

Questionnaire – Demographic Questions

1. What is your current role in the company?

a senior software developer

2. What kind of tasks do you usually do in your work?

Traditional software engineering tasks like software design and software architecture, but also the implementation. And also a bigger part: mentoring for, you know, for other people.

3. Given enough time, can you understand the architecture of an application system that is described using an IaC script of an IaC technology you are familiar with?

Yeah. Or I must say, I hope so.

4. For how many years have you worked on tasks associated with IaC tools?

(Since) 2012

5. How large is the company you currently work for?

< 50

Questionnaire – Compliance Rule Modeling and Checking

6. How do you check the compliance of the software applications of your company?

Manual(ly) by checking the IaC scripts.

7. Do you use well-defined models for the compliance rules applicable to the software applications of your company?

(Yes)

- a) If so, how do you define them?

Well, we have a kind of a, I wouldn't say comprehensive, but a, but a semi comprehensive [Excel], [Excel sheet] that lists traditional or, or, or let's say common security issues when it comes to running software in the cloud or, or over the Internet, and there are certain aspects covered and we analyze each, each aspect and, and ask ourselves if, if we have this, this aspect covered on the one hand in our application software(,i.e.,) in our software code as well, as well, as well as in our IaC scripts, or in our, our final deployment.

8. Do you think having a well-defined and machine-readable format for compliance rules reduces the complexity associated with checking them?

Yes, of course. Yeah. At least the time consuming factor would, would, I would say, hopefully decrease. (...) At the moment, we, we, we do not have a, have a fixed schedule where we check these rules. We, we just do this from time to time when we, when we think, okay, now it's, it's, yeah, it's again time to check them, and if there would be something machine readable or, or something that we can automate, we, we could do it with every release, for example, you know, this would be a big benefit.

9. Do you think having a well-defined and machine-readable format for compliance rules reduces the uncertainty associated with interpreting them?

Yes, of course, yeah.

10. How often do you have to deal with new compliance rules?

Very rarely, not very often. I would, I would, I would have (to say) a year

11. How much do you agree with the following statement: *using IACMF reduces the effort associated with defining and checking compliance rules?*

(Do) we assume that these plugins are already there? or, it's a part of defining (the compliance rule) also writing, writing these (refinement) plugins?

So just defining the rule in the UI seems pretty simple, and if you got the script parts ready and stuff like that, and as I said, if you have the, the right amount of plugins available, like script execution plugin, or I don't know if there's a plugin for Kubernetes, then I would say, I totally agree that it reduces the effort of defining and checking. (However,) it depends on the number of plugins. If you have to implement your plugins with every new compliance rule, I don't know, then it's, it's rather a (two).

Okay. You, you have one time the effort and, and then you have to benefit off of automation, but yeah, it's still effort...

Yeah, the thing is, our [Excel] list is, is very, as I said, it's, it's kind of high level. So there, there are different aspects: There's, there's a public list from the, from the [Owasp] group where they, where they publish their security concerns (...). For example, things like that, the communication to the outside should be never unencrypted and stuff like that. And, and, and, and from that very high-level point, you need, of course, you have to break it down for yourself in, in, in, in some smaller pieces that apply to your architecture.

(Furthermore,) from what I saw, (...) there are pretty good machine readable and also executable compliance rule for the infrastructure as code layer, but when it comes to, I don't know, a software-as-a-service configurations or, yeah, a platform-as-a-service configurations (...) I don't know how to deal with, with such compliance rules.

12. How much do you agree with the following statement: *using IACMF reduces the complexity associated with defining and checking compliance rules?*

I think (...) it depends on the compliance rule. I think for each different compliance rule, you may need a different person to define the rules behind and also to implement the plugin. So I think (from) a complexity perspective, it's highly specific on what kind of compliance you want to check: as you have shown in the video, if you go on to the infrastructure layer, you need people that are aware of (bash) scripting and aware of all these nuances in Ubuntu or in Debian or in whatever distribution. (It) could be the case that the compliance with Ubuntu seems totally different than for, I don't know, (Solaris), (...) so, I don't see that it reduces the complexity much. Therefore, I would say it's a, it's a, it's a two. Yeah.

I see benefit into automation, to, to be honest, if you, if you once have it ready, then, yeah, exactly, exactly, you know, yeah. Then I would say, yeah, totally agree.

13. How much do you agree with the following statement: *using well-defined models for compliance rules reduces the uncertainty associated with interpreting them?*

Then I would say, yeah, totally agree.

Questionnaire – Architectural Reconstruction

14. How do you reconstruct the architecture of running application instances you need to understand?

I would create from an instance model, a C4 model. It's like a UML (has) different perspectives, so a C4 model is like, you create layered architectures or diagrams: You start on the first level where you, where you highly shape your overall system and you identify your upstream systems, (i.e.,) the systems you do not have under your own control, and then you go one step further and you identify your main components that are shaping your overall application system, and then you may go on one step deeper (in which) you can create a class diagram of each component and to visualize how, how the components itself are constructed and stuff like that, and then the fourth layer is then the actual source code. The whole idea behind the C4 approach is that you can also generate models from your existing code (...) automatically, and vice versa. So if you specify the model beforehand, you could be able to generate some (stubs) and, um, skeletons out of it, so that, and at every point in time you have a link between the actual source code and the, and the overall architecture.

(Interviewer: but in this case you depend on the application design not on the instance model itself. Correct?)

Yeah, it's rather the software design instead of the actual (...) deployment aspects of it. You (do) capture some aspects like "This is a container or this runs on an Ubuntu system", but you do not capture "on what port" or stuff like that.

(Interviewer: so if you need to capture information about the running instance, what would you do?)

I would do it manually or by checking if there are already an IaC scripts in place: I would check the IaC scripts and check it side by side for each deployable component.

15. Do you use any (semi-)automated tools for this purpose?

No, no.

16. How much do you agree with the following statement: *using IACMF reduces the effort associated with reconstructing the architecture of running application instances?*

I would say (four) I'm not sure if this is always valuable to see the actual running application instances, for some degree, I agree.

Questionnaire – Compliance Violation Fixing

17. What do you do if you find out that a running application instance violates a compliance rule?

In high level terms, we create an issue and, um, and fix it: (...) we (either) deploy a new version of the application component itself or we apply a compliant configuration manually.

18. Do you use any (semi-)automated tools for this purpose?

We, we do everything completely automated. We have a pipeline (that applies the concept of) immutable architecture, So we are not able to, to, to change any configuration at, at runtime, because on the next deployment, it, it will get overwritten. So we do everything (declaratively), and use the IaC tool to calculate the new diff and to deploy these into our environments.

(Interviewer: But what do you do when, for example, the fixes affect a database component that hosts user data, just like the second use-case in the tutorial?)

Yeah, for stuff like that, (which can be summarized as) managed services, we apply manual configurations, like cloud services from Google or AWS or whatever.

19. How much do you agree with the following statement: *using IACMF reduces the effort associated with fixing compliance violations?*

If you, if you do have a plugin, yeah, I totally agree. But if you do not have a plugin, then (it would be) two. I'm just thinking about this database user use case you showed in the video: if we would apply this to our architecture, this would mean that we would need (...) a plugin to understand what kind of database we use in Google Cloud, and the plugin needs to read the users from the configured database, and (when it comes to fixing) needs to alter the user (using) the Google APIs.

20. How much do you agree with the following statement: *having well defined models for compliance jobs reduces the uncertainty associated with handling detected compliance violations?*

Yeah, absolutely, I totally agree.

Questionnaire – General Questions

21. How do you evaluate the novelty of the framework?

Hard to tell because I have never researched about compliance frameworks, so it seems nice, but from a conceptual perspective, I don't know if this is a novel approach. I don't know, but it seems you use or incorporate well-established concepts from the instance model or deployment model stuff, and combine that with this compliance layer and the way you designed the tool in a way that also (makes) the rules and the plug-ins kind of decoupled, (which) makes it very flexible, but also on the other hand, not that usable, you know, from a user or from a simplification perspective, I would say.

22. How do you evaluate the extensibility of the framework?

So, from what I have seen in the video, it is very good and very extensible from what I saw, so you can define your own rules, you can extend it with new rules, (and) you can also extend it with new plug-ins that enable you to run or incorporate, let's say, more advanced or more complex rules. **However, the plugins (make or break) the whole (framework).**

23. Would you use the framework in your work?

Yes

a) If so, in which areas?

I think. We would at least do a proof of concept if there would be plug-ins for Kubernetes and Google Cloud

24. What is your general impression?

My general impression is it's a nice tool (but) from what I saw, it would be a lot of effort, I think, to instrument. We would require, let's say, **a basic set of compliance rules** applicable for our

tooling before we really start investing in it or start using it because our time is unfortunately very limited. For example, a repository would be nice where you have the compliance rules, according to the Owasp, top 10 or top 20 security common issues. if these things are already captured in compliance rules for Kubernetes, for example, this would be very, very, very beneficial for us.