Magdeburger Journal zur Sicherheitsforschung

Gegründet 2011 | ISSN: 2192-4260
Herausgegeben von Stefan Schumacher und Jörg Sambleben
Erschienen im Magdeburger Institut für Sicherheitsforschung

This article appears in the special edition »In Depth Security Proceedings of the DeepSec Conferences«.
Edited by Stefan Schumacher and René Pfeiffer

IT Security Compliance Management can make sense

Adrian Wiesmann

What kind of internal and external controls from regulations and other sources are there? What is IT-Risk and IT-Compliance management? Why and for whom does it matter? How can we handle it and how does compliance aggregation fit into the picture?

We will then look at the SOMAP.org project which is an Open Source project working on tools to handle IT-Compliance aggregation and IT Security compliance management in general. We will discuss why compliance management is not only about hot air but can make sense when done right.

Citation: Wiesmann, A. (2015). IT Security Compliance Management can make sense. *Magdeburger Journal zur Sicherheitsforschung*, 2, 661–666. Retrieved November 24, 2015, from http://www.sicherheitsforschung-magdeburg.de/publikationen.html Version 2015/11/22 13:46

1 Introduction

1.1 What to expect

In late 2011 I gave a talk about IT Security Compliance Management at the DeepSec conference in Vienna. The presentation contained an introduction to the Security Officers Management and Analysis Project (SOMAP.org). SOMAP.org is a non profit organisation which develops tools around the topic of IT Security Compliance Management.

Back then I worked on the SOMAP.org project's few tools and documents. At DeepSec I was invited to present our latest thoughts on the topic and our idea where Compliance Management should head to.

Since then, many months went by and while a few things remain the same, we learned new things and changed plans here and there.

In this article we will look into what IT Security Compliance Management was for us back then (and still is today), what we think is wrong with it in general and how we tried to change that in 2011, as well as we try today.

After a look back, we will then talk about what the SOMAP.org project learned on the way up until today. And what the project probably would do different today.

But first we have to quickly define a few terms.

1.2 Terms used

Authority Document An Authority Document can be a statute, regulation, audit guideline, best practice and any other document containing one or multiple Controls which are relevant for your organisation and environment.

Control A Control is a requirement from either you or a body of authority. Controls need to be either implemented or need at least to be considered. Depending on the authority body issuing these Controls and your role.

Control Aggregation The process of taking all Controls from the relevant Authority Documents, to remove duplicates and to unify the rest of the Controls. This sometimes is also referred to as multi compliance.

2 The problems with Compliance Management

2.1 Amount of Controls

One of the problems with classic Compliance Management is that it only works theoretically. You have an amount of Controls from a specific amount of Authority Documents. You try to comply to every single one of them. But when the amount of Authority Documents grows, so does the amount of Controls and so does ultimately the complexity.

Looking at an average company, that company will have to follow some industry regulations, probably a few laws, especially regarding their bookkeeping and probably regarding some IT standards. That company actively decides on complying on specific Authority Documents and Controls either because the company thinks it makes sense or because it has to - say - follow ISO/IEC 27001 since it works for a third company which requires it's outsourcing partners to comply with the ISO/IEC 27001 standard. Banks and insurance companies are among the first which require compliance with specific standards from their partners. Not because it is a legislative need, but because they decided so themselves. Other companies started to follow that lead.

The number of Controls to be compliant with starts somewhat above hundred when using the ISO/IEC 27001 standard. Adding more Authority Documents to to your list of relevant Controls will quickly add to that sum.

2.2 The Disorder

With the different Authority Documents come different structures of documents, different structures of Controls, different wordings and sometimes however similarities between the Controls. Different Authority Documents may cover similar topics. In full or just in parts. Choosing multiple Authority Documents may force you to follow hundreds of Controls. Some of which are basically the same, some of which are worded slightly different or probably slightly diverge in substance and some of which are completely different to each other.

Which means, adding more Authority Documents to the mix will significantly raise the chance for you to have Controls which are completely the same, which have similar Controls or - probably as the worst case - which have Controls which are contradictory to each other. The fact that these Controls are probably issued by different issuing bodies and standards agencies will not work in your favour. As example: Controls can be about the same topic, written completely different and require the same.

Adding Authority Documents to your library of Controls means, you have to check every single Control with the Controls you already have in the library to make sure that you have a clean library without similar or contradictory Controls.

Of course you can live without that hassle. Just throw the relevant Authority Documents at your asset owners, telling them to follow what's written in there. Go ahead, try it out, you might be lucky to get away with it.

2.3 Compliance Management is not cool

The cool boys tend to make fun of Compliance Management. Of how to be compliant does not make you

secure. That we should invest in security and not in compliance. That Compliance Management is something done by boring old accountants and IT auditors.

Oh how we laughed about those jokes. But completely missed the point.

Compliance is not cool because it is often done for wrong reasons or done too much in the way of bookkeeping. Or both. Probably all of us heard stories about companies installing Web Application Firewalls (WAF) because somebody told them that only with those they are PCI-DSS compliant. While completely wrong - at least at the time of writing this article - it is also a recipe for disaster. Just throwing hard- or software at an infrastructure without understanding and maintaining that new piece will most probably end in tears. And so does Compliance Management when done wrong. Besides costing your company a heap of money without you gaining anything.

3 How to do Compliance Management

It is still our strong believe: If Compliance Management would be done right and with the right mind set, it would not only make sense but could be a bit of fun in the process as well. Authority documents can contain good practice, can help you to not forget anything relevant. But they are definitely not the excuse to stop thinking.

So here are a few points which we think are important when doing IT Security Compliance Management.

3.1 Do not reinvent the wheel

There are technologies and even products out there which can help you in achieving your Compliance Management goals. We will be talking about Compliance Aggregation in a bit, but let us just state here, that there is always a tool or technology out there which you can build upon. There are tons of technologies out there which you can use to automate things.

Take Asset Management as an example. To manage and report on your compliance level, you need to know your assets. You can either manage them on your own, in your own tool, with your own resources. Or you leverage already existing Asset Management tools and resources and concentrate on your compliance part.

Why should you try to get a hopefully complete list of assets in your company, when there are teams out there, which should know what assets there are. Your IT operations most definitely already has some Asset Management system in place. Use their data to learn about assets. Do the same with Facility Management to learn about rooms and facilities.

If you try to catalogue your assets on your own, we guarantee you that you miss parts of your landscape. And if operations finds out, they will surely be annoyed that you did not ask them for help.

3.2 Make things simple

Many Compliance Management tools we have seen and used before, seem mostly to be copies of the same same. There seems to be some central European based belief, that Compliance Management tools should contain tree based Asset Management functionality and percentage based to-do lists.

This is wrong on so many levels.

While managing assets in a tree and monitoring the degree of realisation of a task may work in smaller environments, it definitely does not work in medium to huge enterprises. As mentioned before: Doing Compliance Management is complex enough. Reuse the work of others. Keep things simple as long as possible.

Having to-do lists in a Compliance Management tool is twice wrong. First because percentage based to-do lists are not following the make-it-simple approach. Say, you have a Control which asks for logging of events, which your software does not do. Therefore you are not compliant on that Control. But to what percentage? Does that really matter? You are not compliant, that's the important part.

Second, why should your Compliance Management software contain a to-do list? Does your company not already run some kind of task management software? Where everybody has an account, which everybody knows how to use, which is integrated in the already existing infrastructure. What about using that instead of some proprietary solution?

When talking about making things simple: Always a good strategy to follow is to automate stuff. You do have an asset list in your company? Use some form of automation to import that into your tool-chain for further usage. It makes no sense whatsoever to use your precious time to copy and paste around data from one tool to the next.

3.3 Think outside the box

Thinking outside the box is something we were already talking about. Many tools and books tell you how to do Compliance Management wrong. Do they know your company and your environment? Think for yourself and find a way which works better for you. Better in the sense of easier, quicker, makes more sense in your company and which is generally less costly.

Let us explain that point with a short example.

When we started with our Compliance Management tool we played around with topic maps. Topic maps are a way to connect information in a way, that the connections between topics contain relationship information. Since everything can be a topic - buildings,

persons, organisations, countries, you name it - topic maps allows to represent data in a structured way. As an example, if you have an Author X and a Book Y, you can put them into a two-way relationship: Author X wrote Book Y. Book Y was authored by Author X.

Applied to assets in an organisation was resulting in many interesting thoughts. Theoretically we were able to automatically inherit attributes of assets. If a program knows that a database system contains sensitive data, it can inherit data protection Controls from the database to the server it runs on, to the room the server is put in. But it can also inherit Controls to the users having access to the database, to the server, to applications using that data. All of that without a tree based approach.

While we had many interesting thoughts regarding automation we unfortunately had to concentrate on just a few concepts. Topic maps back then were not as standardised as they are today. And there were not that many tools and libraries written yet to manipulate topic maps.

3.4 Compliance Aggregation

Mixing multiple Authority Documents quickly results in a mess, as discussed before. Compliance Aggregation is the strategy which in our opinion makes most sense and which we decided to focus on. Compliance Aggregation is the concept of taking all Controls from all the relevant Authority Documents and then to remove duplicates and to unify the rest of the Controls into a neat single catalogue of Controls.

We worked on a data model which allows us to have both: Non-aggregated, original Controls as well as aggregated Controls. All Controls from every relevant Authority Documents can be put into our database. This is a simple import and transformation process where you bring every relevant Authority Document into the structure of our database.

We then added another layer with all the aggregated Controls. Every aggregated Control knows from which Control or Controls it is coming from. Looking at the aggregated Controls catalogue gives you the full view of which parts of what Authority Documents are overlapping, and which parts are not.

As an example, you have a Control which states that you have to log authentication attempts against an application. Your aggregated Control »knows«, that it is based on a Control from CoBIT as well as on a Control from ISO/IEC 27001. A nice effect of this is, that being compliant with your aggregated Control automatically tells you that you are compliant with the linked Control from CoBIT as well as the one from ISO/IEC 27001. Asset owners only need to follow the aggregated controls, you immediately get the compliance level of all the relevant Authority Documents.

This makes working with Controls simpler but still leaves the option to understand where a Control is coming from and what the original authors idea was. When talking about Controls and how it is relevant to an organisation, you often need to know where it is originally coming from. Being able to see on which Authority Documents an aggregated Control is based on, helps much in understanding the relevance of said Control.

Oh and before you think about writing and aggregating your own aggregated Control catalogue. Do some Internet search and consider buying before making. While such a catalogue unfortunately is not in the reach of a hobbyists budget, it will save you from a huge amount of time aggregating Controls in a corporate environment.

3.5 Meta data model

While designing the data model which knows aggregated Controls, we realised another point. Controls from Authority Documents as well as your Aggregated Controls will not change that often. In the case of ISO/IEC 27001, as example, there is a new catalogue version about every 5 years.

We wanted to consider that fact and therefore decided on using a two tier database.

In the »meta tier« we store the non volatile meta data for every kind of stable data. We understand meta data as data which is not case specific, non volatile, not instance data. Meta data can be shared among users of different organisations. Meta data contains Authority Documents and Controls. It has a descriptive character in general.

In the other tier, the case specific »instance tier«, we store all the user specific and sensitive data. Instance data is quite personal to every user. It contains his or her assets, the compliance level of an organisation, data you would not necessarily want to share with anybody outside your organisation.

This means that Controls from Authority Documents are part of the meta data tier, which normal users probably never will have to touch. The idea is that a group of volunteers works on the meta tier and shares changes with all the users. Removing that hassle from normal users: We as an OSS project can work on descriptive data, generate good practice datasets and share that among all the users. So that everybody can benefit from each other. Good practice as it's meant to be.

This data model has some nice side effects. The meta model allows to automate some things. One of which is type inheritance. You attach your asset (the web server running your e-commerce store) to the asset type of »web server« (which is stored in the meta data tier we just talked about). You then link the respective custodian to his or her web server. That person will instantly know which Controls are relevant for him or her.

It is a bit like programming. In the meta tier, you write a model description of an object, how it reacts to which situation, what are its features, and such. When you run the program you create an instance

of such a model object. Every instance of that same object »knows« and reuses the information from the meta model object.

With such a construct you remove the need for your co-workers to read through all the Authority Documents and read more than they really need to know. Just ask them for which asset they are a custodian for and you can give them a list of relevant Controls and have them concentrate on those only.

3.6 Don't do silly calculations

When doing compliance checks, avoid to do silly calculations and estimations. Many Controls are not formulated in a way, that you could state the level or any percentage of compliance. Take the compliance level of a Control concerning logging, as example. Either you do log data in a specific format or you do not. There is no »I do it halve«, »I have it planned« or »I already do log some parts but not all«. Either you do log as described in a Control, then you are compliant, or you do not, then you are not.

Which means, it does absolutely not make any sense to rate compliance levels in percentages. This Control is implemented to 23

There are also tools and concepts out there which contain some calculations where risk is calculated on not implemented Controls. This only works if you can base your calculations on some facts. Measurable facts which are not based on gut feelings.

Some methodologies describe some kind of calculations with magic numbers here, magic factors there and a result of one, two or three bombs. What does that actually mean? What do two bombs mean to your company? Should you not better think about questions like: What is a high risk for my company in the current situation?

When thinking about formulas and results from such calculations, do they really reflect the risk landscape and risk appetite of your company?

From what we learned, silly calculations taken out of thin air do not work. While they might look good, you wont gain much from them. Calculations only work if you get a value at the end which is not green, three bombs or the value »low«, but of which you can understand what's your situation and risk you take in not being compliant with whatever Authority Document you chose to use.

4 What did change since 2011?

It is always fun to look back to what we did, how we did it and what we thought back then. In retrospection we probably learned these most important lessons.

4.1 Explain yourself

If you do not do what everybody else does, you have to explain yourself. Compliance Management is not a point and click matter, although some vendors make you think it is. Doing Compliance Management, you need to know your stuff, you need to understand concepts and most of all, you need to know your environment.

Technology based the SOMAP.org project was sound. We chose and built quite a few technologies which help to quickly add features to our application. But the main problem was, that we did things differently, so we needed quite a bit of explanation. Instead of developing more features, we had to explain where we were and why we were there.

It was difficult to follow the release-early-release-often paradigm since the application we wrote was not necessarily a single user tool, but needed a bit of installation on a server. Normal users do not install server software, they download and try apps.

4.2 Corporate Funding of an OSS project

The other problem was the lack of corporate sponsoring, or sponsoring in general. A project of our size definitely needs some form of sponsoring. While we were lucky enough to get some funding here and there, in the long run this was not enough.

It took quite an effort to work on some aggregation of Authority Documents to get started. Nowadays there is at least one catalogue out there which does exactly what we tried to achieve as a side project. But when we started, we lost quite some time with trying to do that as well.

First we thought it were a good idea to write an app running on the users computer. Somewhat in the middle of the progress we realised that corporate environments do not work like that. There you have central systems, connected to whatever infrastructure you have in that environment. Web based tools are preferred, because they can be integrated more easily, they do not need roll-outs and architecture does not get in the way of network topology (like database access through firewalls).

So the OSS project was constantly evolving, learning about and working on Compliance Management methodologies and writing software. The changes and the complexity of the projects scope, together with the lack of corporate funding finally drained it's power, lowering it's momentum. Not to a halt, but to a very slow pace.

4.3 SOMAP.org is not dead yet

The SOMAP.org project is not dead yet, it is merely resting. Waiting for the right event which lifts it off the ground again, giving it back it's momentum it once had. We think the project still has a good idea, there is still much potential in having a community work

on better tools for all things IT Security Compliance Management.

Looking at today's state of the industry, not that much has changed. While a few think the same as we do, many still do risk calculations by magic and annoy their custodians with interview based compliance checks. Which means that there is still a space for a community run tool and methodology.

If you are interested in the topic, in getting the SOMAP.org project and it's tools off the ground again, get in contact with us. You can reach the project via it's website.

5 About the Author

Adrian is working as an IT Security Officer for a Swiss financial institute. His dayjob is to bother, to pester and to annoy. Every single day he works hard to bring these qualities of his to perfection. With a background in software engineering he focuses on application security and software demolition but enjoys a fine hardware hack or a well executed social engineering stunt as much as everybody else does. He is one of the founders of SOMAP.org, a non-profit organisation which is authoring and publishing documents and tools for analysing and managing IT security risk and compliance with regulations and standards. Adrian holds a masters degree in information security from the Royal Holloway, University of London.