John Leach Information Security

TBSE Technical Description

(Introductory Extract)

Version 1

25 January 2018

John Leach Information Security

John Leach is responsible for the contents of this document.

Contact Details:

Dr John Leach
Director, John Leach Information Security
Tel.: (+44) (0)1225 332 134
Mob.: (+44) (0)7734 311 567
e-mail: john.leach@jlis.co.uk

VERSION	DATE	COMMENTS
1.0	25 January 2018	Extracted from V2.1 of the full TBSE Technical Description document

TBSE Technical Description (Introductory Extract)

Version 1

Contents of the full TBSE Technical Description

Introduction

JLIS is proposing Threat-Based Security Engineering (TBSE) as an analytical method for:

- understanding analytically rather than purely descriptively the dynamics that give rise to security risk,
- quantifying security risk objectively so its absolute magnitude and importance in each situation can be understood by stakeholders and those accountable for managing it, and
- quantifying the various things that go into creating security risk so informed decisions can be made by those charged with the day-to-day management of security risk.

TBSE is currently undergoing professional review in the UK to establish its originality, validity, analytical soundness and technical merits. To assist those reviewers, and to assist others who might have an interest in making early use of TBSE, JLIS has written a 40pp *TBSE Technical Description* that explains how TBSE works 'under the covers'. That document is available upon request (under cover of a signed NDA).

This document is an extract from that Technical Description document and is available freely via the JLIS web site (www.jlis.co.uk). It provides a slightly amended version of the opening three sections of the full document to explain to interested parties what TBSE is about.

What is TBSE?

TBSE is an analysis methodology specific to addressing security risk quantification problems. It takes the form of a paradigm (a way of looking at how attacks engage with targets) and the analytical methods needed to go with that paradigm, appropriate for analysing the dynamics behind risk and calculating absolute numerical values for risk-relevant variables.

An analysis methodology

TBSE is an analysis methodology. It is not a product or a piece of technological machinery that could be patented and built. It is also not a single, one-size-fits-all, generic model that applies to all security risk management questions. It is an analysis methodology for addressing a specific class of questions: security risk questions that involve the quantification of security risk and/or the quantification of the things that go into the creation of security risk.

The desire to quantify an enterprise's security risk, or perhaps to quantify the benefit of a particular security improvement proposal, can be thought of as business questions management might want answers to. TBSE is the methodology for answering such security risk questions. In a nutshell, TBSE enables someone (a security analyst) to structure the problem correctly, work out what analysis needs to be done, determine what data they need to gather, and then do the calculations that will give them the answer to the question that was asked.

Management's security risk questions

Many, if not most, of the security risk questions management might want answers to require the quantification of one or other aspect of security risk or of the things that go into the creation of security risk. For example, some of the security risk question TBSE might be used to answer are:

- How big or important is this threat, not just in terms of how prevalent is it but how potent is it?
- How good is that security control at interceding in the progress of this threat, either at blocking the threat, or at containing the operational impact of the threat, or at containing its cost impact?
- Risk is caused by what remains after a control has done whatever it does, not by how much the control blocks. So, how much of the threat is still active after that control has done its stuff?
- And by how much could I reduce that remaining threat activity if I were to change this or that aspect of the way that control is implemented and operated?
- How much has that particular threat cost me each quarter in lost productivity? And is the variation in my quarterly losses the result of a variation in the threat, a variation in the effectiveness of my controls, or just a variation in my luck?
- If I were to spend £100,000 of effort on improving that security control, how much could I expect to save in improved productivity or reduced breach losses? What is the pay-back period?
- Would I be better off putting my effort into improving this control or improving that control?
- Do I need this security product AND that security product? How much extra risk reduction do I get by having both products in place? To what extent is the second almost redundant given what the first one does?
- Given how this particular threat I am under seems to have changed in the past three months, how much has my risk changed as a result? If that trend in how the threat has changed continues, what will my risk look like in one, two, three quarter's time?
- How good are my security controls at reducing what my risk could be? Am I under-protected or over-protected? Should I award myself an A* or an F?
- Given the nature of my business and the way my productivity relies on my use of technology, how good do I need my security controls to be?
- Which controls are critical to my success and which are secondary? Which of my controls are stretched to their limits and which have spare capability I am not making full use of?

The form TBSE takes

TBSE is built around a paradigm, a particular way of thinking about how security risk arises. And it embodies a particular analytical approach appropriate to that paradigm, a particular style of analysis so security risk quantification problems can be analysed in the way they need to be analysed.

In a bit more detail:

- TBSE starts with a paradigm, a way of abstracting, thinking about and defining the components involved as security threats make progress within an enterprise's technology estate, and the dynamics by which they interact with vulnerabilities, get countered by controls, and end up generating an amount of security risk.
- It enables analysts to build conceptual models with which to understand, calculate and measure what is going on, from a risk perspective, at any point along a threat's pathway. It enables them to calculate absolute values for a variety of risk-relevant results, values that are specific to the threats that impinge on the target, the characteristics of the target, and the way protection is provided for that target. It also enables them to combine results to form aggregate risk results.
- It enables analysts to calculate specific risk results at any point within the threat pathway as a function of the amount and nature of the threat activity taking place there and the vulnerabilities and the controls that are in place. This includes, for example, being able to calculate the effect a particular control has on the progress of a particular threat at a particular point along its pathway, the amount of threat activity that the control intercepts and the amount of threat activity that survives engagement with that control. This last capability is an important one, as it is the threat activity that survives engagement with a control that carries the threat's remaining potential to generate security risk. Being able to quantify that amount is necessary for quantifying the resultant security risk.
- Being able to perform these types of calculation means an analyst can calculate the effect on the resultant risk that proposed changes to one or more controls would have. They can quantify, in £, $ or €, how much extra risk reduction benefit would be obtained for the amount (in £, $ or €) that it would cost to make those proposed changes.
- This then enables Security Managers to build business cases for proposed security investments. It enables Security Managers to demonstrate to business leaders how effective the enterprise's security activities are at providing security protection to their operations and to justify the budgets they need. It enables Business Leaders to set upper limits to the amount of security risk they are prepared to take on (aka the enterprise's 'security appetites'), and Security Managers to configure the enterprise's security controls to keep the enterprise's security risk within those set limits.
- And, finally, it makes security risk manageable for Business Leaders and puts Business Leaders in control. Business Leaders can get the insights and

analytical results they need for making informed and reasoned risk management decisions. And it enables Business Leaders to be held accountable by stakeholders for the level of security risk their enterprise takes on.

The analytical approach required

TBSE is built around a paradigm and a set of analytical methods. Before we look at that paradigm, we need first to understand what type of problems security risk quantification problems are and what type of analytical approach is required for their analysis. Security risk quantification problems require a non-deterministic approach, not a deterministic one. TBSE embodies a non-deterministic approach.

The source of security risk

Security risk arises as a result of the potential, if not actual, interactions between threats and the entity being threatened (the target). It is not an intrinsic characteristic of a target, and it is not a constant for the target. It is a varying consequence arising out of the effects those threats have, or could have, on that target.

The magnitude of the risk created varies in line with the threat. All other things being equal, the greater the expected number of attacks, the greater the risk; the greater the potency of those attacks, the greater the risk; the greater the expected impact of those attacks, the greater the risk. The magnitude of the risk also varies in line with the characteristics of the target, the technologies it uses, the controls it has in place and the vulnerabilities present.

Therefore, to calculate the magnitude of the risk that can be expected given the specifics of the threat and the characteristics of the target, one has to be able to structure and analyse how the threat and target engage with each other and how the interactions between the two play out. One has to be able to conceptualise the dynamics that take place when threats and targets engage, and one has to be able to analyse those dynamics in a manner that enables outcomes to be quantified.

As we shall now see, these analyses require a non-deterministic rather than a deterministic analytical approach.

Deterministic vs. non-deterministic problems

If one is seeking to analyse a single specific attack and its engagement with a specific set of vulnerabilities and controls, and if everything that might have an influence on the outcome of that engagement is known, then it might be possible for the engagement to be analysed deterministically. However, in reality, even for just a single attack on a target, not everything needed for a deterministic calculation will be known: the skill of the unknown attacker; the particular exploits the attacker is going to try to use; the speed with which they will alight upon a suitable open vulnerability; whether that will be a zero-day vulnerability, one for which the patch is just about to be applied, or one for which the patch has just been applied; the speed and expertise of the responders trying to anticipate the attacker's next move and lock unlocked doors; etc. etc. Efforts at deterministic analysis quickly become bogged down by the presence of all the unknowns.

Then, if one wants to broaden the analysis to look not just at a single attack but at a threat, where the threat is a stream of attacks arriving over a period of time, each attack being slightly different from any other, some using a newer or older exploit coming in just ahead of or behind a control update, with who-knows which different staff members responding to the different individual attacks, then the deterministic approach immediately grinds to a halt.

Add to that the fact that not every control is a Preventative control, and that the 'attacker / defender in hand-to-hand combat' way of conceptualising engagements doesn't apply to all types of control, and the deterministic approach simply breaks down. Analysts have tried to scale up deterministic analytical approaches for decades (going back to even before the original creation of CRAMM) without substantial success. Analysing security risk dynamics is not the type of problem one can just throw extra processing power at. The difficulty, fundamentally, comes down to the unsuitability of the approach people have been trying to use. Determinism just doesn't work for security risk interactions. A non-deterministic approach is what's required.

Security risk is not unique in being unsuited for a deterministic approach. For example, a meteorologist doesn't calculate the risk of flooding by trying to follow each drop of water as it circulates in the clouds and falls to the ground and flows downhill and joins the swollen river that threatens to break its banks. It isn't just that there are too many drops of water to follow, it is that there are too many determining factors that simply cannot be known. The meteorologist cannot calculate precisely when and where each drop will fall, and how soaked the immediate piece of ground the drop falls on will already be, and how that drop's water will work its way over stones and around vegetation and through the soil to get to the river, and cannot then add all the billions of drops together to aggregate the results.

Consequently, meteorologists analyse the risk of flooding by building circulation and flow models, and by calculating the probabilities and rates for the different outcomes that might arise. Then they measure the rainfall patterns in the vicinity of interest, and the catchment areas and retention and latency of the local water courses, and relate the risk of flooding to those characteristics.

Land Managers then decide what flood defences they need for managing their flood risk given the possible and predicted levels of rainfall. They select which defences they might want to build, they model and analyse the effect those defences would have on the dynamics by which flooding arises (how they divert catchment, how they increase retention and latency), they decide what settings they would need for each defence (how deep the stream would need to be dredged, how high the floodwalls would need to be built, etc.), and they select and build their defences. Then when the rains come, they measure the amount of rain that falls and the effectiveness of the defences they have built, and they adjust their defences if needed.

Security risk is similar, and it needs to be analysed and managed in a similar sort of a way. We need to build 'circulation and flow models' for security risk that will let us calculate probabilities and rates for the different outcomes that might arise. We need to measure the security equivalent to 'the rainfall patterns in the vicinity of interest', and the relevant characteristics of the target we want to protect, and then deduce the security risk given those measurements and characteristics.

The paradigm behind TBSE is the security risk equivalent to thinking in terms of circulation and flow. It is a non-deterministic paradigm, and the analytical methods that go with it are the non-deterministic methods for building and analysing the flow models we need. For a security analyst to be able to quantify security risk, they need to learn how to think about security interactions non-deterministically. The TBSE paradigm is the way to do that.