Some people looking at TBSE for the first time have commented that it reminds them of the Lockheed Martin Cyber Kill Chain (CKC). I have been remiss at not paying particular attention to the CKC before, and these comments have given me cause to take a look. From my reading of Lockheed Martin's website and the various documents provided there, the CKC looks to me to be a framework for organising one's defences but not for quantifying one's risk. Any similarity between TBSE's Threat Pathway and the CKC 7-step attack chain is only superficial, and TBSE provides the defender with a range of capabilities that the CKC does not.

In my view:

- The CKC is a framework for addressing cyber intrusions by intelligent and adaptable adversaries. But it is not limited to that, and it can be generalised easily enough to cover other threats, including those not driven by intelligent and adaptable adversaries.
- It is a descriptive framework for identifying ways to intercede in the progress of a particular threat. It encourages defenders to address each threat they are concerned about at as many of the 7 steps as they can, under the familiar rubric of 'Defence in Depth'. It also promotes the intuitively sensible idea that the more one does to intervene in the progress of the threat early in the chain, the better leverage one is likely to get on the end objective: minimising the number of attacks that achieve success.
- The framework is limited in its scope to the technical end of the threat chain. It covers the upstream steps by which threat agents generate attacks and attacks are successful at creating security breaches. It does not address what happens downstream of that. It says nothing about the steps by which security breaches generate operational effects or operational effects harm the business or mission. Along with some other methodologies, it covers only the easy half of the story. It stays within the technologist's comfort zone, addressing the narrow perspective of stopping threats creating security breaches, not the fuller business-centric perspective of protecting a business or enterprise from harm. It doesn't help bridge the divide between the IT Security function focussed on technology events and business leaders (the customers of the IT Security function) who are focussed on protecting the business at a reasonable $$ cost.
- The framework is purely descriptive. It does not address how to analyse the effectiveness of the controls one might use, other than by suggesting defenders collect data over many incidents and identify which controls seemed to be effective in the defeat of different attacks. It says nothing about how to quantify the effectiveness of the controls, how to measure controls' effectiveness, or how to improve the effectiveness of any control found to be deficient or critical. It says nothing about how to relate the effectiveness of a control to the way the control is configured and operated, and it says nothing about how to optimise a set of controls to obtain the best overall effect for the minimum financial or operational cost other than by 'trial and error'.
- The literature that comes with the CKC redresses the often perceived imbalance between attacker and defender. For a while the balance has been perceived by some to be in the adversary's favour: the defender has to win 100% of the time, the attacker has to win only once. The CKC proposes a different perspective: the attacker has to be successful across all seven steps in the chain, the defender can break that chain at any one step and has multiple opportunities to do that. This is no doubt attractive from a sales perspective.

TBSE is quite different in concept and purpose from the CKC and provides capabilities that the CKC does not.

- TBSE is not a product or a 'one-size-fits-all' technical solution. It is a paradigm for understanding the dynamics that give rise to security risk, and a methodology that enables those dynamics to be modelled, measured and quantified, and a wide range of risk-relevant insights to be calculated.
- TBSE models the dynamics that underpin security risk stochastically. A stochastic problem is one in which the dynamics that lead to the end result (in this case the dynamics by which security risk is generated by a threat) are driven by processes that have a significant random component to them. Usually this randomness arises because one cannot know everything one would need to know about the specific details of the entities that are interacting via those dynamics (in this case attacks, vulnerabilities and controls). As a result, the dynamics cannot be analysed deterministically. The InfoSec industry has tried for many years to quantify security risk using deterministic methods and it has failed every time. Quantifying security risk is not a problem that can be solved simply by throwing more processing power at it. It is a problem that is not amenable to deterministic methods. It requires a different mode of thinking and a different type of analysis: stochastic analysis.
- As with the CKC, TBSE can be applied to any security threat, not just cyber security threats, and any relevant controls. A major distinction is that TBSE addresses the whole cyber threat chain, from soup to nuts, from threat agents to the business harms that result. It covers the upstream half of the chain, from threat agents to the security breaches they create, but continues within the same holistic paradigm to cover the downstream half of the chain, the operational and business harms that result from security breaches. It is, after all, these operational and business harms that are the primary concern of business leaders. These are the things the business is trying to protect itself against. They can't be ignored.
- TBSE is a fully flexible analytical approach, not just a descriptive framework. It enables the analyst to model and analyse any particular interaction within the threat pathway and to generate whatever risk results might be relevant at the time. How the analyst applies TBSE, the models they build and the calculations they perform, will depend on the particulars of the situation they are analysing. The results they get will be specific to the situation, reflect the risk arising from the specific threat being experienced at the time, the particular vulnerabilities in place, and the particular settings they have for their security controls.
- TBSE enables an analyst to calculate in absolute terms, not just in relative H/M/L terms or on a scale of 1–10, the effectiveness of a control. Effectiveness is calculated as a function of the way the control is implemented and operated, and as a function of the profile of the particular threat the control is protecting against. It enables the analyst to calculate how that control's effectiveness would change if the profile of the threat were to change or if the control were to be implemented or operated differently.
- TBSE does not produce a single overall answer for what the risk might be based on any single piece of analysis. The threat pathway for any threat will have many steps within it, and TBSE is applied to each step of interest. If an analyst wants an overall answer for the risk that survives at the end of a threat's pathway, then they will need to conduct a sequence of analyses covering the entire pathway for that threat. That would be an extensive piece of work and is not how TBSE would normally be used. Instead, TBSE will be used to do particular individual pieces of risk analysis to generate particular individual results. For example, to generate risk metrics that will enable the strengths and weaknesses of individual controls to be understood much better, using those insights to inform whatever risk management decisions need to be made. Or to compare the different RoIs for different security improvement strategies. Or to rebalance the effort put into each of a group of controls relating to one or more Top Threats, to improve the cost-effectiveness of a security strategy.

For an introduction to TBSE, what it is and what it can do, see www.jlis.co.uk.

