Assessing the Security of Internet Connected Critical Infrastructures (The CoMiFin Project Approach)

Hamza Ghani*, Abdelmajid Khelil*, Neeraj Suri*, György Csertán†, László Gönczy†, Gábor Urbanics† and James Clarke‡
*Technische Universität Darmstadt, Germany
Email: {ghani, khelil, suri}@cs.tu-darmstadt.de
†OptXware R&D Ltd., Hungary
Email: {csertan, gonczy, gabor.urbanics}@optxware.com
‡Waterford Institute of Technology, Ireland
Email:jclarke@tssg.org

***Abstract*—As the Internet of Things (IoT) pervasively extends to all facets of life, the "Things" are increasingly extending to include the interconnection of the Internet to Critical Infrastructures (CI) such as telecommunication, power grid, transportation, e-commerce systems, etc. The objective of this paper is twofold: (i) addressing IoT from a CI protection (CIP) and connectivity viewpoint, and (ii) highlighting the need for security quantification to improve the quality of protection (QoP) of CI's. Using a financial infrastructure as an example, a CIP and trust quantification perspective is built up in the EC CoMiFin project [5]. To this end, we are developing a novel security metrics-based approach to assess and thereon enhance the CIP. We focus on the communication level of the CI where IoT is playing an increasingly important role with respect to sensing and communication across CI elements. CI monitoring and notification get a special consideration in our approach. Determining the security and dependability level of the communication over the CI constitutes a basic precondition for assessing the QoP of the whole CI, which is needed for any efforts to improve this QoP. Furthermore, the parameters defining the required level of the QoP determined in terms of Service Level Agreements (SLA) need to be taken into consideration. Thus, monitoring and measuring quantitatively the dependability and security using appropriate metrics is essential for realizing the target-performance comparison of the QoP. As metrics play a central role for such quantification, this paper develops their QoP usage from an IoT perspective.**

I. Introduction

The technological growth of "connected" systems is resulting in multitudes of sensing, monitoring and communicating devices as part of the environment. The resultant IoT extends from the home and work environments to complex e-commerce, utilities, medical and transportation sectors. The IoT potential is reflected in the new trends opening novel application fields especially in the area of security monitoring, surveillance and, in a more generic way, sensing every measurable parameter/aspect/event (physical or not), which entails relevant importance to the considered application. Nevertheless, the new IoT trends and the interconnectedness of things provide new security challenges. New security threats can appear with the increasing IoT connectivity: as the new sensing entities in critical environments represent attractive targets for attackers. Examples are sensors on ATMs, which can lead, in case of a compromise, to potential financial damage. On the other hand, IoT opens new horizons for security monitoring: advanced recognition of malicious behavior and attacks (e.g., through tracking mobile entities in a malicious exploit of mobile banking vulnerabilities). Such sensors and software agents are utilized in several CI's for monitoring and controlling purposes. Modern societies entail an increasing reliance on CI's such as the power grids, communication, financial and transportation networks. These CI's are increasingly interconnected and therefore highly dependent on computing and communication infrastructures. Logically, the need exists to secure CI's against threats (operational or deliberate) arising from the whole spectrum of interconnected entities as their disturbance can cause considerable material, financial and even (in extreme cases) human loss. As a consequence, the accurate and quantitative assessment of the security level of these CI's constitutes a key objective for both the public and private sectors. The need to protect CI's reveals the necessity to quantify trustworthiness (i.e. dependability and security) metrics to determine the exact trustworthiness level [17]. A first step towards CI protection (CIP) is accomplished through CI monitoring and control, provided by a monitoring infrastructure using multiple sensing nodes. A comprehensive trustworthiness evaluation of IoT-based CIP mechanisms is of equal importance as the expected CIP enhancement, in order to weigh off the benefits and risks of IoT-based CIP. As CI's interconnect, the communication networks connecting them garner increasing attention for their resilience properties. Intrusion detection systems (IDS) [13], [16], distributed firewalls [12], spam detection [4], etc. all utilize a responsive and resilient communication framework. Amongst the varied communication techniques, the use of Peer-to-Peer (P2P) overlays [14] is increasingly being pro-

Research supported in part by EC FP7 Strep CoMiFin.

posed to protect CI's [12], [4] because of their inherent properties such as robustness and redundancy. This makes the protection overlay a critical component requiring careful design with guaranteed and measurable properties. The vision of adopting a cooperative approach for protection is much more beneficial for the CI's. As each CI component needs to be secured, synergy effects can be generated and a higher collaborative level of the security relevant information quality can be achieved, which can be beneficial for all participating CI components. A P2P protection layer is similarly based on the concept of cooperation between participants in the overlay network that is dedicated to improve the CI security with respect to the detection of attacks and potential cyber threats. Usually, a dedicated middleware for collecting, processing and disseminating security relevant data is used. The monitoring of this middleware, being a CI itself, is crucial for maintaining the required service level for the participating entities. Defining several dedicated metrics, whose calculation needs partly real time measurement data from the sensing nodes, constitutes the basis for monitoring the security of the CI. Though sensing nodes and overlays are used in multiple scenarios to realize critical protection mechanism, there is a dearth of approaches to quantitatively evaluate the protection enhancement. Quantitative measures are important for designers to maximize the resilience of protection mechanisms, for users to increase their trust in the system, and for managers to assess their investment [1].

Problem Statement: In this paper, we primarily investigate the fundamental question of how to evaluate and assess the trustworthiness of the CIP mechanism deploying sensing nodes and communication overlays. To quantitatively measure the level of QoP, appropriate, application-dependent metrics are needed. Furthermore, metrics need to be defined for both the design phase (Trustworthiness by Design) as well as for run-time (Trustworthiness by Repair). For the latter case, we investigate methods to calculate the appropriate metrics at run-time based on the measurement data gathered by the sensing nodes.

Paper Contributions: For IoT connected critical infrastructures, this paper develops a generic metric-based approach to evaluate QoP at both design and run-time. Specifically, we present:

- an overview of existing QoP metrics taxonomies
- an approach for a metric-based definition of SLAs
- automated generation of the monitoring configuration from the metric and SLA definitions
- multi-level metric evaluation system to handle complexity (plug-in concept) utilizing (i) simple arithmetic evaluators (ii) simple rule based evaluator (iii) complex event processing based evaluator
- a reference implementation for trustworthiness by repair based on run-time Metrics Monitoring

We illustrate our approach using P2P-based protection approaches targeting a Financial Infrastructure (FI) [5]. However, our developed approach is generic and valid for other CI's and can be used for different P2P protection solutions.

Paper Organization: Section 2 discusses architecture models for P2P-based CIP along with the system/overlay architecture considered throughout this paper. Furthermore, we present a case study describing our approach for the critical FI domain while emphasizing the generic character of our approach and its applicability in other different scenarios. The focus of Section 3 is on security metrics and their taxonomies. In Section 4, the role of SLAs in assuring security in the design phase is highlighted. Furthermore, we develop our metrics and SLA-based monitoring framework and present initial implementation assessments. The related work is discussed in Section 5.

II. Architecture and System Model

After discussing the possible architecture models for using P2P for CIP, we present a case study and the corresponding system model that we will follow throughout the paper.

A. Architecture Models

For IoT-based CIP, we distinguish between two fundamental approaches: (a) intrusive or (b) non-intrusive solutions. For intrusive approaches, the protection mechanisms (e.g., IoT-based) are embedded in the CI. Intrusive approaches are not always possible given the proprietary nature of the CI. For the critical FI, no access to the existing CI can be provided rendering intrusive approaches unsuitable. Figure 1 illustrates an example architecture for intrusive protection approaches. In [9], such approach is proposed to protect distributed control systems.

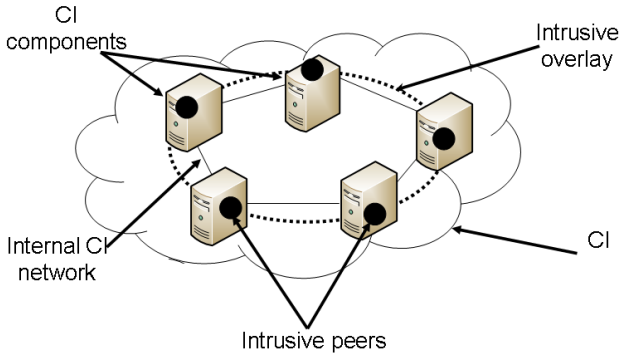

Fig. 1. CI with intrusive protection mechanism

Non-intrusive approaches consist of deploying an additional secure/dependable P2P overlay that is decoupled from the CI. The drivers of deploying a supplemental overlay are twofold (a) meet the specific requirements of non-intrusiveness of the underlying CI and (b) avoid introducing new vulnerabilities (e.g., compatibility, new threats etc) specific to the intrusive approach. Figure 2 illustrates a non-intrusive overlay network on top of a CI.

In Figure 2, the (financial) CI is handled as a black-box. For a meaningful protection, the CI should only minimally interact with the overlay network while having full control on the

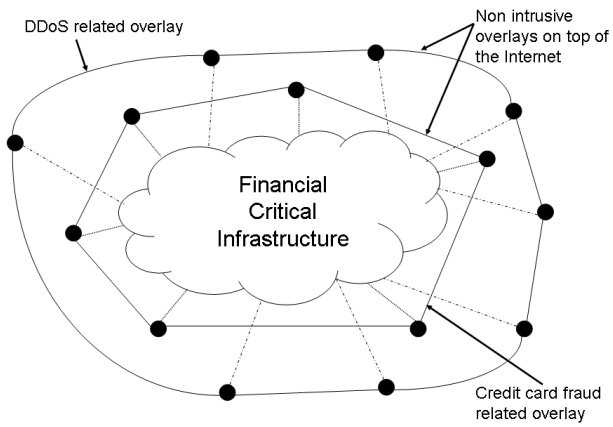

Fig. 2. CI with non-intrusive protection mechanism

exported or imported data. Non-intrusive approaches realize an additional defense line/layer that implements further/new (usually collaborative) security mechanisms in addition to the original security measures taken inside the components/members of the CI.

B. P2P-based Protection of Financial Infrastructure

In this section, we present a case study illustrating the validity of utilizing the properties of P2P overlays to enhance the protection of CI. Financial institutions such as banks, stock markets, insurance companies and rating agencies are tightly interconnected and operate thousands of transactions every day. Most of these transactions are operated electronically. Therefore the FI is highly dependent on the underlying (partly IoT-based) infrastructure. Nevertheless, securing the IoT remains a highly sensitive issue which is tackled by each financial institution on its own. However, a cooperative defense has been proven to be more successful than isolated solutions. Accordingly, the trend is to perceive IT-security as a cooperation issue between FIs and not as a competition field. One fundamental property of the FI is the high rank that privacy has. It constitutes a fundamental requirement that needs to be considered in each IT solution, especially those in relation with the sensitive field of IT security. P2P with its inherent properties can play a central role protecting the FI. This scenario highlights the need of FIs for a secure/dependable (large-scale) monitoring infrastructure to share security relevant information. Accordingly, we consider FIs and propose an approach to facilitate the cooperation between financial institutions to mitigate distributed attacks and disseminate local knowledge about the QoP level of FIs. The EU project CoMiFin [5] is aiming at realizing a better protection of the FIs through the deployment of a P2P overlay following the approach described in the generic architecture of Figure 2. One of the core requirements that the financial institutions willing to collaborate in a shared security mechanism consists in revealing no information about their internal security related processes. Given the strong privacy and non-intrusiveness requirements of financial institutions, only a non-intrusive approach can be considered. Deploying distributed cooperative security mechanisms on the overlay network constitutes a novel approach permitting to benefit from the advantages of a collaborative defense work of different independent institutions. Sharing security relevant information such as attack alerts or potential cyber threats among the financial institutions is beneficial for all participating parties.

III. Trustworthiness Metrics

A widely accepted principle is that an activity is harder to enhance if it cannot be measured. Measuring and controlling IT security through metrics is a less explored research field. Such metrics are a prerequisite for understanding, improving and validating/certifying the security of CI's. As it is almost impossible to survey all the existing security metrics in a section of this paper, we present an insight in the different existing classes of security metrics. First, we present the different classifications of the existing security metrics as a contribution to structure the young research field of security metrics. Then, some of the most representative security metrics will be discussed.

A. Taxonomies of Security Metrics

In the literature there exist some different approaches aiming at categorizing the existing security metrics. Some of these taxonomies have been developed for practitioners. Therefore they do not cover the whole spectrum of existing security metrics as they are industry oriented and try to fulfill the requirements of the market [22], [21]. Others [19] present a high level taxonomy containing metrics for both organizational information security management and product development. The National Institute of Standards and Technology (NIST) [18] has proposed a different security metrics taxonomy from the perspective of an organization. It contains three distinct categories: management, technical, and operational metrics. Each of these categories contains several sub categories (17 in total). The Institute for Information Infrastructure Protection (I3P) [10] has also developed a security metrics taxonomy [21]. Starting from the process control system perspective, it consists of three different metrics classes: technical, operational, and organizational metrics activities. Each of them contains several subcategories. Vaughn et al. [23] have proposed a taxonomy consisting of two categories: organizational security metrics and metrics for technical target of assessment. Seddigh et al. [20] introduce a taxonomy for IT Networks. It consists of three categories: security, quality of service (QoS), and availability. Each of these three categories consists of technical, organizational and operational metrics. The technical metrics consist of subcategories for product rating, incident statistics and security testing. The organizational security metrics include metrics for information assurance program development and resources. And the operational security metrics include metrics for technical readiness, susceptibility and effectiveness. Savola [19] introduces a high-level information security metrics taxonomy incorporating metrics for both organizational information security management and product

development. It begins with the Level 0: security metrics for business management which is the highest category in this taxonomy. It contains five subcategories: (a) Security metrics for cost-benefit analysis containing economic measures such as Return on Investment (ROI) (b) Trust metrics for business collaboration (c) Security metrics for business-level risk analysis (d) Security metrics for information security management (ISM) (e) Security, dependability and trust (SDT) metrics for ICT products, systems and services. The Savola's taxonomy is covering the new emerging field of economic driven security metrics, which targets an audience of managers and decision makers without IT security background. The taxonomies we mentioned here constitute a considerable effort bringing order to the myriad of existing security metrics. They consider the existing metrics from different perspectives: organizational management, product development, operational assessment, etc.

B. Illustration of Calculated Metrics within CoMiFin

Security metrics are usually application dependent. Determining the appropriate set of metrics requires a good understanding of the given domain. In order to define the metrics for the FI scenario in CoMiFin we applied the Goal-Question-Metric (GQM) [3] approach, which is a user-centric, widely accepted metrics definition methodology. As a result we identified the following categories of metrics: (a) Resource-level metrics: This category includes elementary resource usage metrics, such as CPU, memory, disk or network usage. The Metrics Monitoring framework is able to correlate these "low-level" metrics with other, "high-level" metrics in order to detect important but otherwise hidden patterns of misbehavior (b) Availability metrics: The group describes metrics that allow for measuring classical availability attributes in the overlay. For example, mean uptime, availability, reliability and mean repair time of each component. These are the most typical metrics to be included into SLAs (c) Communication metrics: Since the information sharing in the overlay itself has tough security requirements (confidentiality, non-repudiation), the attributes of the applied communication mechanisms are essential metrics. These are for example, the strength of the applied encryption, the ratio of encrypted/signed content, or the time required to transfer messages or the latency of the communication system (d) Application specific metrics: This group describes the applications running over the overlay, e.g. the version of the application running, the number of available but not installed security updates for that particular version (e) Overlay specific metrics: This group of metrics includes metrics that describe the attributes of the overlay as a whole. For example, K-connectivity or proximity properties of the overlay are included (f) Trust metrics: This category includes trust level measurements for CoMiFin participants [2].

IV. Metric-based QoP Assessment

The core idea of our approach consists in (a) a metric-based definition of SLAs, and (b) run-time metric monitoring. We mainly follow a stepwise approach:

1. Define application dependent security requirements for the overlay
2. Define a set of metrics in order to monitor the fulfillment of the defined requirements
3. Based on the defined metrics, determine clear and unambiguous SLAs which fulfillment can be monitored at run-time by Metrics Monitoring (MeMo)
4. IoT-based run-time monitoring of the degree of compliance with the defined security related SLAs
5. Any SLA violations can be detected so that appropriate decisions can be taken according to the penalties defined by the SLA

Steps 1 and 2 target protection by design. Steps 2-5 implement an IoT-based approach for protection by repair utilizing quantifiable metrics. Figure 3 illustrates the interaction between the two core components of our approach, namely metrics monitoring (MeMo) and SLA management. In the following

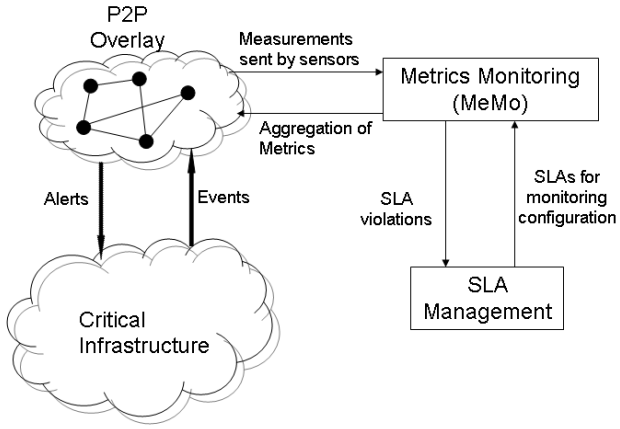

Fig. 3. SLA Management/Metrics Monitoring

we describe the two core components of our approach depicted in Figure 3 presenting their role in enhancing the CIP at design and at run-time.

A. Trustworthiness by Design

In this section, we focus on using metrics in order to define the QoP at design time. Our approach consists in defining metric-based SLA in order to capture the user requirements, to define the guarantees the system is required to provide, and the penalties in case of not reaching the specified guarantees. It constitutes a basis for the process of eliciting meaningful metrics from the stakeholders. At the design phase, the QoP of the CI needs to be clearly and unambiguously determined. In order to deal with this issue, our approach envisages utilizing SLAs as an essential means to contractually define the minimum service level that participants have to guarantee for a determined time period (e.g., minimum of computational resources, k-connectivity, etc.). The SLA Manager (Figure 4) is responsible for creating and managing SLAs. Our approach envisages the implementation of an SLA Manager responsible for managing SLAs throughout their whole lifecycle. The SLA Manager will provide communication interfaces to be used

by the Metrics Monitoring (MeMo) component described in Section IV-B. As a proof of concept, we have implemented an SLA Manager that is interacting with the MeMo in order to monitor the degree of compliance with the predefined SLAs. Especially, the security related part of the SLA is playing a central role in this respect. The interaction with MeMo can be triggered in case that an SLA has been violated. IoT-based monitoring activities of MeMo allow the detection of SLA violations, which can be directly reported to SLA Manager. This notification is accompanied by a suggestion of appropriate countermeasures that can be taken according to the penalties stated in the SLA (e.g., forcing some members to leave CoMiFin). In the following, we describe the design specifications and the prototype implementation of the SLA Manager. The internal architecture of the SLA Manager is shown in Figure 4. Two main subcomponents are responsible for providing the functionalities of SLA Manager. The SLA Lifecycle is responsible for managing SLAs throughout their lifecycle (from creation to disbanding). Furthermore, the Alarm Engine assures the compliance with the SLAs through its interaction with MeMo, which is monitoring the degree of compliance of each participating institution with the signed SLAs. Finally, SLAs are stored in a dedicated SLA repository. Figure 5 presents a component diagram showing the interfaces between the SLA Manager and MeMo. The current prototype implementation of the SLA Manager contains a part of the final functionalities. It focuses on the core functionalities of creating, modifying and deleting SLAs from the repository. The current implementation includes: (i) A web application permitting to manage SLAs throughout their lifecycle (from creation to expiration), (ii) a database component where the SLA Manager can store the SLAs, and (iii) a set of web services assuring the interaction with MeMo. Two categories of web services can be differentiated: a) Web services provided by the SLA Manager: As shown in Figure 5, the web service interfaces SLAViolationNotification and SLAWarningNotification can be called by MeMo in case that the latter has detected an SLA violation for the first interface and in case that MeMo has pro-actively detected that an SLA is close to be violated for the second interface. b) Web services required by the SLA Manager: MeMo provides, as depicted in Figure 5, three web service interfaces that the SLA Manager will call to notify MeMo about any SLA status modification (SLA creation, modification and deletion).

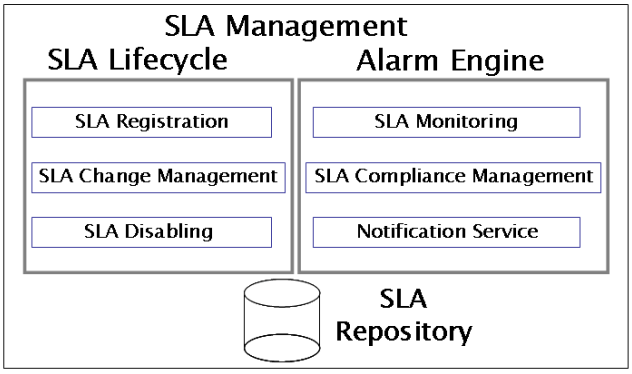

Fig. 4. Building blocks of SLA Manager

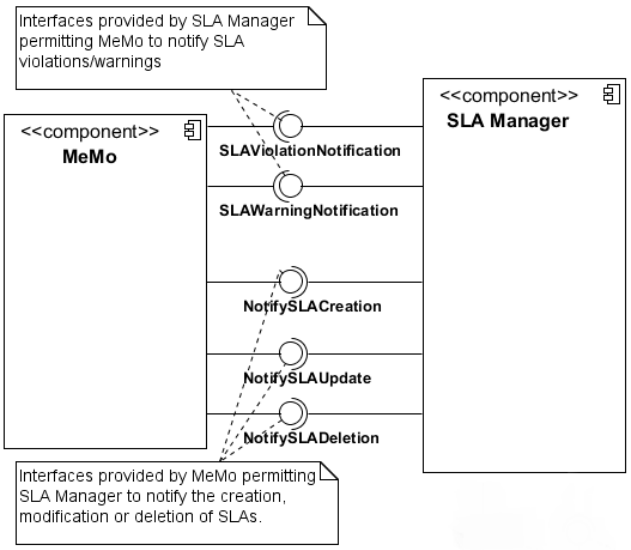

Fig. 5. SLA Manager/MeMo

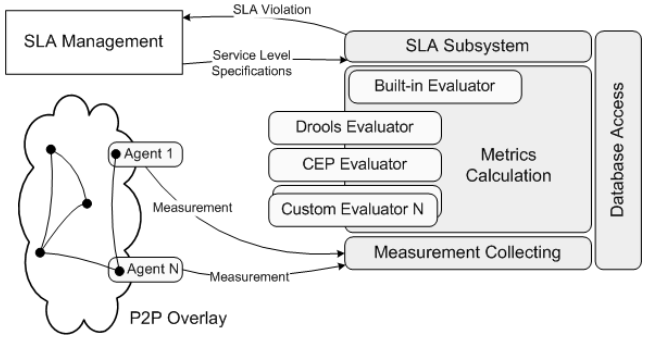

Fig. 6. Reference architecture for metrics evaluation

B. Trustworthiness By-Repair

Now we focus on using metrics to evaluate QoP at run-time and accordingly trigger alerts or overlay reconfigurations to maintain the desired QoP level.

1) Metrics Monitoring (MeMo): Metrics Monitoring is responsible for (a) collecting measurements at run-time, (b) calculating the metrics based on these measurements, and (c) emitting notifications when the violations of Service Level Specifications occur. Figure 6 depicts a layered architecture of the MeMo framework. The first layer, which is labeled ”Measurement Collecting”, receives the measurements from the nodes in the overlay through P2P protocols such as routing, multicast, etc. Each measurement is determined by the metrics definition and is implemented by so called agents deployed to the nodes of the overlay. Agents in this sense are small, simple applications capable of measuring the necessary attribute locally on the node, e.g., a CPU usage agent measures the

utilization of the CPU of the node. Here we emphasize that IoT already provides a flexible and (due to its distributed nature) resistent measurement infrastructure which can contribute to improve the quality and timeliness of measurements. The use of embedded devices as measurement sensors naturally raises security and privacy concerns which are out of the focus of the current paper but are future research activities.

The next layer, which is called "Metrics Calculation", calculates the defined metrics based on the collected measurements. There are two ways how the metrics can be calculated in the MeMo framework: (a) Either built-in mechanisms are used, allowing for simple and relatively fast arithmetic calculations, e.g. totaling, averaging, minimum or maximum calculation or (b) evaluator plug-ins are utilized if advanced or complex functionality is required (e.g., the collected measurements are basis for a complex metric or the amount of data is large). Evaluator plug-ins has access to the collected measurements and the previously calculated metrics. In the current implementation, a *rule engine* based plug-in is used for correlating the received measurements to effectively process metrics and detect meaningful patterns by performing complex event processing. Also a *trust evaluation* plug-in is integrated to score measurement on the basis of trust level of participants (see [2] for details).

"SLA Subsystem", the topmost layer of the architecture, is responsible for collaborating with the SLA Manager. Internally, it is responsible for processing the SLAs (either they are predefined or coming from the SLA Manager) and tracking the adherence/violation of these SLAs. In case of SLA violation, this component notifies the SLA Management system. All three layers use the Database Access subcomponent to persistently store measurements, metrics, SLA descriptions and notifications.

2) MeMo Preliminary Implementation: The preliminary implementation of MeMo integrates several open source components. The Measurements Collecting component is implemented by the Nagios citeNagios monitoring server along with its agent framework to collect specific measurements. The component has been designed in a way that supports the easy replacement of the underlying monitoring system in later development phases against industry standard monitoring tools. The built-in Metrics Calculation functionality is provided by database triggers of MySQL database which provides only simple functionality regarding metrics calculation. On the other hand, calculating with triggers is effective. As mentioned before, a rule engine based evaluator plug-in has been developed which utilizes the Drools [7] rule engine. Communication between the evaluator plug-ins and the Metrics Calculation subcomponent is implemented by Java RMI or through Web Services using the JBoss web service stack.

3) Automated Generation of Monitoring Configurations: As IoT infrastructures are subject to frequent changes, a consistent and coherent model of the system needs to be maintained which is then the basis of IT system monitoring and management. Currently, we use Eclipse Modeling Framework and related technologies to automatically generate Nagios configuration settings. Our current implementation still needs manual intervention yet it is i)under development to help runtime validation of configuration and synthetize configurations on the basis of high level requirements and ii)logically independent from the concrete platform, as stated before.

V. Related Work

P2P overlays are increasingly used as a part of the protection strategy of systems and infrastructures. In the following we will present a short review of the existing P2P protection techniques. The secure overlay services (SOS) [12] approach aims at preventing denial of service (DoS) attacks through the usage of a secure overlay tunneling. It allows communication only between a confirmed user and a target through the authorization mechanism of the SOS overlay. A basically similar approach dedicated for web server protection is called WebSoS [6]. It utilizes overlay networks in order to allow authenticated users to access web servers even if they are under a congestion-based DDoS attack. In this approach, the overlay network immediately surrounding the web servers to be protected filters and blocks all incoming packets from hosts that are not legitimate. P2P architectures are also utilized for collaborative intrusion and malware detection [16]. The proposed approach is based on a decentralized, P2P design that addresses dependability and load unbalance issues affecting existing systems based on centralized and hierarchical schemes. Other proposals based on P2P defensive schemes (e.g., [15], [8]) differ from this paper because their focus is on novel algorithms for anomaly detection that should be facilitated by cooperation. Other P2P schemes (e.g., [25], [24], [11]) are used to disseminate information about malicious IP addresses through some publish/subscribe model. SCADA systems protection can be enhanced through the usage of P2P [9]. P2P techniques allow the design of self-organizing Internet-scale communication overlay networks. The approach presented in [9] emphasizes two inherent resilience mechanisms of P2P networks which are path redundancy and data replication. It shows how SCADA system's resilience can be improved by using P2P technologies.

VI. Conclusions and Future Work

In this paper, we presented an approach to assess the security level of internet connected CI's. This approach is IoT-centric and is based on metrics covering both design and runtime phase. CI's are currently utilizing own internal monitoring systems. Our metrics-based approach can be used to assess the level of the existing protection mechanisms. To this end, SLAs are used to help determine a level of protection to be reached. Furthermore, a collaborative approach for protection is more beneficial for CI's than individual protection mechanisms merely. Considering the financial CI, which is characterized by tough privacy issues, a non-intrusive, external monitoring overlay is proposed. The contribution of this external system to the level of protection has to be measured and validated utilizing our IoT-centric, metric-based approach. Although the

used P2P overlay network is non-intrusive, it still contains and provides security relevant information for the CI, on which the CI relies. The quality of its robustness and protection mechanisms should also be measured and validated (against SLAs). Our IoT-based approach plays a central role in this regard. The novelty of our approach can be summarized in the following: (i) we present metric based definition of SLAs (ii) semi-automatic generation of the monitoring configuration out of the metric and SLA definitions (iii) a multi-level metric evaluation system to handle complexity (the plug-in concept described in Section IV-B). We believe that future work needs to be done to develop formal models of security measurement and Metrics. Furthermore, as the protection of privacy is a crucial aspect, appropriate metrics need to be defined.

REFERENCES

[1] A. Acquisti. Essays on privacy, anonymity, and tracking in computer-mediated economic transactions. PhD thesis, UC Berkeley, 2003.

[2] R. Baldoni et al. Trust management in monitoring financial critical information infrastructures. In *MOBILIGHT*, 2010 (to appear).

[3] Basili et al. The goal question metric approach. in Encyclopedia of Software Engineering, pp. 528-532, Wiley, 1994.

[4] A. Brodsky et al. A distributed content independent method for spam detection. In *Proc. HotBots'07*, p. 3., 2007.

[5] http://www.comifin.eu, 2010.

[6] D. Cook et al. WebSoS: protecting web servers from ddos attacks. In *Proc. of ICON'03*, pp. 461-466, 2003

[7] http://www.jboss.org/drools, 2010.

[8] C. Dumitrescu. A peer-to-peer approach for intrusion detection. in *Proc. of CCGRID'06*, vol. 1, pp. 89-92, 2006.

[9] D. Germanus et al. Increasing the resilience of critical scada systems using peer-to-peer overlays. In *Intl. Symposium on Architecting Critical Systems*, LNCS 6150, pp. 161-178, 2010.

[10] Institute for Information Infrastructure Protection http://www.thei3p.org/, 2010.

[11] R. Janakiraman et al. Indra: a peer-to-peer approach to network intrusion detection and prevention. In *Proc. WET ICE'03*, pp. 226–231, 2003.

[12] A. Keromytis et al. SoS: An architecture for mitigating ddos attacks. *IEEE Journal on Selected Areas of Commn.*, Vol. 22, 176-188, 2004.

[13] M. Locasto et al. Towards collaborative security and P2P intrusion detection. In *Proc. IEEE Information Assurance Workshop*, pp. 30-36, 2005.

[14] E. Lua et al. A survey and comparison of peer-to-peer overlay network schemes. *IEEE Commn., Surveys, Tutorials*, Vol. 7, pp. 72-93, 2005.

[15] D. J. Malan and M. D. Smith. Host-based detection of worms through peer-to-peer cooperation. In *Proc. of WORM'05*, pp. 72-80, 2005.

[16] M. Marchetti et al. P2P architecture for collaborative intrusion and malware detection on a large scale. In *Proc. Intl. Conf. on Information Security*, pp. 475-490, 2009.

[17] S. Naqvi and M. Riguidel. Quantifiable security metrics for large scale heterogeneous systems. In *Proc. IEEE Conferences on Security Technology*, pp. 209-215, 2006.

[18] National Institute of Standards and Technology (NIST) http://www.nist.gov/, 2010.

[19] R. Savola. A novel security metrics taxonomy for r&d organizations. In *Proc. of ISSA'08*, pp. 379-390, 2008.

[20] N. Seddigh et al. Current trends and advances in information assurance metrics. In *Proc. of PST'04*, pp. 197-205, 2004.

[21] M. Stoddard et al. Process control system security metrics, state of practice. Technical report, Institute for Information Infrastructure Protection Research, 2005.

[22] M. Swanson et al. Security metrics guide for information technology systems. NIST report 800-55, 2003.

[23] R. Vaughn et al. Information assurance measures and metrics : State of practice and proposed taxonomy. In *Proc. of HICSS'03*, pp. 331-340, 2003.

[24] V. Yegneswaran et al. Global intrusion detection in the domino overlay system. In *Proc. of NDSS'04*, pp. 1-17, 2004.

[25] C. Zhou et al. A peer-to-peer collaborative intrusion detection system. In *Proc. of Intl. Conf. on Networks*, volume 1, pp. 118-123, 2005.