for

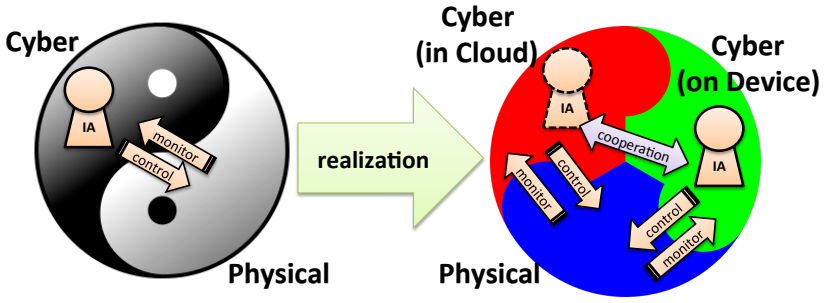

Figure 1: Industrial Agents and Cyber-Physical Systems

expansion of Service-Oriented Architecture concepts down to their layer. Web services are suitable and capable of running natively on embedded devices, providing an interoperability layer and easy coupling with other components in highly heterogeneous shop-floors. Industrial Agents can take advantage of such a SOA-based infrastructure and build more sophisticated approaches on top. However this also increases the requirements for trust, security and potentially also privacy.

- Trust and Privacy: As the infrastructure becomes more complex and we move away from monolithic systems that host the fully-fledged functionalities, towards cooperative and modularly-built systems, so does the dependence on key requirements among them including trust and in some scenarios privacy. Especially the privacy issues have not been adequately tackled when it comes to the Industrial Agent scenarios, mostly because up to now operations were carried out in strongly controlled environments where the majority of data and processes was owned by a single stakeholder. However with the increased generation of data due to the Internet of Things new approaches such as analytics and simulation can be realized based on distributed cross-enterprise real-world data. Hence introducing and enforcing a full policy-driven data lifecycle management remains a grand challenge. For Industrial Agents this becomes relevant as they need to operate on large datasets but also respect policies and privacy-preserving approaches, while in parallel make also sure that their operations do not leak or provide information that might be misused.

- Cyber-Physical Systems: Although the majority of Industrial Agent systems up to now was realized in software with some but limited integration in hardware, the advances in networked embedded devices in industry the last years, indicate that this is already changing. The significant decrease on hardware prices with the parallel increase in the computational and communication resources it may possess, have given rise to several systems that can be largely summarized under the Cyber-Physical Systems (CPS) domain. These, go beyond traditional stand-alone monolithic systems that could have some intelligence and be empowered by Industrial Agents. On the contrary they are multi-faceted multi-layer entities (both in hardware and software) that are highly complex and can operate autonomously but also in cooperation with other systems both on-premise

and out-of-premise. The latter is empowered by the usage of Internet technologies and connectivity, including the cloud paradigm. As such Industrial Agents have assumed new roles in CPS, and not only can execute in-CPS but also rely on external entities e.g. for activities offloading, cooperation and wide-area management.

As depicted in Figure 1, we see a shift to the realization of Industrial Agents. Up to now these were mostly software solutions with some management/control capabilities on the underlying hardware (as shown in the left side of Figure 1). With the prevalence of the Cloud and Internet technologies the intelligence of a single Industrial Agent can now rely both on-device and in-cloud, creating a cooperation link among its different parts (as shown in the right side of Figure 1) that may lead to a better solution. The latter implies that Industrial Agents have now to operate as part of a much more complex system; hence naïve approaches especially related to security are neither contemporary nor realistic. Figure 2 presents an overview of some potential threats within the operational context of Industrial Agents, which we will investigate more closely.

3 Agent Threat Context

Security in software Agents is in general a challenging issue, and several considerations are made (Jansen and Karygiannis, 1999; Karnouskos, 2001; Mcdonald, 2006) including potential dependence on operational conditions, applications etc. The security threats arise (as also shown in Figure 2) due to the special properties agents usually possess and utilize (e.g. autonomy, mobility, code execution etc.) which leads to key threats common in mobile code that transports and executes itself. Although the examples given below are not exhaustive, they should provide a general basis for understanding of threats relevant to Industrial Agents.

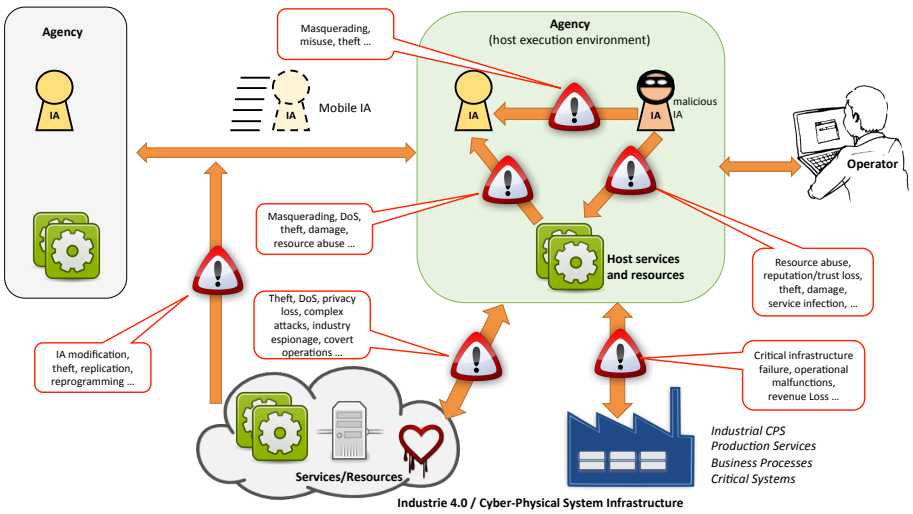

Figure 2: Industrial Agent Threats

3.1 Misuse of agent(s) by the host

All agents execute in an environment installed on a host. Although certain guarantees can be made about the execution and solutions exist, the agent has to place some trust on the infrastructure and services provided to it by the host, which may lead to security compromises and incidents. Some example attack scenarios may include:

- Masquerading: the deception of the agent in order to acquire its internal information. If the agent can not reliably verify the host environment it executes as well as the services offered, it may release information intended for third parties. In some scenarios, even the execution of the agent itself reveals its internal processes which also may be of use (for further attacks or re-engineering).
- Denial of Service: unacceptable delays may be introduced by the malicious host environment during the execution of the agent. The result might be inability to use external services, unreliable or slow operation, and other factors which may induce the purpose of the agent useless. In addition of course other attacks such as suspending the agent or even deleting it might lead to results that defy the purpose of the agent application.
- Eavesdropping: internal and external communication and states may be monitored which may provide direct access to the data and operations of the agent. This opens the door also for further attacks such as the malicious cloning of agents etc.
- Cloning/Replacement: An agent whose internal behavior can be replicated may be replaced by a malicious one, who can participate then in covert operations e.g. collect further data in the system, operate in other inaccessible up to now to the attacker environments and execute malicious commands replacing the original ones etc.
- Agent manipulation: a malicious host may be able to interfere with the normal execution of an agent and manipulate selectively its state and data in order to guide its behavior. In such scenarios the agent may be under the impression that its goal was achieved, which however may not be fully correct as the solution might not be optimal, or the data upon a decision was made might be false (but falsely considered trustworthy).

3.2 Misuse of the host by agent(s)

Malicious agents may scan and identify security weaknesses in the host environment. Subsequently attacks may be performed once the possibilities are analyzed. The latter might include:

- Damage: if given access, the agent may modify/reconfigure resources such as disk files, policies, network access etc. which effectively impacts all other agents executing at that moment.
- Masquerading: the identity of a trusted entity might be claimed, and unauthorized access to data may be obtained. Such misbehaviors may damage the reputation of the host and lead to trust loss as well as further attacks.

- Denial of Service: malicious behaviors may trigger security countermeasures on the host side which will result in potential disruption of offered services and their functionalities, which will have an effect on the operation of the platform and the legitimate users.
- Security breach / Theft: the identification of security holes may lead to further security breaches as well as be the starting point of malware installment which will subsequently "turn" the host to a malicious one where further attacks can be performed to agents and in-network.

3.3 Misuse of an agent by another agent

Malicious agents may pose a threat for other agents executing in multi-agent systems. Such threats may go unnoticed by the host platform and have significant impact on the victim agent as well as the host functionalities. Examples include:

- Repudiation: the malicious agent after negotiation can deny its participation in a transaction or communication it took part. This may result in conflicts and misuse of resources and services.
- Denial of Service: the malicious agent may overwhelm with interactions the victim agent and consume its available resources. The latter might result to inability of the victim agent to function properly and even high costs due to resource usage.
- Masquerading and misinformation: the malicious agent may disguise its identity and perform actions that will effectively beat the purpose of existence of the victim agent and it take the blame. This can result to trust and reputation loss, especially in communities where this matters e.g. electronic marketplaces where price negotiation takes place.

3.4 Misuse of agent(s) or host by underlying infrastructure

Although the most common attacks involve the agents and the host (and their interaction patterns), attacks could also happen outside the agent environments e.g. in the underlying network infrastructure (both at software and hardware level). The later rely on operating system and other layers of abstractions and may be practically undetectable from the agent or its host execution environment. Typical examples of such attacks include monitoring of communication, replay attacks, cloning of agents and host in order to study their behaviors/strategies, modification of agent system data and state etc. Especially the hardware-based attacks are given little attendance. However the last years we have witnessed the rise of several USB based attacks Clark et al. (2011); Davis (2011), as well as others involving the Ethernet card Duflot et al. (2010), or even the battery Miller (2011) etc. Such attacks pose a wide spectrum of potential threats and are not specific only on the agent systems but generally to any software executing on the specific node.

3.5 Complex attacks

Although many other cases could be described, there are several initiatives for trusted code execution in other domains, that strive towards solving these and similar issues (Jansen and Karygiannis, 1999). Most of the aforementioned threats, assume that the attacks are working in standalone mode. However, more complex attacks are usually collaborative and distributed, which makes it much more difficult to detect and react to them. In collaborative attacks two or more entities are working together towards common goals. Such entities might be agents or a combination of agents, malicious hosts and other services.

Complex attacks usually provide a high level of sophistication e.g. may be event triggered. For instance they may start when a specific event such as time, location, agent identity, agent payload, etc. occurs. These threats may not always be identified on-time as scanning of the agent code may only partially help, since the pattern interaction among the agents and other services under specific conditions needs to be considered.

4 Requirements on Industrial Agent Solutions

Industrial Agents may suffer from the security threats common to all software agents, but there are also differences related to the operational context where they are utilized in industrial environments. In Industrial Agents, the emphasis is put on the specific requirements that need to be fulfilled, sometimes at all cost, such as reliability, fault-tolerance, scalability, industrial standard compliance, quality assurance, resilience, manageability and maintainability etc. Depending on the scenario where Industrial Agents are used, these requirements may have varying degrees of importance and the focus is on well-established, stable and proven approaches rather than experimental and not fully tested features. Also industrial solutions need to fully guarantee business continuity as well as compliance to quality and legal requirements posed on the industrial domain where they are utilized. Therefore, technology as such is not the only criterion, but rather the whole operational context and lifecycle of the Industrial Agent solution is considered.

Each Industrial Agent system solution naturally has to support the requirements set by the respective cases. While most functional requirements may be case specific and security should be integrated directly in their design, implementation and operation, there are several other non-functional requirements that usually industry considers. These industrial requirements may differ to the degree in which they are important per case, however, these usually significantly differ from the ones imposed in simple prototypes and proof of concept operations, as they need to be deployed in productive environments and adhere to their operational context. Examples of these include:

- Code Quality: Software companies developing industrial solutions have standards they adhere to, in order to guarantee the quality of the developed solution. While typical development pitfalls can be avoided, such as insecure practices which would enable the Industrial Agent threats mentioned to apply, there are also other motivations such as maintainability, easy refactoring of libraries, consistency of features, configurability, easy logging/debugging etc.

- Maintainability: The solution has to be easily maintainable which implies modularity of the developed approach, incremental updates, minimization of downtime, on-the-fly feature enablement, testability etc.
- Policy Compliance: As any solution used in productive industrial environments, also the Industrial Agent solutions needs to adhere to the policies set by the organization and comply to the requirements. However, matching these policies to the interaction patterns of Industrial Agents is challenging and requires expert knowledge. The balance between security and operational aspects that adhere to the policies needs to be considered already starting from the design phase of Industrial Agents.
- Upgradeability: Industrial Agents solutions have all the advantages as well as the security threats of modern mobile devices and software. As such the unattended upgrades of their functionalities (agents) as well as those of their execution environment (host platform) are of high priority.
- Manageability: Industrial Agents, independently if they are static or roam the network, interact with systems and services and collect, store, and transmit business relevant (and potentially critical) data such as location information, process data, critical infrastructure measurements etc. These should be protected and securely managed e.g. with utilization of encryption or secure communication. They should also be easily integrateable in the existing management infrastructure of the organization.
- Auditability: Industrial Agents perform a multitude of functions e.g. they interact with systems, perform management actions, control physical systems, negotiate contracts, etc. As such the auditability of their operations is often a requirement, especially when interacting with third party systems and services. Not only security but also trust are key issues here.
- Safety: Considering that Industrial Agents have been integrated within or interact and manage physical systems (Mařík et al., 2005), and that the later operate in critical infrastructures or factory shop-floors, safety is considered a high priority requirement. Any security breach or misbehavior of on the industrial agent side, may have real-world consequences and threaten the safety of employees and infrastructure.
- Extensibility/Modularity: Industrial Agents have high negotiation skills and can easily interact with Internet based services, which calls for robust modular approaches that extend their functionalities based on the available services. Although this increases certain qualities of the Industrial Agents, it also creates a dependence on the infrastructure services which may be misused as we have already discussed.
- Performance: for many industrial scenarios, the performance of the agents is critical as it dictates the performance of the system. Especially in cases where production lines are controlled or near real-time decisions need to be made, the performance of the Industrial Agents is one of the highest priorities. Many security and performance tradeoffs may be considered (Zeng and Chow, 2013), depending on the concrete requirements.

- Reliability: Industrial applications have to operate reliably and in a deterministic manner. A crash or misbehaving Industrial Agent, may result to physical damages in the factory and of course to financial impact due to damage, delays, system reconfiguration, maintenance, etc.
- Usability: For solutions interacting with users e.g. operators, engineers etc., several aspects need to be considered as they directly impact productivity, training and support costs, development time and costs, maintenance, customer satisfaction etc. Today, these aspects are largely ignored by Industrial Agents, and the focus is mostly on functionality.
- Energy Efficiency: in specific scenarios, the Industrial Agents operate within resource constrained devices and therefore they must ensure the lowest impact to its resources (computation, communication, memory etc.). However this is challenging as the agent must have also an understanding of its operational environment and the energy impact of its actions.

As we can see, some of the example requirements mentioned (which constitute in no way an exhaustive list), can have a significant impact on the design, implementation, operation and acceptance of Industrial Agent solutions. Security, trust and privacy though, touch directly or indirectly on all of these, and although some overhead might be imposed, not considering them in the solution realization is not an option for systems used in production environments.

5 Discussion

Security is a process and as such (i) tradeoffs are inevitable and (ii) the question is not if an incident happens, but how to timely identify it (Vollmer and Manic, 2014) and effectively deal with it. To this end prioritized security goals and consistent security policies must be in place and be respected by the Industrial Agent solution. Secure activity logging, as well as real-time monitoring, anomaly detection and analytics could help in the early identification of security breaches.

Traditional security measures e.g. protection with firewalls, honeypots, known attack scanning etc. although necessary are not seen as enough. Especially in the era of IPv6 where globally unique IP addresses per device are supported, security and privacy issues should be revisited. In light of the new security, mobility and quality of service features offered by the protocol (e.g. IPsec, enablement of privacy extensions etc.) as well as their utilization in industrial agent scenarios, there is a need to have a holistic understanding of efficient usage of the offered capabilities as well as how they can be misused.

Effective security can be achieved at high degree when security considerations and good practices can be integrated in the lifecycle of the Industrial Agent solution. This includes:

- Industrial Agent systems requirements and use cases: These need to be properly defined (Mead et al., 2009), and weak points should be identified. Detailed scenarios and diagrams should be documented, that provide clarity on the functionalities and actor interactions. Subsequently “threat” cases can be defined, clearly depicting misuse potential in the system and its operations.

- Industrial Agent systems design and implementation: During design and implementation, concrete technologies and interaction patterns come into realization. Detailed system architecture diagrams and attack trees should be defined and documented. Here also technology-specific analysis should be performed to guarantee also that the implementation is not exposing the solution to threats. Typical security actions including code reviews, modular usage of software and incremental updates are examples that could be considered.

- Industrial Agent systems operational threat and vulnerabilities identification: While secure design and implementation may be realized, this does not guarantee also a secure operational phase. As such detailed monitoring, penetration testing and risk analysis should be carried out, identifying additional potential cases for misuse.

- Industrial Agent systems risk analysis, impact and mitigation: the extend of security breaches has to be considered, and the impact on the productive systems has to be assessed, including the business relevant impact. Subsequently mitigation plans have to be put in place that guarantee business continuity and resilience.

Industrial Agent solutions and their operation has to follow common best practices for securing information technology systems. This implies adherence to key elements such as those identified by Swanson and Guttman (1996). Industrial Agent solutions will also need to be largely aware of the operational context and this includes multiple security considerations such as:

- Agent-based security: this includes both agent as well as agent host execution environment relevant aspects. As such considerations should be made towards attack detection (side/covert channels, communication patterns etc.), resilience and availability, code security etc.

- Network security: network services, communication, topology, discovery, routing etc.

- Hardware security: trusted execution hardware platform, firmware attacks, tamper detection, security function offloading, cryptoprocessors etc.

- Data security: including repudiation, trust, integrity, privacy, authorization, lifecycle management etc.

- User security: including awareness of system's capabilities and threats, integration of user feedback etc.

Security safeguards need to be in place, not only on individual CPS hosting the agent or interacting with it, but also on the processes in which they participate (Karnouskos, 2014). This requires system and potentially system-of-system wide behavior monitoring and checks for anomalies (Pereira et al., 2013). Heuristics for estimating behavior deviation may provide hints, which should be assessed and analyzed in conjunction with other metrics. This is challenging but probably achievable to some degree if the process is under the control of a limited number of stakeholders. However, in the envisioned widely collaborative CPS systems-of-systems this is a daunting task.

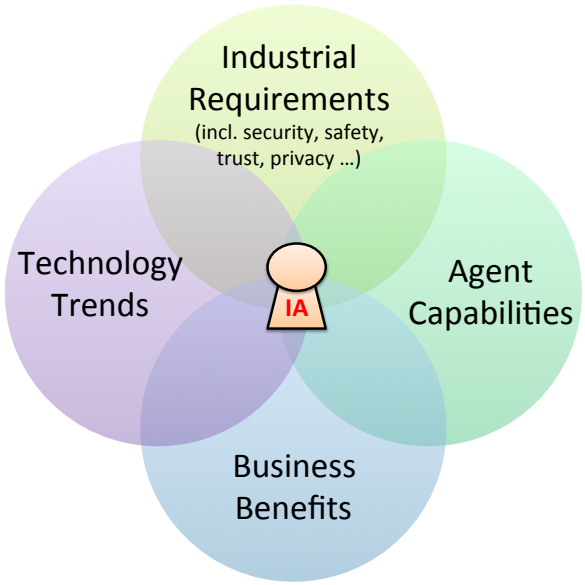

Figure 3: Target space of promising secure Industrial Agent solutions

Software and hardware security are not the only issues to be considered; human users must be included in the process (Karnouskos, 2014). Security clearance on people does not imply security on their accompanying assets. In the Stuxnet case (Karnouskos, 2011), a trustworthy employee with an unknowingly rootkited laptop or an infected USB flash drive would be enough to spread the virus. This could be, for instance, a contractor carrying a personal device, who is assigned to do maintenance on a facility. Perceived trust and risk assessment (Patrick, 2002) are seen as key aspects to be considered when designing, deploying and operating Industrial Agent solutions. Risk assessment should also include a survivability analysis for the threats, mitigation strategies as well as impact analysis e.g. on operational aspects. The latter is also of key importance for industrial systems, as most of them are connected to real-world processes and any malfunction has direct consequences on business processes, operations and finances.

To be able to see the potential misuse, one has to be well acquainted not only with general good practices of security management and coding, but also understand the capabilities and potential of specific Industrial Agent technologies and systems that use it. Failure to do so, will result probably in ineffective enterprise-wide strategies or to the enforcement of constraints which might be ineffective or severely limit the benefits brought by the Industrial Agent solution. The latter can have a significant impact on the acceptance of Industrial Agent solutions overall, as we are still at early stages of its widespread usage in industrial productive systems.

Finally, deciding on the adoption of Industrial Agents has to do with the tangible business benefit it will bring to the production environment where it will be utilized. Hence, the targeted space of promising Industrial Agent solutions are seen in the common space defined by industrial requirements (including

security, safety etc.), agent capabilities, technology trends, and tangible business benefits as depicted in Figure 3. This aspect is pointed out also by Schild and Bussmann (2007), who also mention that "in different industries the same system may have a quite different economic impact". As such we conclude and reinforce the view that a security-enabled holistic view is needed.

6 Conclusions

Agent technologies in general as well as Industrial Agents have been with us quite some time. However, up to now we have seen limited utilization in industrial productive environments, while several use cases have been successfully demonstrated in labs and for research purposes. As we have analyzed, key technology trends and especially cyber-physical systems provide another chance for Industrial Agents, as the latter could act as enablers in several aspects of the emerging Industrie 4.0 infrastructure (Kagermann et al., 2013) and play pivotal role towards achieving that vision. However, to do so, security aspects need to be properly addressed for the whole lifecycle of the Industrial Agent systems. We have already investigated several threats that may arise directly or indirectly with the operation of the technology, and how additional requirements of industrial systems should be considered, if Industrial Agents are to be widely accepted and used in real-world industrial settings.

References

ACATECH (2011), Cyber-Physical Systems: Driving force for innovation in mobility, health, energy and production, Technical report, ACATECH – German National Academy of Science and Engineering.
URL: *http://www.acatech.de/fileadmin/user_upload/Baumstruktur_nach_Website/Acatech/root/de/Publikationen/Stellungnahmen/acatech_POSITION_CPS_Englisch_WEB.pdf*

Cheminod, M., Durante, L. and Valenzano, A. (2013), 'Review of security issues in industrial networks', *Industrial Informatics, IEEE Transactions on* **9**(1), 277–293.

Clark, J., Leblanc, S. and Knight, S. (2011), 'Compromise through usb-based hardware trojan horse device', *Future Gener. Comput. Syst.* **27**, 555–563.
URL: *http://dx.doi.org/10.1016/j.future.2010.04.008*

Colombo, A. W., Bangemann, T., Karnouskos, S., Delsing, J., Stluka, P., Harrison, R., Jammes, F. and Martínez Lastra, J. L., eds (2014), *Industrial Cloud-based Cyber-Physical Systems: The IMC-AESOP Approach*, Springer. ISBN: 978-3-319-05623-4.
URL: *http://www.springer.com/engineering/production+engineering/book/978-3-319-05623-4*

Colombo, A. W. and Karnouskos, S. (2009), Towards the factory of the future: A service-oriented cross-layer infrastructure, *in* 'ICT Shaping the World: A Scientific View', Vol. 65-81, European Telecommunications Standards Institute (ETSI), John Wiley and Sons.

Davis, A. (2011), USB – undermining security barriers, *in* 'Black Hat USA 2011, Las Vegas, NV, USA'.
URL: *http://media.blackhat.com/bh-us-11/Davis/BH_US_11-Davis_USB_WP.pdf*

Duflot, L., Perez, Y.-A., Valadon, G. and Levillain, O. (2010), 'Can you still trust your network card?', CanSecWest 2010, Vancouver, Canada.
URL: *http://www.ssi.gouv.fr/IMG/pdf/csw-trustnetworkcard.pdf*

Jansen, W. and Karygiannis, T. (1999), Mobile agent security, Technical report, National Institute of Standards and Technology (NIST). NIST Special Publication 800-19.
URL: *http://csrc.nist.gov/publications/nistpubs/800-19/sp800-19.pdf*

Kagermann, H., Wahlster, W. and Helbig, J. (2013), Recommendations for implementing the strategic initiative INDUSTRIE 4.0, Technical report, ACATECH – German National Academy of Science and Engineering.
URL: *http://www.acatech.de/fileadmin/user_upload/Baumstruktur_nach_Website/Acatech/root/de/Material_fuer_Sonderseiten/Industrie_4.0/Final_report__Industrie_4.0_accessible.pdf*

Karnouskos, S. (2001), 'Security implications of implementing active network infrastructures using agent technology', *Computer Networks, Elsevier* **36**(1), 87–100.

Karnouskos, S. (2011), Stuxnet worm impact on industrial cyber-physical system security, *in* 'IECON 2011 - 37th Annual Conference on IEEE Industrial Electronics Society', pp. 4490–4494.
URL: *dx.doi.org/10.1109/IECON.2011.6120048*

Karnouskos, S. (2014), 'Security in the Era of Cyber-Physical Systems of Systems', *ERCIM News* (97), 44–45.
URL: *http://ercim-news.ercim.eu/en97/special/security-in-the-era-of-cyber-physical-systems-of-systems*

Leitão, P., Marík, V. and Vrba, P. (2013), 'Past, present, and future of industrial agent applications', *Industrial Informatics, IEEE Transactions on* **9**(4), 2360–2372.

Marrón, P. J., Minder, D. and Karnouskos, S. (2012), *The Emerging Domain of Cooperating Objects: Definition and Concepts*, Springer.
URL: *http://dx.doi.org/10.1007/978-3-642-28469-4*

Mařík, V., Vrba, P., Hall, K. H. and Maturana, F. P. (2005), Rockwell automation agents for manufacturing, *in* 'Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems', AAMAS '05, ACM, New York, NY, USA, pp. 107–113.
URL: *http://doi.acm.org/10.1145/1082473.1082812*

Mcdonald, J. T. (2006), Enhanced Security for Mobile Agent Systems, PhD thesis, Florida State University, Tallahassee, FL, USA. AAI3252145.
URL: *http://www.cs.fsu.edu/research/dissertations/JTM.pdf*

Mead, N. R., Hough, E. D. and Stehney, T. R. (2009), Security quality requirements engineering (SQUARE) methodology, Technical report, Carnegie Mellon. CMU/SEI-2005-TR-009, ESC-TR-2005-009.

Miller, C. (2011), Battery firmware hacking, *in* 'Black Hat USA+2011, Las Vegas, NV, USA'.
URL: *http://media.blackhat.com/bh-us-11/Miller/BH_US_11_Miller_Battery_Firmware_Public_WP.pdf*

Patrick, A. S. (2002), 'Building trustworthy software agents', *Internet Computing, IEEE* **6**(6), 46–53.

Pereira, A., Rodrigues, N., Barbosa, J. and Leitao, P. (2013), Trust and risk management towards resilient large-scale cyber-physical systems, *in* 'Industrial Electronics (ISIE), 2013 IEEE International Symposium on', pp. 1–6.

Porter, M. E. and Heppelmann, J. E. (2014), 'How smart, connected products are transforming competition', *Harvard Business Review* pp. 65–88.
URL: *http://hbr.org/2014/11/how-smart-connected-products-are-transforming-competition/ar/1*

Schild, K. and Bussmann, S. (2007), 'Self-organization in manufacturing operations', *Commun. ACM* **50**(12), 74–79.
URL: *http://doi.acm.org/10.1145/1323688.1323698*

Swanson, M. and Guttman, B. (1996), Generally accepted principles and practices for securing information technology systems, Technical report, National Institute of Standards and Technology Technology Administration (NIST).
URL: *http://csrc.nist.gov/publications/nistpubs/800-14/800-14.pdf*

Vollmer, T. and Manic, M. (2014), 'Cyber-physical system security with deceptive virtual hosts for industrial control networks', *Industrial Informatics, IEEE Transactions on* **10**(2), 1337–1347.

Zeng, W. and Chow, M.-Y. (2013), 'Modeling and optimizing the performance-security tradeoff on d-ncs using the coevolutionary paradigm', *Industrial Informatics, IEEE Transactions on* **9**(1), 394–402.