Editors: Mohamed Kaâniche, mohamed.kaaniche@laas.fr | Aad van Moorsel, aad.vanmoorsel@ncl.ac.uk

Quantifiably Trusting the Cloud:

Putting Metrics to Work

Ruben Trapero, Jesus Luna, and Neeraj Suri | Technische Universität Darmstadt

Cloud computing, with its purported myriad technological and economic advantages, is enabling the spread of innovative applications and services. And yet its fuller uptake is often constrained, mostly because cloud service customers (CSCs) perceive a lack of transparency in the security and privacy assurances of cloud service providers (CSPs).[1] This issue of trust is especially relevant now that a growing number of CSPs offer diverse cloud services, including virtual machines and storage, that enable complex services and workflows, which in many cases leverage multiple CSPs.

CSPs typically employ mechanisms that protect the data at rest, data in transfer, and data being processed to ensure a level of security and compliance with relevant data protection regulations that many small and medium enterprises couldn't otherwise afford. However, this requires the CSCs to yield control of their data to CSPs, leaving it potentially exposed or exploitable. CSCs' concerns about losing control over their data doesn't necessarily imply that providers are intentionally malicious. CSPs might violate the security and privacy commitments, as stated in their service-level agreements (SLAs)[2] and certifications, for many reasons: because of negligence, as the result of an attack, or because the jurisdiction in which they operate forces them to give law enforcement agencies access to cloud data.[3]

The lack of transparency in service provision and in the techniques and tools for obtaining trustworthy assurance of a CSP's compliance with security and privacy commitments symbolizes this "to trust or not to trust" dilemma. Hence, the interplay among security, privacy, and risk is critical to building trust.

Figure 1 highlights the security and privacy facets that formalize the quantitative service aspects related to security, privacy, and performance indicators. These are supported by metrics advocated in the service provision domain, such as draft ISO/IEC 19086-2 for cloud SLA metrics.[2] (For more information, see the "What Are Metrics?" sidebar.)

Naturally, any system that deals with security and privacy has valuable assets to protect, for example, customers' personally identifiable information. Malicious attackers or unexpected events such as power outages can imperil a CSP's capability to protect these assets. Managing the risk associated with these undesirable events is essential to prevent the potential impact of incidents. Evaluating such risks allows CSCs to define their overall trust requirements, or risk profile, by outlining the desired cloud service's pertinent security and privacy attributes.[4]

But how exactly do CSCs assess whether the CSPs are protecting their resources and properly managing their services in accordance

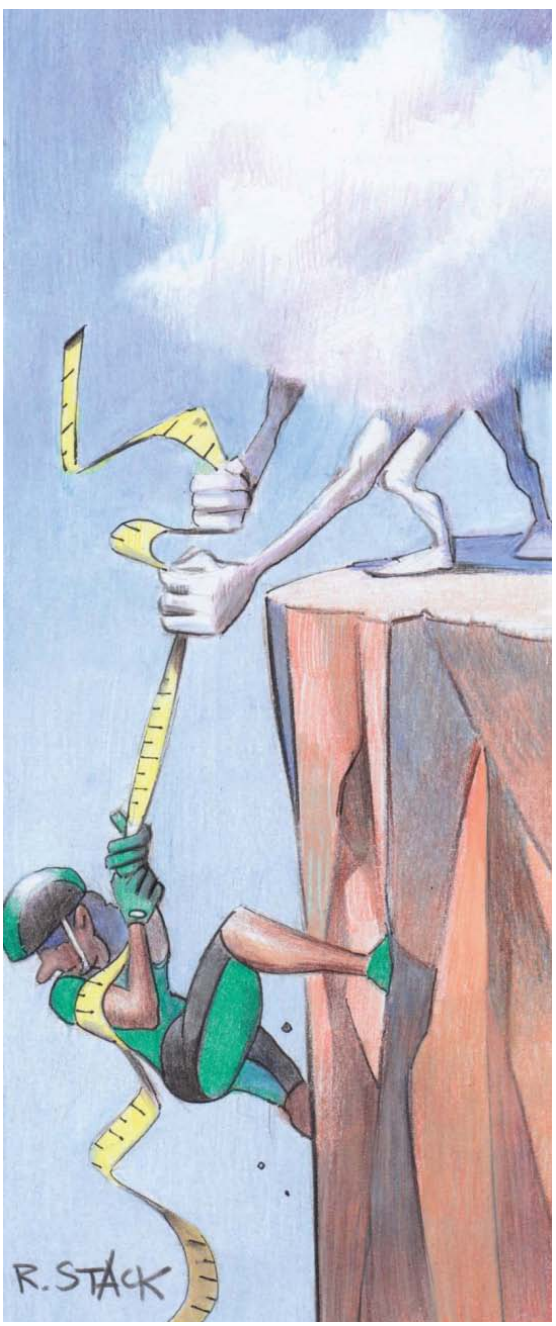

1540-7993/16/$33.00 © 2016 IEEE Copublished by the IEEE Computer and Reliability Societies

What Are Metrics?

According to NIST 500-37 (http://csrc.nist.gov), a metric is a "standard of measurement that defines the conditions and the rules for performing the measurement and for understanding the result of a measurement." Metrics are commonly used to set the boundaries and margins of the levels that cloud service providers (CSPs) are able to provide (along with their limitations). The expected levels for a specific metric are called service-level objectives (SLOs), which are typically included in service-level agreements (SLAs). In the security domain, some security control frameworks, such as Cloud Security Alliance's Cloud Control Matrix, divide metrics into security control groups and these groups into categories. The metric values are obtained via underlying measurements (monitoring) of the system. Table A illustrates an example vulnerability/patch-management control group. This control group is associated with four metrics. For example, scanning report age represents the age of the latest vulnerability report and is given in hours, whereas repository availability is a Boolean denoting whether the repository containing the vulnerability reports is actually available.

Table A. Example security category, group, and metrics.

Control category	Control group	Metric	Value type
Threat and vulnerability management	Vulnerability/patch management	Scanning report age	Integer (hours)
		Repository availability	Boolean
		Vulnerability list availability	Boolean
		Scanner availability	Boolean

with these trust requirements? Do CSCs have the tools and information to determine their CSPs' trustworthiness? These questions place trust as a transversal concept interlinking security, privacy, and risk (see Figure 1).

Quo Vadis?

Security assurance and transparency are the main requirements for instilling CSC trust in CSPs (see the "Metrics-Based Assessment in a Nutshell" sidebar). Addressing the trustworthiness of cloud services, particularly from a CSC's perspective, requires coping with several challenges:

- the challenge of a common reference model to represent security and privacy in the cloud,
- the challenge of continuous assessment techniques to evaluate security and privacy to obtain runtime assurance levels for CSCs,
- the challenge of continuous risk assessment methodologies for security and privacy, and
- the challenge of continuous CSP monitoring.

A Common Reference Model

Current models typically represent security and privacy agreements as a set of textual statements and legal requirements. This is not only unattractive but also quite often incomprehensible to CSCs who aren't security or privacy experts. CSPs must transition from this static concept of textual agreements and policies to machine-readable and automatically manageable SLAs that include security and privacy commitments as well as the requisite metrics to verify their fulfillment.

Continuous Assessment Techniques

Currently, a sprinkling of techniques exists for quantitatively assessing the level of security and privacy that CSPs provide. Although these techniques are promising starting points, they rely on static repositories with security self-assessments—the CSPs themselves declare how they cover different aspects as defined in SLAs (such as the STAR repository; cloudsecurityalliance.org/star). Additional uncertainties arise because there's no way for CSCs to know whether CSPs are actually fulfilling their specified commitments. Third party–based cloud security certification schemes (such as ISO27001 and the Cloud Security Alliance's Open Certification Framework; https://cloudsecurityalliance.org/group/open-certification) constitute the sole source of static, single-point-in-time "trust." In reality, these discrete snapshots of compliance might not correspond to what the CSP is actually providing in its operational environment, resulting in inaccurate CSC expectations.

Continuous assessment techniques, on the other hand, rely on better reliability information extracted not only from static policies such as SLAs but also from

direct monitoring of CSPs. This gives CSCs up-to-date information on how CSPs are using their data and managing their services as well as a way to refine the initial set of security and privacy requirements. CSPs can also use the results of continuous assessment to adapt their offers to potential CSC requirements, thereby customizing the service to CSCs' needs.[5]

Continuous Risk Assessment Methodologies

CSP evaluations should rely on risk assessment techniques associated with the security and privacy metrics requested by CSCs. Risk assessment techniques should evaluate threats, vulnerabilities, and patterns of previous incidents to ascertain the potential impact of incidents associated with the security and privacy metrics enforced by a CSP. The results of risk assessments can help CSCs define and update their security and privacy requirements; for example, a CSC might demand a higher level of security for an area that is at high risk for unfulfillment. These results can also provide useful information that allows the fine-tuning of continuous assessment techniques.

Continuous CSP Monitoring

Leveraging the fulfillment of security and privacy agreements is an important part of CSPs' operational management and depends heavily on their continuous monitoring capabilities.[5] Continuous monitoring requires clearly defined security and privacy metrics as well as nonintrusive techniques for obtaining measurements. The data collected using these continuous monitoring techniques becomes the input required for assessment. Observing the target CSPs is also imperative to identify potential deviations from their agreements, thus leveraging the design of remedial activities to fulfill the agreements.

Metrics-Based Assessment in a Nutshell

Security assurance and transparency are two of the main requirements to leverage cloud service customers' (CSCs') trust in CSPs. The typical assessment comprises three steps: elicitation, quantification, and evaluation. During elicitation, the CSCs' requirements are retrieved; for example, they might have requested scanning report age in hours. In the quantification stage, all possible values of a metric are normalized to a common scale across levels. For example, scanning report age can be considered to have three possible values—24, 48, or 72 hours—that can be mapped to three possible levels—1, 2, or 3, respectively. Similarly, repository availability can have the possible Boolean values of true or false, which can be mapped to the levels 1 or 0, respectively. Once the metrics are quantified to a common scale, the evaluation stage applies a set of aggregation techniques to compare the quantified metric values to CSCs' requirements. The comparative security assessment results can help CSCs choose the best CSP for their requirements or help them ascertain a CSPs' degree of either conformance or shortcoming in meeting the agreed-upon requirements.

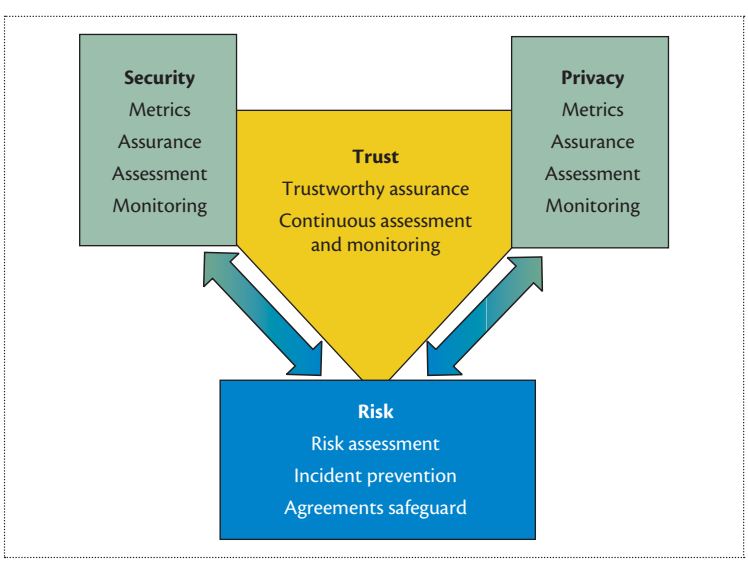

Figure 1. Trusting the cloud: interplay among security, privacy, and risk.

Putting Life Cycles to Work

Discrete solutions typically solve discrete problems. Moreover, trust is an end-to-end attribute. Consequently, our advocacy for trust as a general systems property revolves around two themes: the CSC and CSP interplay as a life cycle (see Figure 2); and the quantification of security, privacy, and risk, projecting trust as the quantifiable degree of reliance one can (or should) put in cloud services. (For more on evaluating security using metrics, see the sidebar.)

The proposed process starts with CSPs declaring the security and privacy agreements to be fulfilled (see Figure 2). These agreements also define the security and privacy metrics to be monitored, along with the CSCs' security and privacy requirements.

The monitoring information collected from the CSP can be used in various ways. For instance, risk assessments use monitoring data to perform real-time estimation of the risk associated with the security and privacy aspects in their policies and

Evaluating Security Using Metrics

Evaluating the quantified metrics is the final step in the security assessment process. Only a handful of methodologies exist to evaluate CSP security levels, with *quantitative policy trees* (QPT) and *quantitative hierarchical process* (QHP) being the prominently used techniques.[1]

Both methodologies use as input the security metrics organized in a hierarchy that groups metrics into controls and controls into categories (see Figure A). The CSC requirements and security metrics are quantified and mapped to the hierarchy's lowest nodes. Logical operations are possible across these quantified values.

Subsequently, a set of aggregation rules are applied progressively to obtain a final assessment score. The specific operations and aggregation rules used are what differentiate QPT and QHP.

QPT defines AND/OR relationships between the branches of the hierarchy (see Figure A1). This allows for the representation of dependencies across the metrics. The metrics that are mandatory due to regulatory compliance are linked by an AND relationship. QPT normalizes the distance between the CSC's requirements and the CSP's offered level with respect to the maximum metric level. These values are aggregated to obtain the final assessment score.

Alternatively, QHP is based on the analytic hierarchy process. QHP creates a pairwise comparison matrix that's populated with the relative ratios between the maximum possible level of the metrics and the levels offered by the CSP. The final score is obtained by using the CSC requirements to calculate a priority vector: the normalized eigenvector of the pairwise comparison matrix. In general, QHP increases the flexibility of the analysis because it allows partial scores at different hierarchy levels (see Figure A2).

Reference

1. J. Luna et al., "Quantitative Reasoning about Cloud Security Using Service Level Agreements," *IEEE Trans. Cloud Computing*, vol. PP, no. 99, 2015; http://dx.doi.org/10.1109/tcc.2015.2469659.

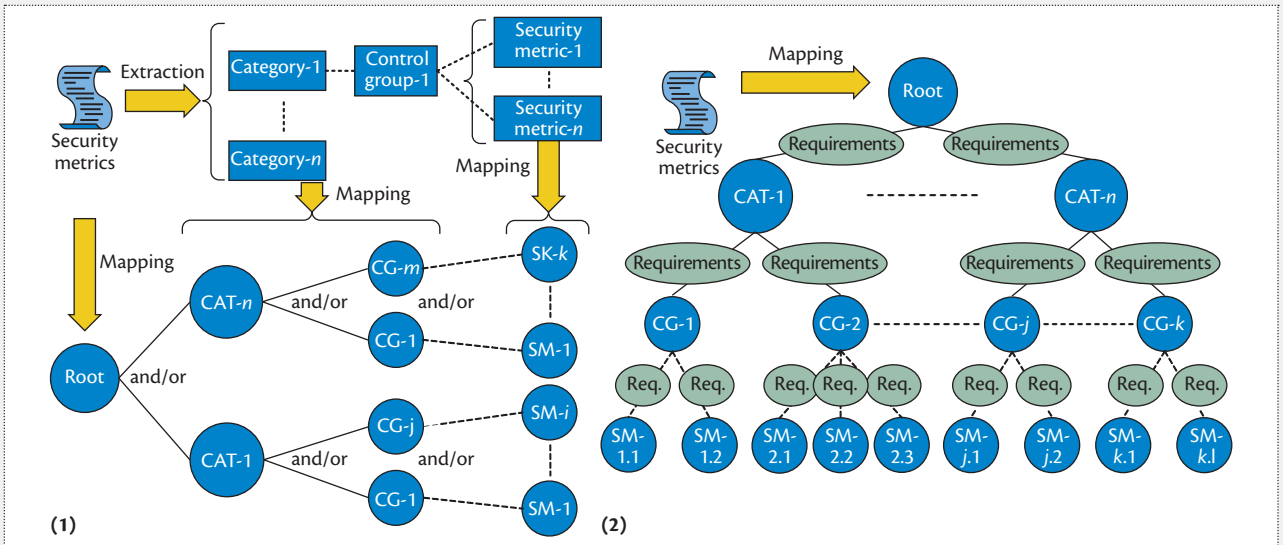

Figure A. Hierarchy tree for security metrics in (1) quantitative policy trees and (2) quantitative hierarchical process.

SLAs. A risk assessment can also help CSCs define their security and privacy requirements. For example, they can receive warnings about the metrics most likely to be affected by threats or other eventualities. The monitoring information also constitutes useful input to support continuous CSP assessment. Such continuous assessment methodologies naturally complement risk assessment by allowing estimation of the impact of threats or other events that might affect security and privacy.

The result of such assessments can help CSCs decide which CSP best matches their requirements or

can allow them to analyze the actual quality of a CSP's services. And, of course, continuous assessment gives CSPs feedback on requirements and threats so that they can address specific security or privacy metrics or update their SLAs and policies.

Emerging solutions for assessing cloud trustworthiness must consider the interactions among security, privacy, and risk. They will require a quantitative mind-set to use, and should be an integral part of life cycles that increase transparency in the provision and adoption of trustworthy cloud services. ■

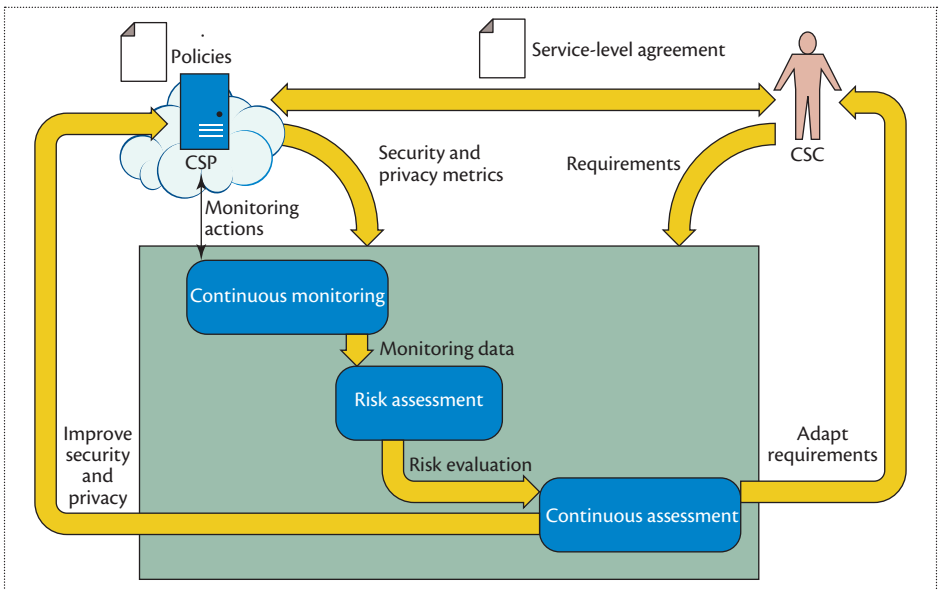

Figure 2. Trustworthy life cycle approach. CSC is cloud service customer, and CSP is cloud service provider.

References

1. K. Giannakouris and M. Smihily, "Cloud Computing-Statistics on the use by Enterprises," Eurostat, Nov. 2014; http://ec.europa.eu/eurostat/statistics-explained/index.php/Cloud_computing_-_statistics_on_the_use_by_enterprises#Factors_preventing_enterprises_from_using_cloud_computing.
2. "ISO/IEC DIS 19086-1 Information Technology—Cloud Computing—SLA Framework—Part 1: Overview and Concepts," Int'l. Org. Standardization, 2016; www.iso.org/iso/home/store/catalogue_tc/catalogue_detail.htm?csnumber= 67545.
3. D. Yadron, S. Ackerman, and S. Thielman, "Inside the FBI's Encryption Battle with Apple," *Guardian*, 18 Feb. 2016; www.theguardian.com/technology/2016/feb/17/inside-the-fbis-encryption-battle-with-apple.
4. J. Luna et al., "Quantitative Reasoning about Cloud Security Using Service Level Agreements," *IEEE Trans. Cloud Computing*, vol. PP, no. 99, 2015; http://dx.doi.org/10.1109/tcc.2015.2469659.
5. J. Luna et al., "Leveraging the Potential of Cloud Security Service-Level Agreements through Standards," *IEEE Cloud Computing*, vol. 2, no. 3, 2015, pp. 32–40.

Ruben Trapero is a senior researcher at Technische Universität Darmstadt. Contact him at rtrapero@cs.tu-darmstadt.de.

Jesus Luna is a senior researcher at Technische Universität Darmstadt. Contact him at jluna@cs.tu-darmstadt.de.

Neeraj Suri is a professor at Technische Universität Darmstadt. Contact him at suri@cs.tu-darmstadt.de.