A Survey on Research Challenges in Privacy and Security for Cloud-Supported Internet of Things

Bharathi Shetty[a] and Umesh Chavan[b]

[a]*Assistant Professor, Department of Information Technology, Walchand College of Engineering, Vishrambag, Sangli, Maharashtra, India. Email: bharati.shetty@walchandsangli.ac.in*
[b]*Assistant Professor, Department of Information Technology, Walchand College of Engineering, Vishrambag, Sangli, Maharashtra, India. Email: umesh.chavan@walchandsangli.ac.in*

Abstract: There has been a growing interest in the ability of embedded devices, sensors, and actuators to Communicate, and create a ubiquitous cyber-physical world. The growth of the notion of the Internet of Things (IoT) and the rapid development of technologies such as short range mobile communication and improved energy-efficiency is expected to create a pervasive connection of "things." This will inevitably result in the generation of enormous amount of data, which have to be stored, processed, and accessed. Cloud computing has long been recognized as a paradigm for big data storage and analytics. The combination of cloud computing and IoT can enable ubiquitous sensing services and powerful processing of sensing data streams beyond the capability of individual things, thus stimulating innovations in both fields. There is an urgent need for novel network architectures that seamlessly integrate them, and protocols that facilitate big data streaming from IoT to the cloud. The data security, privacy, and reliability, are critical concerns during the integration. Our contribution is to analyze the current state of cloud-supported IoT to make explicit the security considerations that require further research work.

Keywords: Internet of Things (IoT), cloud, law, privacy, security.

1. INTRODUCTION

Duringthe last decades of the Twentieth Century, there was much research into sensor and communications technologies. At that time, sensor-based systems tended to be developed in "silos," being localized, application- and technology-specific. It became evident that sensor data could potentially be used for many diverse purposes if a means of sharing could be devised. The term "Internet of Things" (IoT), first coined in 1999 by Ashton at MIT, 1 came to be used to capture this aspiration: (1) based on ever-wider connectivity of sensor/actuator-based systems, more general data sharing would become possible than within the specific applications for which those systems were developed and (2) computers would become autonomous, able to collect data and take decisions based on them, without human intervention. Moreover, IoT represents a broader move to the vision of pervasive or ubiquitous computing .

Recently, the IoT concept has captured imaginations within government and commerce, as a technology capable of supporting immense growth [1], [2]. However, systems aiming at this wider vision are in their infancy. Sensor/actuator-based systems have been developed independently of the IoT vision of open data sharing. It is crucial that the security, privacy, and personal safety risks arising from open access to data, across and beyond these systems, are evaluated and addressed. IoT potentially covers a wide range of applications, including smart

home systems, smart street lighting, traffic congestion detection and control, noise monitoring, city-wide waste management, real-time vehicle networks, and smart city frameworks [3]. At the individual level, personal health and lifestyle monitoring systems are being integrated with general health care services [4]. Such application scenarios tend to be sensor/actuator-based, each developed for a single purpose. In contrast, the IoT philosophy is the wide-scale integration of potentially all technology, including individual devices, applications, servers, and so forth, in addition to sensors/actuators, i.e., the data from a range of different sources is capable of diverse potential application and should be developed with broad usage and wide availability in mind.

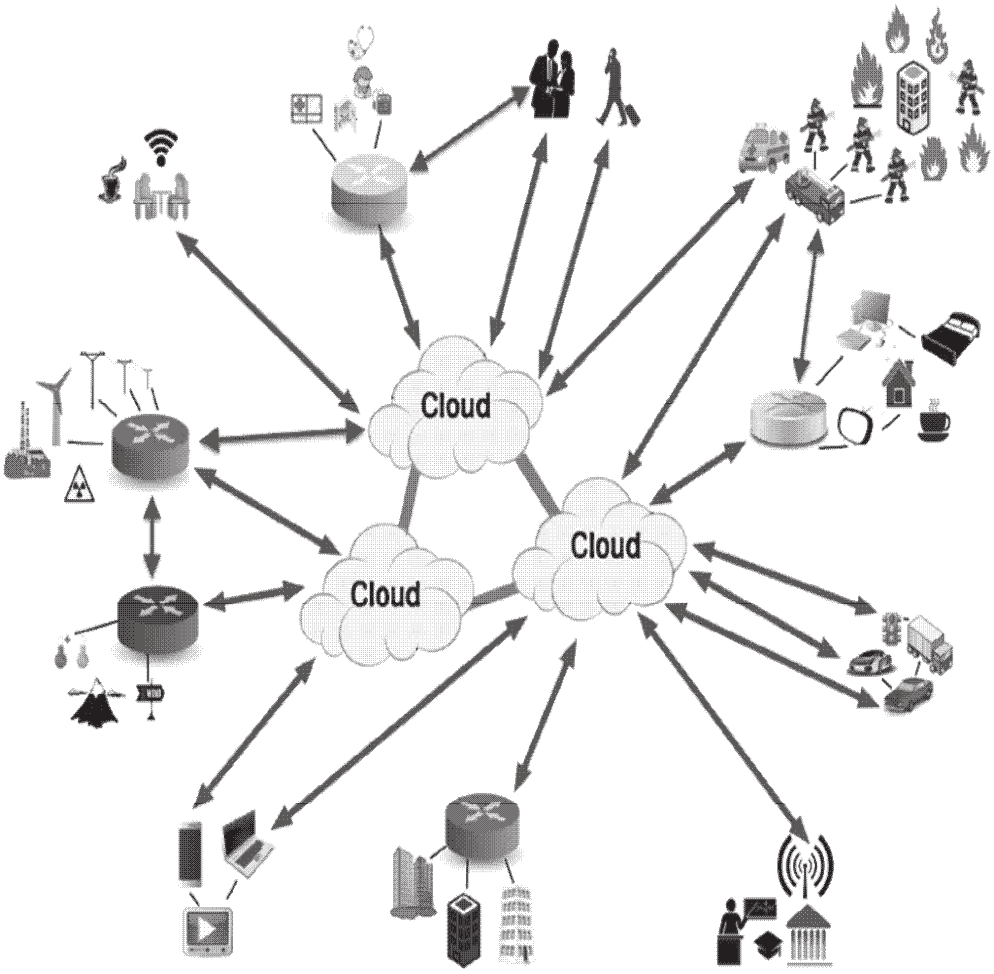

Figure 1: The interactions within an IoT-cloud

The cloud is an obvious technology for achieving this open sharing. Cloud computing has evolved to manage, process, and store *big data*, that, e.g., has arisen from services such as search engines. Data analytics has become an essential complement to cloud-hosted web services. Similar services can be used for large-scale data from IoT systems (including those that are mobile), making them independently shareable and widely available. The cloud is an ideal component in an IoT architecture. First, because cloud services can operate across a range of systems, services, and devices, it provides the natural point for: (1) data aggregation and analysis and (2) the management, control, and coordination of the range of systems and services. Furthermore, (3) cloud services offer benefits in terms of resource management, as clouds are always ON, can scale to meet demand, and can allow the off loading from constrained hardware of data (for computation [5] and storage) and management specifics. In this paper, we use *IoT-cloud* to refer to IoT architecture that incorporates cloud services.

The support for connectivity and open sharing via cloud services allows, e.g., emergency services to interact with traffic control, power (utility) providers, we present *some key security considerations* for IoT-cloud. Each section operates to encapsulate a number of considerations.

We first consider issues accessing the cloud (Section II), exploring issues of secure transport (1) and cloud access controls (2), before considering the range of data management concerns in Section III, We then discuss issues of identity management. Dealing with malicious "things" and associated attacks is explored in Section IV, The focus then turns to the integrity of cloud services, considering certification and trustworthiness Section V.

2. ACCESSING THE CLOUD

A. Secure Communications

There are two motivations for securing communication: (1) secrecy: preventing eavesdropping and data leakage and (2) integrity: protecting data from corruption/interference. Note that here we do not consider communication within subsystems, but rather are concerned with the interaction of "things" with cloud services. Secure communication is required to prevent unauthorized access to data (or metadata) that might be sensitive. Transport layer security (TLS) [6] uses cryptography to establish a secure channel to protect transmissions (including metadata such as protocol state, thus limiting side-channels) from both eavesdropping and interference. TLS employs a certificate based model, relying on public key infrastructure (PKI) and certificate authorities for authentication. TLS is a common feature of cloud-provider offerings, and can be used to secure the confidentiality and integrity of communications between "things" and the cloud provider. With a general view to making secure communication more commonplace, there is recent work on enabling TLS over protocol stacks other than TCP/IP to better suit the requirements of "things," in terms of complexity and resource requirements. Examples include datagram transport layer security (DTLS [7], [8]) for datagram oriented protocols such as User Datagram Protocol (UDP), and LLCPS [9] that applies TLS over the near-field communications Logical Link Control Protocol (LLCP). Depending on the deployment, architecture, and interfaces to cloud services, these technologies could facilitate new forms of secure "thing"–cloud interactions. Apart from TLS, there are, of course, other mechanisms of securing "thing"–cloud communication. Data can be encrypted by applications, which protect data not only in transit butalso beyond. Sharing secrets naturally entails management and engineering considerations [10].

Aside from any vulnerability inherent in the approach, the protection offered by any secure communication mechanism is only as good as its implementation. For example, the recent Heart bleed vulnerability in the widely used OpenSSL library is estimated to have left 24%–55% of TLS protected endpoints open to attack [11]. Extra care and consideration must be given to the newer schemes and implementations currently being developed to support IoT, especially those that may not havebeen widely scrutinized or deployed.

B. Access Controls for IoT-Cloud

It is important that (external) access to cloud resources is regulated. *Access controls* [10] operate to govern the actions that may be taken on objects, be they accessing particular data (a file, record, and data stream), issuing a query, performing some computation, and so forth. Controls are typically *principals* focused, in the sense that control policy governing a particular action is defined to regulate those undertaking the action, enforced when they attempt to take that action. There are two aspects to access control: (1) *authentication* and (2) *authorization*. Authentication refers to verifying who a principal is, i.e., are they who they say they are? Authorization rules follow authentication; once a principal is identified, what are their rights and privileges; what actions are they authorized to undertake?

In a general cloud context, the provider will offer access controls to ensure that only the correct tenants/

users (the *principals*) access the appropriate data and services. Cloud providers often have login/credential-based services for authenticating tenants/users. Authorization policy will be enforced as a principal attempts to take an action, based on their level of privilege, which might allow them to access storage and files held by the provider, initiate computation services, etc. The precise controls will depend on the specific service offering, but often include *access control lists*, *role-based access controls*, *capabilities*, etc. See [10] for an overview of a number of security engineering techniques.

In an IoT context, a challenge for any access control regime is accounting for the fact that the interactions between "things" may involve encounters with "things" never before seen, or owned and operated by others. Toward this, *trusted platform modules* (TPM) [12] offer promise by providing strong guarantees, e.g., with respect to device identity [13] and configuration [14], which access control mechanisms can leverage.

IoT-cloud poses extra challenges. The first concerns authentication, given the size and scale of the IoT vision, correctly identifying the "things" and determining the relevant cloud services/tenant applications is a real concern; Section V is dedicated to issues of identity. There are also difficulties in the fact that infrastructure and data may be shared. Currently, cloud policy is focused: authorization rules are to ensure that a tenant accesses only its own resources, i.e., their files, VMs, databases, etc. However, for the IoT-cloud, the lines are blurred. The data and resources of a tenant may be relevant to a number of different principals, and/or may control and coordinate a number of "things." Policy must be able to be consistently defined and applied across both of these dimensions.

Access controls may be contextual, e.g., people may in general only access data concerning themselves. In exceptional circumstances, such as medical emergencies [15], wider access may be desirable, as specified by "break-glass policies." Mechanisms are required to enable flexible access control policies to be defined by different parties, while also being able to identify and resolve potential policy conflicts. Such concerns are nontrivial, and will likely require some external constraints, such as ownership or economic incentives (e.g., those paying for the service) to help make access control policy more manageable.

3. DATA MANAGEMENT IN THE CLOUD

A. Identifying Sensitive Data

In an IoT context, it will often be the case that data is considered sensitive. This is because data will encapsulate various aspects of the physical environment, including highly personal information about individuals, groups, and companies, and can also have physical consequences, e.g., actuating commands. It is, therefore, important that security mechanisms are designed to take account of the potential sensitivity of the data. A recent example illustrating a failure of such involved a baby monitor, where an OS device on the local network could listen in without being subjected to access controls [16]. Furthermore, any device that had ever accessed the monitor could then *remotely* listen in, anytime, anywhere. This represents a clear failure to recognize and/or account for the sensitivity of the audio feed.

Identifying the "thing" that produces data may not always be sufficient to determine how sensitive its data. For example, a location sensor may be considered as generating sensitive data when representing the movements of a particular person, but the data produced by the same sensor may be less sensitive when itis attached to freight in transit. Furthermore, sometimes, only specific items/data-instances are highly sensitive, even when produced by the same "thing," e.g., a facial recognition device in a public space could provide the current location of the Prime Minister, thus having national security implications.

B. In-Cloud Data Protection

This concerns the cloud provider protecting data within their service, by preventing data leakage: (1) during transmission; (2) during processing; and (3) when data is stored "in the cloud." In all cases, data should not flow to unauthorized parties, including cloud insiders as well as cloud users [17]. With respect to communication, some cloud providers now apply TLS internally within their infrastructure, including data centers to protect against any internal threats or security breaches. This appears largely in response to recent highly publicized security breaches, such as those carried out by the U.S. National Security Agency.

The business model of cloud service provision is based on economies of scale, through services that share resources. For example, tenants may share the same physical machine by running above separate VMs during processing. Therefore, cloud providers ensure strong *isolation* between cloud tenants/users to prevent the leakage of data between them. This isolation can occur at different levels, including the OS (containers) [18], VM (hypervisor) [19], and in hardware. If storage is provided, depending on the level of isolation, the service offering might implicitly segregate all resources from others. Other levels of isolation may involve shared data storage infrastructure and software, such as shared databases, and thus rely on standard access control technologies

Cloud providers invest significant resources into ensuring strong access controls and complete isolation. Some are important for IoT-cloud, as well as for cloud service provision in general. Concerns over the extent of provider access (by cloud insiders) do not only concern datathat can objectively be considered highly sensitive. Rather, there may be laws that lead to particular data management obligations. Or, simply, there may be little trust in the cloud provider, e.g., if they reside in a jurisdiction that lacks a robust data protection regime. In this case, the "thing" may decide to encrypt the data it uploads to the cloud.

4. MALICIOUS THINGS

A. Malicious "Things" Protection of Provider

The cloud provider will maintain various accesses, and other controls, to protect against specific attacks, e.g., a rogue "thing" attempting to exploit the service, perhaps through some sort of injection attack. Even if attacks are successful, cloud isolation mechanisms offer containment, limiting their fallout. Such attacks are not unlike the security concerns of the cloud as it is today. Previous sections have explored how IoT dramatically increases scale, where there is the potential for a vast number of "things" to interact directly with a provider. Thus, one clear IoT-cloud vulnerability is cloud *denial of service* (DoS), which could potentially be launched from a large number of compromised "things." Cloud services are naturally *elastic*, designed to rapidly scale up/down resources in response to increases in demand, but still remain vulnerable to DoS [21]. Therefore, there is a need to explore more advanced DoS techniques in light of the fact that IoT greatly increases the scope of such an attack, particularly as "things" become more integrated with cloud services.

B. Malicious "Things" Protection of Others

Since the cloud can operate as a mediator and coordinator between "things," it offers potential in terms of improving security across the IoT ecosystem. This is because the cloud provides a natural "choke-point" between "things," in which security policy can be implemented and enforced. Requiring input data to pass through a validation process allows the cloud to effectively disconnect (or ignore inputs from) "things" that are detected as compromised. This also helps ensure data integrity, as only valid data (in terms of rate/format) rather than that from a faulty, compromised, or inappropriate (but perhaps non malicious) "thing" can enter a (possibly shared) database or flow to others via the cloud. Furthermore, there is scope for the cloud to be used more proactively,

e.g., by issuing control messages to "things" to turn them off (or adjust some parameters) where necessary, or perhaps to trigger software/firmware updates.

A fundamental consideration is in determining the "things" that have been compromised. This will be relevant at different levels, depending on the circumstances for instance, approaches could involve determining the malicious or untrustworthy nodes in a network [22], analyzing the data outputs, patterns of behavior or reputation of a "thing" [23], or perhaps involve human intervention, e.g., reporting a device as stolen. Work is required on developing such techniques, in line with new developments in technologies and their uses.

5. TRUST IN THE CLOUD PROVIDER

A. Trustworthiness of Cloud Services

A general concern is how much trust can be placed in a cloud service provider; i.e., that they will properly: (1) secure their service; (2) ensure it is correctly configured; (3) report leakages/issues; and (4) use data only for their intended purposes. Key to building trust is providing some degree of visibility/transparency over the cloud service. How this might be enhanced through audit, including when using external, third party cloud services and controlling where data is located in order to abide by regulations.

Recent developments in hardware technologies [24] enable new levels of trust, providing TPM [12] and remote attestation for cloud computing [25]. These can work to increase the level of trust that tenants have in the provider; for instance, by enabling data integrity and confidentially to be guaranteed regardless of the platform on which the data is processed [26],or to provide guarantees concerning the physical location of data [27]. Such techniques are reaching maturity, e.g., IBM is rolling out a scalable TPM-based cloud platform [25], [28].

It is likely the case that end-users are more willing to trust well-established and known cloud providers, rather than those with little history or reputation, such as startups offering cloud hosted applications. Several projects have focused on preventing the misuse or leakage of data by cloud applications through complex isolation mechanisms [29], which enables the control policy to be attached to data (potentially by "things," tenants, or providers) in order to control the flow of data, and to generate audit logs. More generally, having mechanisms that limit data mismanagement are crucial to enabling the wide-scale vision of information sharing underpinning the IoT.

6. CONCLUSION

Cloud computing and Internet of Things (IoT) is an emerging technique getting growing interest in networking research community. The support of cloud to IoT places different security and privacy issues. In this paper we have identified and described security-related considerations within the following broadrange of concerns: issues of data transport to/from cloud services and data management in the cloud, issues associated with identity management, issues arising from malicious "things" and issues of certification, trust, and compliance with regulations and contractual obligations. This work is aimed serve as an introductory material to people who are interested in pursuing research in this area.

REFERENCES

[1] P. Middleton, P. Kjeldsen, and J. Tully, Forecast: The Internet of Things, Worldwide. Gartner, 2013 [Online].

[2] S. Jankowski, J. Covello, H. Bellini, J. Ritchie, and D. Costa, "The Internet of Things: Making Sense of the Next Mega-

Trend," Goldman Sachs, New York, Tech. Rep., 2014. [Online].

[3] A. Zanella, "Internet of Things for smart cities," IEEE Internet Things J., Vol. 1, No. 1, pp. 22–32, Feb. 2014.

[4] J. Bacon et. al., "Personal and social communication services for health and lifestyle monitoring," in Proc. 1st Int. Conf. Global Health Challenges (Global Health'12) IARIA DataSys 2012, Venice, Italy, 2012, pp. 41–48.

[5] M. Kovatsch, S. Mayer, and B. Ostermaier, "Moving application logic from the firmware to the cloud: Towards the thin server architecture for the Internet of Things," in Proc. 6th Int. Conf. Innov. Mobile Internet Serv. Ubiq.Comput. (IMIS), 2012, pp. 751–756.

[6] T. Dierks and C. Allen, "The TLS protocol version 1.0," IETF, Tech. Rep., RFC 5246, 1999 [Online].

[7] D. McGrew and E. Rescorla, "Datagram transport layer security (DTLS) extension to establish keys for the secure real-time transport protocol (SRTP)," Internet Engineering Task Force (IETF), RFC 5764, 2010 [Online]. Available: http://datatracker.ietf.org/doc/rfc5764/

[8] S. L. Keoh, S. Kumar, and H. Tschofenig, "Securing the Internet of Things: A standardization perspective," IEEE Internet Things J., Vol. 1, No. 3, pp. 265–275, 2014.

[9] P. Urien, "LLCPS: A new security framework based on TLS for NFC P2P applications in the Internet of Things," in Proc. IEEE Consum. Commun.Netw. Conf. (CCNC), 2013, pp. 845–846.

[10] R. J. Anderson, Security Engineering: A Guide to Building Dependable Distributed Systems, 2nd ed. Hoboken, NJ, USA: Wiley, 2008.

[11] Z. Durumeric et. al., "The matter of heartbleed," in Proc. Internet Meas. Conf. (IMC), 2014, pp. 475–488.

[12] T. Morris, "Trusted platform module," in Encyclopedia of Cryptography and Security. New York, NY, USA: Springer, 2011, pp. 1332– 1335

[13] C. Lesjak, T. Ruprechter, J. Haid, H. Bock, and E. Brenner, "A secure hardware module and system concept for local and remote industrial embedded system identification," in Proc. Emerg. Technol. Factory Autom. (ETFA), 2014, pp. 1–7.

[14] M. Hutter and R. Toegl, "A trusted platform module for near field communication," in Proc. Int. Conf. Syst. Netw. Commun. (ICSNC), 2010, pp. 136–141.

[15] J. Singh and J. Bacon, "On middleware for emerging health services," J. Internet Serv. Appl., Vol. 5, No. 6, pp. 1–34, 2014.

[16] N. Dhanjani. (2013). Reconsidering the perimeter security argument [Online]. Available: http://www.dhanjani.com/docs/Reconsidering%20 the%20Perimeter%20Security%20Argument.pdf, accessed on Jul. 21, 2015.

[17] J. Singh, J .Powles, T. Pasquier, and J. Bacon, "Data flow management and compliance in cloud computing," IEEE Cloud Comput.Magaz. SI Legal Clouds, 2015, to be published.

[18] S. Soltesz, H. Pötzl, M. E. Fiuczynski, A. Bavier, and L. Peterson, "Container-based operating system virtualization: A scalable, highperformance alternative to hypervisors," in ACMSIGOPS Operating Syst. Rev.. New York, NY, USA: ACM, 2007, Vol. 41, No. 3, pp. 275–287.

[19] P. Barham et. al., "Xen and the art of virtualization," in SIGOPS OperatingSyst. Rev., 2003, Vol. 37, No. 5, pp. 164–177.

[20] I. Anati, S. Gueron, S. P. Johnson, and V. R. Scarlata, "Innovative technology for CPU based attestation and sealing," Intel Corporation, 2013 [Online]. Available: https://software.intel.com/sites/default/files/article/413939/hasp-2013- innovative-technology- for- attestation -and-sealing.pdf.

[21] S. Yu, "DDoS attack and defence in cloud," in Distributed Denial of Service Attack and Defense. New York, NY, USA: Springer, 2014, pp. 77–93.

[22] T. Moore, J. Clulow, S. Nagaraja, and R. Anderson, "New strategies for revocation in ad-hoc networks," in Proc. 4th Eur. Conf. Secur. Privacy Ad-Hoc Sens. Netw., 2007, pp. 232–246.0

[23] P. Michiardi and R. Molva, "Core: A collaborative reputation mechanism to enforce node cooperation in mobile ad hoc

networks," in Proc. IFIP TC6/TC11 6th Joint Work. Conf. Commun. Multimedia Secur. Adv. Commun. Multimedia Secur. Norwell, MA, USA: Kluwer, 2002,pp. 107–121.

[24] Intel, Software guard extensions programming reference, Tech. Rep. 329298-001US, 2013 [Online]. Available: https://software.intel.com/sites/default/files/329298-001.pdf, accessed on Jul. 21, 2015.

[25] S. Berger, R. Cáceres, K. A. Goldman, R. Perez, R. Sailer, and L. van Doorn, "vTPM: Virtualizing the trusted platform module," in Proc. Secur.Symp., 2006, pp. 305–320.

[26] A. Baumann, M. Peinado, and G. Hunt, "Shielding applications from an untrusted cloud with haven," in Proc. 11th USENIX Conf. Oper. Syst. Des. Implement. (OSDI), 2014, pp. 267–283.

[27] K. R. Jayaram, D. Safford, U. Sharma, V. Naik, D. Pendarakis, and S. Tao, "Big ideas paper: Trustworthy geographically fenced hybrid clouds," in Proc. ACM/IFIP/USENIX Middleware, 2014, pp. 37–48.

[28] S. Berger, K. Goldman, D. Pendarakis, D. Safford, E. Valdez, and M. Zohar, "Scalable attestation: A step toward secure and trusted clouds," in Proc. Int. Conf. Cloud Eng. (IC2E), 2015, pp. 185–194.

[29] S. Lee, E. L. Wong, D. Goel, M. Dahlin, and V. Shmatikov, "πBox: A platform for privacy-preserving apps," in Proc. 10th USENIX Symp. Netw. Syst. Des. Implement., 2013, pp. 501–514.