




BANK OF ENGLAND



CBEST Intelligence-Led Testing

An Introduction to Cyber Threat Modelling


Version 2.0



Copyright notice

© 2016 Bank of England

This work is licensed under the Creative Commons Attribution 4.0 International Licence. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, 444 Castro Street, Suite 900, Mountain View, California, 94041, USA.



You are free to:

- Share — copy and redistribute the material in any medium or format.
- Adapt — remix, transform and build upon the material for any purpose, even commercially.

The licensor cannot revoke these freedoms as long as you follow the licence terms.

Under the following terms:

- Attribution — you must give appropriate credit, provide a link to the licence, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.
- No additional restrictions — you may not apply legal terms or technological measures that legally restrict others from doing anything the licence permits.

Notices:

- You do not have to comply with the licence for elements of the material in the public domain or where your use is permitted by an applicable exception or limitation.
- No warranties are given. The licence may not give you all of the permissions necessary for your intended use. For example, other rights such as publicity, privacy or moral rights may limit how you use the material.

Contents

Executive summary	3
<hr/>	
1 Introduction	4
1.1 Purpose of this document	4
1.2 Terms of reference	4
1.3 Information sources	4
1.4 Structure of this document	5
1.5 Legal disclaimer	5
<hr/>	
2 Threat modelling overview	6
2.1 Introduction	6
2.2 Background	6
2.3 Requirements of the CBEST threat model	6
2.4 Overview of the CBEST threat model	7
2.5 Relation to existing approaches	7
<hr/>	
3 Threat entity goal orientation	9
3.1 Introduction	9
3.2 Identity and interests	10
3.3 Motivations	11
3.4 Intentions	11
<hr/>	
4 Threat entity capabilities	13
4.1 Introduction	13
4.2 Resources	13
4.3 Skill, prowess and maturity	14
4.4 Resolve	14
4.5 Access to target	15
4.6 Risk sensitivity	15
<hr/>	
5 Threat entity modus operandi	16
5.1 Introduction	16
5.2 Operational stages	16
5.3 Method and activity profiles	17
5.4 Threat artefacts	18
5.5 Transmission media	18
<hr/>	
6 Applying the model	19
6.1 Introduction	19
6.2 Populating the model	19
6.3 Threat assessments	19
6.4 Threat scenarios	20
6.5 Conclusions	21
<hr/>	
7 References	23

Executive summary

This document defines an analytical model of cyber threat intelligence in terms of a threat entity's goal orientation, the capabilities it uses to pursue its goals and its modus operandi. The model is intended to act as a common guiding template for conducting a cyber threat assessment for use by penetration testers to define a set of realistic and threat-informed cyber attack test scenarios.

For the purposes of this document a threat is defined as a harmful outcome resulting from an entity's actions in pursuit of its goals. Given this definition, the CBEST threat model defines a threat entity's potential in terms of:

- the threat entity's **goal orientation**, which explains the existence, and potential seriousness of a threat entity;
- the **capabilities** the threat entity uses to pursue its goals, which explains the threat potential exhibited by a threat entity;
- the threat entity's **modus operandi**, which explains how malicious activity unfolds and what malicious code artefacts are created.

None of the concepts involved in cyber threat modelling are new; all of the building blocks can be found in open sources, many of them academic or governmental in origin. However, the novel aspect of the threat model described in this document lies in the specific mixture of concepts and categories and how they have been organised to achieve a tailored solution to the CBEST programme's unique demands.

CBEST stakeholders will use the threat model as follows:

- **threat intelligence analysts** will use the model to supplement their existing methods of cyber threat assessment, collecting and analysing new information where necessary and fusing the results into a coherent whole;
- **penetration testers** will use the model developed by the assessors to define a set of realistic and threat-informed cyber attack test scenarios;
- **regulators** will use the model as a guide to conducting cyber threat assessments within individual financial institutions;
- **Firms/FMIs (Financial Market Infrastructure)** will use the model to produce threat assessments that satisfy the requirements of the CBEST programme, are cost-effective and provide outputs that are their security teams find useful.

As well as defining the CBEST threat model, this document also presents some initial remarks, in advance of a formal use case specification, regarding its application.

As CBEST threat assessment activities proceed they will reveal insights about the methodological and conceptual value of the threat model presented in this report. This practical application of the model will provide valuable feedback by which practitioners may further improve it.

1 Introduction

1.1 Purpose of this document

This document defines an analytical model of cyber threat intelligence in terms of a threat entity's goal orientation, the capabilities it uses to pursue its goals and its modus operandi. The model will act as a common guiding template for conducting a cyber threat assessment that will be used by penetration testers to define a set of realistic and threat-informed cyber attack test scenarios.

1.2 Terms of reference

As already discussed in the sister report, *Understanding Cyber Threat Intelligence Operations* (CBEST (2016)), the Bank of England is developing CBEST. This is a framework for developing intelligence-led cyber threat vulnerability tests against financial institutions' critical systems. These tests mimic the actions of groups and individuals who are perceived by Government and commercial threat intelligence providers as posing a genuine threat to systemically-important financial institutions within the Critical National Infrastructure.

CBEST encompasses the following groups of stakeholders:

- **regulators:** governmental authorities bearing official responsibility for the stability of the UK financial system;
- **Firms/FMIs (Financial Market Infrastructure):** financial-sector organisations recognised by CBEST as critical to systemic stability;
- **service providers:** CBEST-accredited service providers who provide threat intelligence and penetration testing services.

Under CBEST these stakeholders will work together to assess cyber threats that pertain to the resiliency of the UK financial system as a whole rather than any single organisation. A threat model is needed in order to determine what threats might be of relevance to systemic resiliency, what malicious entities are behind these threats, why some threats are more important than others and how to assess defensive and responsive postures relative to these threats. Without a common model shared across the programme's stakeholders it will be difficult to combine the results to form an overall picture of threats at a systemic level.

But the model must also be practically useful. The basic characteristics required of the CBEST threat model are therefore analytical utility, consistency and versatility. The model is intended to facilitate analysis, not to constrain it. It should not be a rigid template that cyber threat intelligence providers simply follow when performing their assessments. Rather, it should provide a common interface through which all stakeholders can integrate and compare the results of diverse assessments undertaken by intelligence providers.

The above stakeholders will therefore use the threat model defined in this document as follows:

- **threat intelligence service providers** will use the model to supplement their existing methods of cyber threat assessment, collecting and analysing new information where necessary and fusing the results into a coherent whole;
- **penetration testing service providers** will use the model developed by the assessors to define a set of realistic and threat-informed cyber attack test scenarios;
- **regulators** will use the model as a guide to conducting cyber threat assessments within individual financial institutions;
- **Firms/FMIs** will use the model to produce threat assessments that satisfy the requirements of the CBEST programme, are cost-effective and provide outputs that are their security teams find useful.

1.3 Information sources

Information for this report was gathered from online open sources and discussions with industry professionals. A full set of references appears at the end of this document. Information was also derived from various CBEST meetings and workshops attended by the representatives of the Bank of England, CREST and the Cyber Working Group during the first quarter of 2014. In 2015 the Bank of England Cyber Sector Team commissioned a review and update of this document during which various stakeholders were canvassed for their input.

1.4 Structure of this document

The remainder of this document is structured as follows:

- **Section 2**, *Threat modelling overview*, describes the background to threat modelling, the requirements of the CBEST threat model, an overview of the model and an explanation of how it relates to existing approaches;
- **Section 3**, *Threat entity goal orientation*, presents the threat entity goal orientation component of the CBEST threat model which explains the existence, and potential seriousness, of a threat entity;
- **Section 4**, *Threat entity capabilities*, presents the threat entity capability component of the CBEST threat model which explains the threat potential exhibited by a threat entity;
- **Section 5**, *Threat entity modus operandi*, presents the threat entity modus operandi component of the CBEST threat model which explains how malicious activity unfolds and what malicious code artefacts are created;
- **Section 6**, *Applying the model*, presents some initial remarks, in advance of a formal use case specification, regarding the application of the CBEST threat model;
- **Section 7**, *References*, lists sources of information used in the production of this document.

1.5 Legal disclaimer

The information and opinions expressed in this document are for information purposes only. They are not intended to constitute legal or other professional advice, and should not be relied on or treated as a substitute for specific advice relevant to particular circumstances. The sponsors and authors of this document shall accept no responsibility for any errors, omissions or misleading statements in this document, or for any loss that may arise from reliance on the information and opinions expressed within it.

2 Threat modelling overview

2.1 Introduction

This section describes the background to threat modelling, the requirements of the CBEST threat model, an overview of the model and an explanation of how it relates to existing approaches.

2.2 Background

All threat assessments, however simple or complex, are based upon an underlying structured, rational model that guides assessors as they distinguish and compare their subjects. Assessors will often attribute their success to 'gut feel' and other intuitive abilities. Articulating these 'buried' implicit models in the form of an explicit conceptual model makes them transparent and subject to critique and improvement.

Cyber threat modelling (Goldsmith and Siegel (2012)) is a specialisation of a more general modelling method that is known, in various fields such as psychology and social science, as 'strategic interaction analysis', 'adversary modelling', 'conflict analysis', 'force modelling' or 'actor profiling' (OMG (2014)). They are all concerned with understanding the choices and behaviours of agents with conflicting interests. All such approaches are ultimately rooted in socio-psychological theories of goal formation, behavioural competence and strategic learning in organisations. This was originally termed 'expected utility theory'.

As a result, none of the concepts involved in cyber threat modelling are new; all of the building blocks can be found in open sources, many of them academic or governmental in origin. However, the novel aspect of the threat model described in this document lies in the specific mixture of concepts and categories and how they have been organised to achieve a tailored solution to the CBEST programme's unique demands.

Despite their advantages, models are not sacred. The sole criterion for judging a model should be how useful it is to the analyst. The CBEST threat model, like all other models, is therefore best viewed as a tool, or a toolbox. As mentioned in the CBEST *Understanding Cyber Threat Intelligence Operations* report, all models are approximations of reality and should be treated as such.

It is also important to remember that accurate forecasting of the behaviours of complex organisations is one of the most difficult kinds of analysis. Even with the most promising models, it is not always possible to collect data of sufficient quality to test them. Even when the data is available such models can only, at best, predict only the probability distributions of different outcomes. This explains the lack of universally accepted models for predicting complex threats. That said, it is also clear that much harmful behaviour is, to some degree, explainable and therefore, to some degree, predictable.

Because of this, responsible threat assessment, indeed any type of behavioural modelling, involves a mix of art and science, continuous reflective scrutiny, experimentation and adaptive refinement. These considerations have influenced the development of the CBEST threat modelling framework presented in this document.

2.3 Requirements of the CBEST threat model

The primary users of the CBEST threat model will be:

- **assessors** who will use it as a template for cyber threat assessments;
- **testers** who will use developed models for defining cyber attack test scenarios.

In order to be of practical use to its users, the threat model must provide a systematic means of enumerating and organising the key components of a cyber threat scenario. The model must be precise and rigorous enough to ensure the kind of accuracy expected of a professional intelligence analysis. At the same time it needs to be clear, simple and flexible enough so that analysts with different backgrounds and experience levels can use it (subject to reaching CBEST approved status).

Each CBEST threat assessment will be different due to the variety of Firms/FMIs involved in the assessments and the range of threats each one faces. The threat model must therefore be able to accommodate significant variations in the quality of the input

data relating to a threat entity and its activities. Key quality criteria include scope, type (qualitative or quantitative), volume, structure, detail and accuracy.

Depending on the quality of the input data, threat assessments developed using the model will range from basic but nonetheless valuable profiling of an incident, the threat entity and its TTPs, (or, as described in the CBEST *Understanding Cyber Threat Intelligence Operations* report, reporting *what* happened and *why* it happened) through to more sophisticated predictive analysis (ie what *will* happen). Adopting a consistent modelling approach that can accommodate variations in input data will enable regulators to make consistent and straightforward comparisons across diverse threat assessments.

2.4 Overview of the CBEST threat model

The CBEST *Understanding Cyber Threat Intelligence Operations* report defines a threat as follows:

Threat

- an expression of intent to do harm, ie deprive, weaken, damage or destroy;
- an indication of imminent harm;
- an agent that is regarded as harmful;
- a harmful agent's actions comprising of tactics, techniques, and procedures (TTPs).

For the purposes of this document a threat is viewed as a harmful outcome resulting from an entity's actions in pursuit of its goals. Given this viewpoint, the CBEST threat model defines a threat entity's potential in terms of:

- the threat entity's **goal orientation**;
- the **capabilities** the threat entity uses to pursue its goals;
- the threat entity's **modus operandi** (including TTPs).

Modelling the threat entity's goal orientation and capabilities enables us to understand which threats may occur, what entities will cause them and why some threats are more important than others. Modelling the threat entity's modus operandi enables us to understand how any malicious activity will unfold and any characteristic tactics or malicious code artefacts of which defenders should be aware.

Each of the three strands of the CBEST threat model listed above decomposes into two levels of granularity. Each level can be expanded into further levels of detail and can ultimately be expressed as quantitative measurements or qualitative, concrete facts. The level of detail will depend on the needs of the modellers and the availability of intelligence; in all cases, qualitative assessments at higher levels of generality may be derived in the absence of concrete, detailed measurements.

Regarding 'intelligence', the term is already defined in the CBEST *Understanding Cyber Threat Intelligence Operations* report but it is worth repeating here for the sake of clarity. The definition is as follows:

Intelligence

Information about threats and threat actors that provides sufficient understanding for mitigating a harmful event.

Intelligence is therefore a particular kind of information, where information is data in context, or a higher-level abstraction or viewpoint made on the basis of one or more elementary data items.

2.5 Relation to existing approaches

The creators of the CBEST threat model have reviewed over fifty academic, military, governmental and industry publications (see References) spanning the past 30 years to devise a configuration that is conceptually sound, analytically useful and capable of integrating outputs found across the widest range of other approaches.

Of particular relevance are sources pertaining to previous efforts by the US Department of Defense, the US Department of Energy, NATO, the EU's ENISA and the RAND Corporation. The most mature examples of intelligence-led cyber security assessments are also the earliest precedents for CBEST, namely the US DoD's internal protocol for information assurance, DoD requirements for contractors participating in the DIB CS/IA (Defense Industrial Base Cyber Security/Information Assurance) programme and DHS's Enhanced Cybersecurity Framework. Since at least 2010 these initiatives all began to apply current intelligence about likely adversaries into the 'playbooks' of organisations tasked with emulating attackers, often known as 'red teams' (Parks (2010); IDART (2009); Department of Defense Science Board (2003)).

The CBEST threat modelling approach neither conflicts with nor competes against any of the numerous 'threat intelligence frameworks' currently under development. These are used to construct machine-readable data feeds that pipe threat intelligence into Security Information and Event Management (SIEM), anti-virus software, firewalls, intrusion prevention systems (IPS) and intrusion detection systems (IDS). They are based on a proprietary or open standard and each is promoted by an industry patron. Examples include OpenIOC (Mandiant), STIX (Mitre), CIF (REN-ISAC), IODEF (Internet Engineering Task Force) and VERIS (Verizon). They are discussed in further detail in the CBEST *Understanding Cyber Threat Intelligence Operations* report.

The CBEST threat model accommodates any kind of data input whether or not it has been pre-structured by any of the above threat intelligence frameworks since it is organised at a higher level of abstraction. Categories in the CBEST model can therefore either map directly to equivalents in other models or subsume them.

More fundamentally, the CBEST model is an analytical model, whereas the threat intelligence frameworks are simply conventions for labelling data for easier automated dissemination and processing. Although they are important, such frameworks are not formal analysis models that analysts can use to explain patterns, trace causality or interpret significance.

The References section at the end of this document contains a full set of references that were reviewed when developing the CBEST threat model.

3 Threat entity goal orientation

3.1 Introduction

This section presents the threat entity goal orientation component of the CBEST threat model that explains the existence, and potential seriousness, of a threat entity.

The extent to which an entity constitutes a threat depends upon the potential harm entailed in the outcomes resulting from the entity's pursuit of its goals. Modelling the relationship of threatening outcomes to goal pursuit requires understanding how entities judge the value of outcomes, the pressures or incentives that drive them and how they intend to achieve their goals. These aspects are collectively defined as a threat entity's goal orientation, which is modelled as a function of interests and values, motivations and intentions as outlined below.

Identity and interests

A threat entity's identity and corresponding interests determine the range of all potential future outcomes that the entity could expect. This is on the basis of interaction, threatening or otherwise, with other entities (MIT (2008)). Identity and interests help reveal the possibilities allowed by an entity's internal makeup, the relationship to its strategic environment, its range of available options for acting and its perceptions of relevance and value. It is always through this lens that threat entities encounter motivations and react by forming goals, intentions and plans to pursue them.

Because they are so fundamental, factors of identity and interests inform a threat model in ways that are subtle and complex but nevertheless pervasive. Moreover, these factors tend to remain stable over time and are ascertainable through many open sources. Intelligence analysts with experience researching particular entities will be able to provide inputs for these factors without extensive requirements for new collection and analysis.

A threat entity's identity is a function of its internal composition, how it functions, its basic outlook and its position relative to other entities within the wider strategic environment. For example, authoritarian regimes such as North Korea tend to perceive the strategic environment differently and make decisions that reflect only the ruling cadre's interests, as opposed to states founded on more open and democratic principles of political organisation. States that are ensconced deeply in alliance institutions or economically oligopolistic organisations, such as NATO and OPEC, respectively, make decisions that account for the interests of a larger number of external parties.

Along yet another dimension, some political entities are institutionally and behaviourally bound to maintaining a particular political or security status quo, while other entities, often called 'revisionists', are constitutionally opposed to the status quo. For example, Russia is strongly opposed to the continued pre-eminence of NATO, which is the status quo of the European security environment, while the United States and the EU member states almost all strongly support the status quo. Thus, the ways that an entity may devise threatening behaviour will often strongly reflect the presence of the structural, positional and perspectival factors classified below under the 'Identity and interests' heading.

A threat entity's interests consist of the range of potential outcomes and impacts that the entity perceives as most important in terms of its future viability and wellbeing. They are the foundation for understanding whether, and to what extent, an entity may pose a threat. Interests refer to potential future events or contexts that are of sufficient relevance to an entity to warrant the risks and costs associated with that entity threatening others. Values are the principles that specify the ordering of entities' preferences. Together, these determine the reasons why an entity would engage in threatening activity and, just as importantly, why an entity would not.

For example, the United States, Israel and many other nations have a vested interest in preventing Iran from obtaining a nuclear capability but not in seeing a full blown civil war engulf the country which would destabilise the entire region. As such, any efforts to curtail Iran's nuclear programme will tend to be circumscribed by the countervailing interest to maintain basic political stability for Iran. Moreover, the United States and EU states have a stronger interest in a stable Iran than does Israel, implying that Israel's strategic behaviour regarding Iran is less restrained from reaching higher levels of overall political disruptiveness than comparable

efforts by other interested states. Such an analysis could help explain why the Stuxnet (aka 'Olympic Games') campaign targeted only the nuclear enrichment capabilities of Iran rather than, say, the command and control infrastructure of the Iranian Revolutionary Guard Corps.

The list of factors or attributes that could inform a model of threat entities' identities and interests is vast. Those enumerated below reflect the most successful efforts from the fields of strategic studies, international relations and practical intelligence analysis for conflict modelling.

Motivations

Motivations are factors that capture how the expected impacts of potential outcomes translate into the incentives and inhibitions that guide a threat entity's choices and behaviours, including those detrimental to others' interests. Motivations therefore provide the direction, urgency, flexibility and prioritisation criteria that inform how an entity determines its goals, level of effort and orientation towards other entities.

For example, given a fixed configuration of identities and interests, motivations for intensified threat activity often reflect relatively sudden alterations in a strategic scenario that may embolden or inhibit the seriousness of a threat actor's pursuit of its goals. For example, the United States had no interest in placing troops in Afghanistan until 11 September 2001. In a longer-term example, the militaries of most Western European nations could no longer justify massive investments in standing armies after the fall of the Soviet Union. Adjacent to the security realm, we can find many examples of economic motivations competing with security priorities, such as the prospect of capitalising on foreign demand for weapon systems, although doing so may imperil a country's military personnel years in the future.

Intentions

Intentions denote what an entity wants to achieve and how it plans to act in order to do so. Intentions are especially sensitive to context, which is the basis for most of the categories in this section of the threat model.

Some notable examples of intentions include the formal pronouncements for creating military cyber commands by various nations or resolutions to increase the cyber security co-operation ties between countries. Less formally, we can analyse long-term trends in cyber infiltration activity to discern with great confidence that some nations systematically encourage, and intend to persist in encouraging, cyber espionage by private sector or even rogue entities, although no such intention is ever formally stated in government documentation.

The above three components of the threat entity goal orientation model are described below.

3.2 Identity and interests

Identity

- Entity Name
- Entity Type (eg nation-state, state proxy, hacktivist, organised crime)
- Entity Size and Complexity
- Entity's Internal Power/Influence Structures
- Entity's Developmental Tendencies (eg erratic, stable, sustainable)
- Entity's Perceptions of Itself and the Strategic Environment
- Roles and Relationships to Other Relevant Entities

Interests

- Existential (a requirement which, if not met, would result in the entity's demise)
 - Inflows of Key Resources
 - Internal Cohesion/Stability of Control
 - Financial Viability
 - Key Relationships With Other Entities
- Security/Conflict Interests
 - Number of Actual or High Likelihood Conflicts
 - Severity
 - Internal — External Dimension
 - Security Relationships with other entities (dependencies, allies, rivals, enemies)

- Security Problems (issues or trends)
- Economic/Financial
 - Economic Governance Profile
 - Trade Relationships
 - Economic Development Profile
 - Technical Sophistication and Innovation
 - Financial Factors
 - Key Dependencies/Potential Substituting Options
- Socio-Cultural
 - Religious
 - Ethnic
 - Historical Communities of Interest
 - Normative Orientation (rule of law, reciprocity, democratic, liberalism, etc)
 - Cosmopolitan or Technocratic Elites
 - Functional Orientations (eg hi-tech, media-centric, criminalised)

3.3 Motivations

Types of motivation

- Habitual (eg bureaucratic inertia, continuity with previous practices)
- Perceived Threats (harm, punishments or losses)
- Perceived Opportunities (advantages or gains)
- Systematic Errors, Oversights, Biases, and Misperceptions
- Susceptibility to Outside Influences
- Normative Parameters

Desired types of value

- Politico-strategic
 - Security, Power, Dominance, Hegemony
- Status
 - Autonomy, Acceptance, Belonging, Influence
- Wealth
 - Development — Stable Growth — OECD Status — Sustained Enrichment
- Orientation towards system in which entity is a key player
 - Status Quo, Gradual Revisionist, Radical Revisionist

Threat entity's relationship to targeted entities

- Type and Attributes of Relationship
 - Valence (cooperative or conflicting)
 - Intensity of Interaction (strong or weak)
 - Consistency (enduring/persistent or prone to shifts)
 - Complexity (single-issue or multiple and interdependent)
- Compatibility of Interests
 - Mutually Exclusive (zero-sum)
 - Variable
- Crisis Potential

3.4 Intentions

- Specific Intentions
 - Overtly Stated Plans
 - Known or Inferred (including confidence level)
- Preparedness (capabilities already allocated and mobilised)
 - Means
 - Funds
 - Time
 - Dependencies on Other Entities

- Characteristics Inherent Across Set of Intentions
 - Number of Outcomes Sought
 - Explicitness/Formality of Stated Intentions
 - Coherence/Consistency Across Goal Set
 - Rigidity or Malleability of Goal Set
 - Time Orientation and Term of Goals (long or short-term)
- Specificity of Goals Relative to the Target
 - Dependencies Beyond Entity's Control
 - Inclusivity or Exclusivity of Expected Benefits (ie how zero-sum are the expected outcomes of goal attainment?)
 - Availability of Other Options

4 Threat entity capabilities

4.1 Introduction

This section presents the threat entity capability component of the CBEST threat model which explains the threat potential exhibited by a threat entity.

Each of the attributes listed below can, to some extent, substitute for one another assuming that none of them are zero or close to zero. For example, operational effectiveness can result from high prowess with sparse resources (as long as the resources are not nil). Similarly, a combination of high resolve and risk acceptance can work to offset relatively lower prowess or resources.

This is important because different attribute mixes will in turn produce different threat scenarios derived from the threat model. They will also have different implications for the relative efficacy of deployed countermeasures.

There are five main components of the threat entity capability model, namely:

- Resources;
- Skill/Prowess/Maturity;
- Resolve;
- Access to Target;
- Risk Sensitivity.

These are described below.

4.2 Resources

The resources available to an entity determine the scope, intensity, sustainability and diversity of the total set of actions that entity can take. Resources are, in a sense, the raw materials and fuel that can be applied or expended. In terms of the metaphor of a one-on-one fight, resources would correspond to an opponent's size and strength.

Any organisation's resources comprise, at a minimum, a mixture of people, technology and finances. Importantly, any assessment of resources' relevance to overall capability should reflect not merely the addition of the volume of each resource type but also the balance among them and how they are organised. Balance and organisation are important because they together determine what resources are actually available and useful rather than those simply possessed.

Human

- Number of Personnel
 - Leadership
 - Specialists
 - Non-Specialist
- Personnel Quality
 - Education and Training
 - Experience
- Prevalence of Key Innovators ('game changing' talent)
- Prevalence of Key Connectors (having many trans-organisational relationships)
- Balance of Composition ('teeth to tail' ratio, etc)

Materiel/technical

- Physical
 - Equipment/Systems
 - Hardware Infrastructure
 - Facilities

- Informational/Virtual
 - Malicious Codebase
 - Vulnerability and Exploit Research
 - Software, Tools, Kits and Scripts
 - Data/Intelligence Reservoirs
 - Virtual Infrastructure (eg domains, email addresses)
- External-Interactive
 - Access to Assets of Affiliated Entities
 - Access to Criminal Services (eg money laundering)
 - Access to 'Dark' Fora

Budget/financing

- Budget/Revenue
- Credit
- Capital/Invested Assets

4.3 Skill, prowess and maturity

This sub-category consists of factors that are less tangible but no less crucial than those found in the Resources sub-category above. If the Resource sub-category relates to an opponent's size and strength then Skill, Prowess and Maturity relates to an opponent's cunning and experience. This sub-category reflects the extent to which a threat entity operates with awareness, intelligence, strategic acumen, learning potential, creative problem-solving, decision-making coherence and operational experience. This will depend on the threat entity's composition, organisational features and its relationship with other, allied entities.

Organisational performance sophistication

- Properties of Structural Complexity (eg clustering coefficients, eigenvector centrality)
- Properties of Dynamic Complexity
 - Communicative Patterns
 - Functional Differentiation Paired With Holistic Integration Over Time
- Any Other Evidence Of Organisational Learning Over Time

Operational and tactical ingenuity

- Novel Activities Per Operational Undertaking Per Unit of Time
- Ratio of Spending on R&D to Total Budget
- Ratio of Person-Hours Spent on R&D to Total
- Indicators of Innovativeness in Performance

Authority

- Formal Authorities Listed in Open Documentation
- Informal Authorities Wielded in Practice
- Access to Sensitive or Secret Information Based on Authority
- Impunity or Insularity From Negative Consequences

Trans-organisational integration

- Integration Relative to Peer Organisations (for state actors, this would be 'whole of government' integration, for example)
- Cross-domain Integration (for state actors an example would be PPPs)
- Integration with Allies/proxies (for nation-state or state proxies only)

4.4 Resolve

This sub-category consists of factors that determine how assiduously or aggressively a threat entity will incite and sustain malicious actions against a target. This is, in effect, a measure of the willingness, cost-tolerance and determination that a threat entity can muster and hold, especially when confronted with a target's efforts to detect, deter or defend against malicious activity. It is necessary to separate this sub-category from the others because it reflects attributes, influences and causal pathways not reducible to other factors. Some of these include intangible yet demonstrably real features of an entity's strategic

culture, the expected utility calculus in different situations and the socio-psychological aspects of hostile strategic interaction over time.

- Capacity to Focus Malicious Action (ie intensity of the pursuit of a strategic goal, relative to ancillary demands or influences)
- Persistence in Response to Failure and Setbacks
- Persistence in Response to Detection/Attribution
- Strategic or Operational Autonomy

4.5 Access to target

This sub-category provides a means of measuring the previous successes and the additional, complimentary support or augmentation capabilities that some threat entities may marshal in addition to their cyber espionage and cyber conflict efforts. That is, the factors below reflect the second-order paths by which threat entities may exploit a targeted organisation's systems.

- Extent of Previous Success in Infiltrating Target (entails familiarity with targets and ready-to-hand exploitability)
- Extent of Access to Supply Chain and Adjacent Targets
- Privileged Positioning (ie relationships between target and threat entity or its proxies)
- Presence and Influence of Threat Enablers (ie high-risk insiders, etc)

4.6 Risk sensitivity

Risk Sensitivity is similar to Resolve yet remains distinct enough to warrant its own discrete sub-category. It refers to how much potential danger or harm a threat entity will risk facing. Resolve, by contrast, refers to how much danger or harm the threat entity can incur while still maintaining its hostile activity. Because Risk Sensitivity ultimately reflects the aggregation of complex socio-psychological processes involving learning and error-correction, there are countless ways to model it. The most commonly applicable factors are listed below.

- Factors of Risk Acceptance
 - Responses to Detection
 - Stealth Use vs Target Status and Defensive Ability
- Factors of Risk Aversion
 - Time Delay After Unsuccessful Operation
 - Enhanced Security Around Other Operation When One Is Detected
 - Withdrawal Of Activity After Detection

5 Threat entity modus operandi

5.1 Introduction

This section presents the threat entity modus operandi component of the CBEST threat model that explains how malicious activity unfolds and what malicious code artefacts are created.

Practically every potential threat entity with any interest in disrupting the UK financial system will be an organisation, most commonly a nation-state unit or hacktivist group. Potentially malicious entities pose threats by virtue of their goal orientations and capabilities but these do not result in harm until the entity takes action. As well as direct harm on the target, disruption (caused by cyber criminals, for example) can also be inflicted on organisations of all sizes beyond the immediate target.

Having assessed a set of entities that pose potential threats and the relative severity across that set, analysts can form an understanding of how their threat behaviours will manifest themselves as identifiable actions and malicious code artefacts. A significant portion of a threat entity's actions will reflect the penetration testing techniques used by CBEST penetration testers. However, it is also necessary to model threatening actions at the functional level, ie at the level of human and organisational (rather than technical) performance.

Each threat entity will exhibit a particular mix of technical and behavioural characteristics. At the most general level, a threat entity's modus operandi include the following:

- Operational Stages;
- Method and Activity Profiles (including TTPs);
- Threat Artefacts;
- Transmission Media.

These are described below.

5.2 Operational stages

Operational Stages reflects two organising principles: temporal and functional. Because these two principles are conceptually distinct but practically interrelated this often proves to be a source of confusion for security professionals. More importantly, conflating the temporal and functional principles is a persistent conceptual weakness in many malicious activity models that are based on these twin principles and fail to recognise or acknowledge that feature. This is one of the fundamental flaws with the popular 'kill chain' model popularised by Lockheed-Martin analysts (Hutchins, Cloppert and Amir (2011)). Because of the kill chain's minor systemic inconsistencies, the CBEST operational stage model has been designed to be more conceptually fundamental and consistent, thus allowing the CBEST model to exploit information expressed using the terminology of the kill chain vocabulary or others like it.

A model based on dual temporal-functional dimensions is conceptually sound as long as the dual principles are recognised and accounted for analytically, in other words:

- begin by testing which of the model's variables must occur in sequence rather than parallel;
- assume that all remaining variables are purely functional, ie, no time dependence;
- analyse the degree of temporal dependency by looking for patterns of sequential and parallel occurrences by time interval.

Analysts may then use the data to perform a time dependency analysis using simple event sequencing for qualitative data. More powerfully, they may also use inferential statistics, including simple time series (often used to test for causal correlations that require time for the cause to yield effects) and autocorrelation (correlation of a variable to itself at different times). The goal of such analysis is to assign a time dependency score (or uncertainty score if inconclusive) to any pair of variables exhibiting the relevant sequence. The operational stage variant developed for CBEST testing is set out below.

Pre-attack research, reconnaissance and target selection

- General reconnaissance across potential target population
- Referencing previous successful infiltrations

Planning and preparation of attack components

- Malicious code development
- Vulnerability exploit acquisition
- Threat infrastructure instantiation
- Preparation of delivery vehicle

Infiltration campaign

- First occurrence
- Total count over time

Post-infiltration entrenchment, reinforcement and maintenance of persistence

- Patterns of lateral spreading
- Preferred locations for persistence enhancements (ie backdoors)

Identification, exfiltration or manipulation of data

- Balance of identification, exfiltration vs manipulation
- Types of data as prioritised by attackers
- Methods of searching/indexing
- Means of exfiltration
- Content exfiltrated (relative to searching/indexing)

Exploitation of operational results

- Evidence that exfiltrated data has been exploited
- Time lag from exfiltration to exploitation
- Evidence of access in future operations

5.3 Method and activity profiles

A method or activity profile is a second dimension by which to organise data that can either complement the Operational Stages dimension or, less ideally, serve in place of it if necessary.

Tactics, techniques and procedures

- Research activity
- Pre-attack planning and testing
- Infrastructure setup and maintenance

Stealth and defensive countermeasures

- Stealth
- Non-crypto obfuscation
- Encryption
- Infrastructure hardening
- Infrastructure dynamism

Operational tempo

- Time between attacks
- Time from infiltration to entrenchment
- Time from entrenchment to goal-attainment
- Time to notice success by defence
- Time to launch new attack after success by defence
- Activity types by range of clock time (timezone indications)

Collective vs solo orientation

Note that this refers to the entity's relations to others who might not necessarily be individuals

- Known instances of solo attack (and frequency where possible)
- Known cases of collaborative attacks (and frequency where possible)
- Instances of reciprocal interaction for attack-support (and frequency where possible)
- Instances of exchange-based interaction for attack support (and frequency where possible)
- Instances of solo research (and frequency where possible)
- Instances of research-based interaction (and frequency where possible)
- Instances of non-malicious interaction (and frequency where possible)

Identifying characteristics

- Common mistakes
- Uncommon (ie novel or statistically improbable) combinations of TTPs
- Linguistic indicators
- Signatures of activity (eg patterns or sequences found in traffic data)
- Signatures derived from malicious code or other supporting digital artefacts

5.4 Threat artefacts

Malicious code

- Samples by version #
- Frequency of use
- Use modality (use for attack, sell, reverse to modify or learn)

Supporting code

- Kits
- Stealth tools
- Scripts for analysis of data or traffic

Data/research reservoir

This refers to the information that the threat entity has available. This will not always be possible to gauge in practice but when this is possible it can be valuable.

- Evidence of relevant intelligence obtained and held by a threat entity
- Information sources that threat entity is known to have accessed (malware scanners, staff lists, etc)

Operational command and control infrastructure

- Number of nodes
- Functional differentiation across nodes
- Domains
- IPs
- Registrant info and proxy registration mechanisms
- Botnet structure and indicators of known nodes

Support infrastructure

- Stealth Infrastructure (eg VPNs, TOR)
- Testing platforms (eg VirusTotal)
- Backup infrastructure
- Access to forums and channels

5.5 Transmission media

- Name of Transmission Media
- Type of Transmission Media (email client, IM, website, USB drive, etc)
- Reliance on Ancillary Transmission Media

6 Applying the model

6.1 Introduction

This section presents some initial remarks, in advance of a formal use case specification, regarding the application of the CBEST threat model.

6.2 Populating the model

It is not necessary for an intelligence provider to provide precise inputs to every one of the categories of the CBEST threat model. Instead, each threat assessment should seek to discover and apply as much information as possible, which can then be organised in the categories to maximise convenience and cross-assessment consistency.

Of course, it is relevant to build threat assessments as completely and accurately as possible, but the point still stands that adequate fulfilment of the threat intelligence function for any CBEST exercise will not require full specification of the model in this document.

Value can be derived from only a handful of, or even one, key pieces of reasonably specific intelligence assessed reliably as high confidence. For example, if one obtains a malware sample that shows strong evidence of authorship by a state actor of an important country not normally accused of cyber espionage or CNA, then it can form the basis of an entire scenario, given fewer indicators about the target systems or the attackers' probable goals or interests. Similarly, finding a link between a known but secretive hacktivist group and the covert threat infrastructure they are developing or testing can be enough intelligence to build most scenarios.

Furthermore, a lot of detail on the modus operandi can make for a passable scenario. For example, identifying and analysing a malware sample and associated activity logs from the Firm/FMI client's network often provides enough for a scenario made acceptable but mediocre by reliance on low-confidence guesses of potential actor types and national identities or fictitious 'could have been' attack activity profiles.

Unless a threat intelligence provider is already in possession of extensive profiling information, many threat assessments will be limited by the requirements of the threat scenarios. That is, a typical limit sets in after a threat intelligence provider has specified:

- Goal orientation:
 - the actor's sub-type of identity (nation-state IC or military or LEO; hacktivist issue-driven or general anti-mainstream);
 - many interest and identity fields will follow from knowing an actor's sub-type;
 - 2–3 Motivations;
 - 1–2 Intentions (even inferred intentions are often acceptable, if a good basis exists);
- Capabilities:
 - 2–3 Resource Indicators (one at least should be a 'Virtual' sub-type; none necessary for Budget);
 - 2–3 Skill/Prowess/Maturity Indicators (one should be Skill and Knowledge sub-type, one should be Resolve sub-type);
- Modus Operandi:
 - 2 Actions (either from the Operational Stages model or the Methods and Activities);
 - 1 malware sample (or else two other Artefact types);
 - 3 Artefacts (two should be from either Infrastructure sub-type or each).

6.3 Threat assessments

On the basis of the CBEST threat model, each threat assessment should draw from all available inputs to identify what threatening events are likely, the relative priority of those threats, which entities are responsible for each and the malicious activities that occur as threat entities pursue their objectives. These are the foundational components upon which any realistic threat assessments will be based.

The primary outputs for CBEST Firms/FMIs and regulators will ideally consist of brief reports that enumerate and rank-order the set of potential threats relative to one another at the outset, and which then provide the intelligence findings and analyses which support that enumerated set of threats. The CBEST threat modelling framework provides a useful way of structuring such intelligence for reporting, although different intelligence providers may prefer to provide additional or supporting information in other data formats such as XML or CSV files.

An assessment of each threat will provide explanations of:

- the negative impact on UK financial system resiliency that may result from malicious actions;
- the threat entities harbouring some intention or potential interest in such outcomes;
- why such intentions or potential exist;
- how capable the threat entity is, therefore how likely the outcome will be;
- the sequences of actions and supporting artefacts by which the threats manifest themselves.

Of course, each threat assessment should include requisite discussions of a threat intelligence provider's research methods and confidence levels assigned for each component of their assessment.

6.4 Threat scenarios

The CBEST threat model will act as a common guiding template for conducting a cyber threat assessment that will be used by penetration testers to define a set of realistic and threat-informed cyber attack test scenarios.

These threat scenarios form the link between 'threat intelligence-led' and 'penetration testing-led' regimes. Each threat assessment contains all of the necessary information to inform the development of a threat scenario and ancillary guidance to CBEST penetration testers.

The threat scenarios should be brief stories that narrate the flow of interaction between threat entities and their potential targets, in this case, the participating CBEST financial organisations. Each scenario should consist of the following core features:

- setting: time and place (which will include cyberspace);
- a threat entity;
- optional: specification of identity (or a set sharing characteristics that justify their coexistence in the scenario);
- a threat entity's formulation of its goals, explaining how they denote potential malicious activities;
- the event sequence reflecting threat entities' preparations and conduct of initial infiltrations;
- the event sequence reflecting threat entities' post-infiltration activities;
- discussions of potential impacts resulting from variants of threat entities' successes (this implicitly specifies the ranking of that particular threat).

Each scenario will, thus, contain a high-level summarisation of the more detailed information that comprise the threat assessment outputs described above. CBEST penetration testers will use these scenarios to prioritise their plans for penetrating participants' networks, specifically by allowing them to most easily determine what systems are more important than others, given the objectives and capabilities of actual threat entities. In this way, some penetration testing activities can be safely discarded for some exercises, while others will feature as essential and others still can be listed as optional or conditionally preferable.

For example, a threat scenario may narrate how a major nation state's intelligence services will seek to infiltrate a UK bank's high-frequency trading algorithm platforms in order to cause maximal disruption in the event of a future crisis. The threat entity prefers to coerce the UK government by threatening to cause the disruption without actually having to execute it. Therefore the threat will prioritise multiple, redundant, highly undetectable points of access into the targeted system and the threat entity will seek to maintain such channels of access over time.

This scenario, in turn, specifies that the foreign threat entity's cyber espionage operators will prefer some methods of infiltration and access reinforcement over others. For instance, they may prefer not to exploit existing infiltrations and may target legal staff attached to the high-frequency trading operations rather than the bank's general help desk or customer information database administrators.

Specifying which threat entity is involved will allow threat intelligence providers and penetration testers to identify and prioritise the specific types of malicious code, vulnerabilities and TTPs are most likely to characterise the threat entity's actions. While the penetration testers should by no means be prevented from using their own methods, it will be useful to explore how well a Firm/FMI client's defenders perform relative to threat indicators that most realistically emulate the actual threat entity of concern. Thus, in addition to the threat scenarios, the threat intelligence providers performing a threat assessment should provide lists of indicators of threat entities' TTPs, typical malware and threat infrastructure usage.

In terms of CBEST threat model, including both 'goal orientation' and 'capabilities' will allow for flexibility and redundancy to allow analysts to make threat assessments even when high-confidence indicators are lacking. For example, one threat intelligence provider may have high-confidence indicators, found through intercepting communications, of an entity's intentions to attack part of the UK financial system. Another threat intelligence provider may use different information to analytically infer a similar assessment of the same threat entity's intentions. The model will allow both assessment approaches to fit together.

Similarly, even though one may lack explicit indications of a threat entity's intention to cause systemic disruption in the future, one may safely infer such an intention from a combination of motivations to seek advantage over the United Kingdom plus findings from previous threatening activity against other targets. Furthermore, many pieces of intelligence residing within the model will be the same for any given threat entity, regardless of the potential target in question.

Where possible, the scenario should employ a narrative structure. The main reasons for preferring a narrative format whenever feasible are three (although there are many more):

- narrative employs an intuitive structure which increases engagement and memory retention;
- a story-based structure allows for easier arrangement of fixed information and demarcated spaces for improvisation;
- narrative reflects the 'real' salience of an event, ie it conveys more 'meaning' than 'fact'.

A narrative format consists of, at minimum, the following phases:

- exposition (current equilibrium);
- disruption and rising action:
 - introduction of problem (often as mystery or anomaly);
 - tension builds as problem spreads;
 - tension peaks as imminent catastrophe;
- climax and resolution;
- new equilibrium order and falling action:
 - new synthesis emerges;
 - synthesis spreads.

6.5 Conclusions

The CBEST Threat Model reflects a conceptual foundation and common terminological baseline for guiding threat assessment activities and maintaining comparability of results across them. Crucially, the threat model is not intended to be a checklist of information fields to be filled in mechanistically nor is it expected that intelligence inputs sufficient to cover every suggested field or category will be available for any given test. Indeed, one can successfully determine the relative threat potential of most entities without filling in all or even most fields although some, such as the entity's identity and some malicious activity indicators, will always be necessary.

The CBEST threat model is meant to facilitate analysis and to enable the analysts performing it, not to reduce the analytic process unrealistically to an algorithm or to prohibit the flexibility, innovation, critical thinking and creativity that improve assessments over time. To that end, the model provides an information 'wish list' that threat intelligence analysts can use when performing threat assessments. Those fields that are left blank can be listed in a rough order of priority as intelligence gaps for future efforts. In any case, referencing the CBEST threat model will help the programme's threat assessors communicate their findings, analytic outputs and insights in a common vocabulary. Furthermore, the explanation of the model above will also serve as the point of departure for future efforts to improve or modify it.

Cyber threat analysts can engage the CBEST threat model's structure with confidence because it has been developed with reference to the most rigorous and successful approaches available from non-classified sources. None of these previous efforts

has been perfect enough to incite universal adoption across even a significant minority of the threat intelligence profession and, for this reason, the CBEST model neither copies wholesale nor ignores previous successes. The model's elements were chosen and organised with a strong prioritisation of logical rigor and conceptual clarity but not to the detriment of the model's practical usability.

This document has specified the CBEST Threat Model only in terms of the two to three highest levels of conceptual generality. Put differently, each category or data field found in the model can be expanded down to the level of observable, measurable concreteness, configurable as a vast number of options including composite indicators, weightings, indices and metrics of statistical inference. Indeed, some threat intelligence organisations will have already achieved this level of detail for some types of threat indicators. However, this overview document is not the appropriate vehicle for explaining the many ways of applying threat modelling concepts as part of the analytical process. For those interested in more detailed explanations, readers should reference the many sources referenced in this document or contact the CBEST programme officials.

As CBEST threat assessment activities proceed they will reveal insights about the methodological and conceptual value of the threat model presented in this report. This practical application of the model will provide valuable feedback by which practitioners may further improve it.

References

- D'Amico, A, Buchanan, L, Goodall, J and Walczak, P (2010)**, 'Mission impact of cyber events: scenarios and ontology to express the relationships between cyber assets, missions and users'. In proceedings of the International Conference on Information Warfare and Security, Dayton, Ohio, April.
- Atkins, W (2010)**, 'Read teaming: it's good to be bad'. Missouri S&T ACM SIG in Security, available at http://acm.device.mst.edu/security-files/2010-02-10-Red_Teaming.ppt. Sandia Corporation.
- Badgers (2011)**, 'Proceedings of the first workshop on building analysis datasets and gathering experience returns for security', available at <http://iseclab.org/badgers2011/badgers2011-proceedings.pdf>. Association for Computing Machinery, Inc.
- Bejtlich, R (2009)**, 'TaoSecurity: incident phases of compromise', available at <http://taosecurity.blogspot.com/2009/06/incident-phases-of-compromise.html>.
- Billo, C (2004)**, 'Cyber warfare: an analysis of the means and motivations of selected nation states', available at www.ists.dartmouth.edu/docs/cyberwarfare.pdf. Trustees of Dartmouth College.
- Bishop, M (1995)**, 'A taxonomy of UNIX system and network vulnerabilities'. Report CSE-95-10, available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.33.5712>. University of California at Davis.
- Brown, B, Cutts, A, McGrath, D, Nicol, D, Smith, T and Tofel, B (2003)**, 'Simulation of cyber attacks with applications in homeland defense training'. In Sensors and Command, Control, Communications and Intelligence (C3I) Technologies for Homeland Defense and Law Enforcement, pages 63–71.
- CBEST (2016)**, 'Understanding Cyber Threat Intelligence Operations', Bank of England.
- Cebula, J and Lisa, R (2010)**, 'A taxonomy of operational cyber security risks'. CMU/SEI-2010-TN-028, available at www.sei.cmu.edu/library/abstracts/reports/10tn028.cfm. Carnegie Mellon University/Software Engineering Institute.
- Chapman, I, Leblanc, S and Partington, A (2011)**, 'Taxonomy of cyber attacks and simulation of their effects'. In proceedings of the 2011 Military Modeling and Simulation Symposium, pages 73–80, Boston, Massachusetts.
- Chuvakin, A (2014)**, 'Delving into threat actor profiles', available at http://blogs.gartner.com/anton-chuvakin/2014/03/14/delving-into-threat-actor-profiles/?utm_medium=twitter. Gartner, Inc.
- Cohen, F (1999)**, 'Simulating cyber attacks, defences and consequences'. Computers and Security, pages 479–518. Elsevier Science Ltd.
- Costantini, K (2007)**, 'Development of a cyber attack simulator for network modeling and cyber security analysis'. Unpublished manuscript, available at <https://ritdml.rit.edu/bitstream/handle/1850/5440/KCostantiniThesis10--2007.pdf?sequence=1>. Department of Industrial and Systems Engineering, Rochester Institute of Technology, Rochester, New York.
- Cowan, C, Arnold, S, Beattie, S, Wright, C and Viega, J (2003)**, 'Defcon capture the flag: defending vulnerable code from intense attack'. In proceedings of the DARPA Information Survivability Conference and Exposition, April.
- Davis, P and Arquilla, J (1991)**, 'Thinking about opponent behavior in crisis and conflict: a generic model for analysis and group discussion', RAND Note N-3322-JS, available at www.rand.org/content/dam/rand/pubs/notes/2007/N3322.pdf. The RAND Corporation.
- DefenseNews (2012)**, 'Building better cyber red teams', available at www.defensenews.com/article/20120614/TSJ01/306140003/Building-Better-Cyber-Red-Teams. Gannett Government Media Corporation.
- Department of Defense Science Board (2003)**, 'Task Force Report on the role and status of red teaming activities', available at www.au.af.mil/au/awc/awcgate/dod/dsb-redteam.pdf. Office of the Under Secretary of Defense For Acquisition, Technology and Logistics.
- Duggan, D and Hutchinson, R (2004)**, 'Red teaming 101', available at www.cs.nmt.edu/%7Ecs491_02/RedTeaming-4hr.pdf. Sandia National Laboratories.

- Duggan, D, Thomas, S, Veitch, C and Woodward, L (2007)**, 'Categorizing threat: building and using a generic threat matrix', available at http://idart.sandia.gov/methodology/materials/Adversary_Modeling/SAND2007-5791.pdf. Sandia National Laboratories.
- Elster, J (2007)**, 'Explaining social behavior: more nuts and bolts for the social sciences'. Cambridge University Press.
- Eriksson, J and Noreen, E (2002)**, 'Setting the agenda of threats: an explanatory model'. *Uppsala Peace Research Papers No. 6*, available at www.pcr.uu.se/publications/UPRP_pdf/uprp_no_6.pdf. Uppsala University.
- Ezell, B, Farr, J and Wiese, I (2000)**, 'Infrastructure risk analysis of municipal water distribution system', *Journal of Infrastructure Systems*, Vol. 6, Issue 3, pages 118–22. American Society of Civil Engineers.
- Fontenot, G (2005)**, 'Seeing red: creating a red-team capability for blue force', available at www.au.af.mil/au/awc/awcgate/milreview/fontenot.pdf. Military Review.
- Fovino, I, Coletta, A and Masera, M (2010)**, 'Taxonomy of security solutions for the SCADA sector'. Deliverable D2.2, Version 1.1. European Network for the Security of Control and Real Time Systems.
- Frazier, T (2009)**, 'Lessons from eight years of government experiments in cyber warfare research and development', available at <http://webhost.laas.fr/TSF/IFIPWG/Workshops&Meetings/56/workshop/04.frazier.pdf>. BAE Systems.
- Gallegos, F and Smith, M (2006)**, 'Red teams: an audit tool, technique and methodology for information assurance', available at www.isaca.org/Journal/Past-Issues/2006/Volume-2/Pages/Red-Teams-An-Audit-Tool-Technique-and-Methodology-for-Information-Assurance1.aspx. ISACA.
- Gladman, B (2007)**, 'The 'best practices' of red teaming'. DRDC CORA TM 2007-29. Centre for Operational Research and Analysis.
- Goldsmith, D and Siegel, M (2012)**, 'Systematic approaches to cyber insecurity', available at <http://ecir.mit.edu/images/stories/Goldsmith%20Siegel%20Modeling%20Cyber%20Dynamics%201%2025%2012%202.pdf>. MIT Sloan School of Management.
- Hansman, S and Hunt, R (2004)**, 'A taxonomy of network and computer attacks', *Computers and Security*, Vol. 24, Issue 1, pages 31–43, available at <http://dx.doi.org/10.1016/j.cose.2004.06.011>.
- Herrmann, R and Fischerkeller, M (1995)**, 'Beyond the enemy images and spiral model: cognitive-strategic research after the Cold War', *International Security*, Vol. 49, No. 3, pages 415–50.
- Howard, J (1997)**, 'An analysis of security incidents on the Internet 1989–1995'. Doctoral dissertation, available at www.cert.org/archive/pdf/JHThesis.pdf. Carnegie Mellon University.
- Howard, J and Longstaff, T (1998)**, 'A common language for computer security incidents'. Sandia Corporation.
- Hutchins, E, Cloppert, M, and Amin, R (2011)**, 'Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains'. In proceedings of 6th Annual International Conference on Information Warfare and Security. Lockheed Martin Corporation.
- IDART (2009)**, 'The Information Design Assurance Red Team (IDART™)', available at www.idart.sandia.gov/index.html. Sandia Corporation.
- Joint Chiefs of Staff (2006)**, 'Joint Publication (JP) 3–13, Information Operations, 13 February, available at www.carlisle.army.mil/DIME/documents/jp3_13.pdf. Department of the Army.
- Joint Chiefs of Staff (2012)**, 'Joint Publication (JP) 3–13, Information Operations, 27 November, available at www.dtic.mil/doctrine/new_pubs/jp3_13.pdf. Department of the Army.
- Killourhy, K, Maxion, R and Tan, K (2004)**, 'A defense-centric taxonomy based on attack manifestation'. Presented at the International Conference on Dependable Systems and Networks, Florence, Italy.
- Kjaerland, M (2006)**, 'A taxonomy and comparison of computer security incidents from the commercial and government sectors'. *Computers and Security*, Vol. 25, Issue 7, pages 522–38, available at <http://dx.doi.org/10.1016/j.cose.2006.08.004>.
- Kuhl, M, Kistner, J, Costantini, K and Sudit, M (2007)**, 'Cyber attack modeling and simulation for network security analysis'. In proceedings of the 39th Conference on Winter Simulation: 40 years! The Best is Yet to Come, Washington D.C.

- Lambe, P (2006)**, 'Defining taxonomy', available at www.greenchameleon.com/gc/blog_detail/defining_taxonomy/. Straits Knowledge.
- Lauder, M (2009)**, 'Red Dawn: the emergence of a red teaming capability in the Canadian forces'. *Canadian Army Journal*, Vol. 12.2. Canadian Army Publishing.
- Liljenstam, M, Liu, J, Nicol, D, Yuan, Y, Yan, G and Grier, C (2006)**, 'RINSE: The Real-Time Immersive Network Simulation Environment for network security exercises (extended version). *Simulation*, Vol. 82, No. 1, pages 43–59.
- Mateski, M, Trevino, C, Veitch, C, Michalski, J, Harris, J, Maruoka, S and Frye, J (2012)**, 'Cyber threat metrics'. Sandia Corporation.
- McGannon, M (2004)**, 'Developing red team tactics, techniques and procedures'. *Red Team Journal*.
- McGannon, M and Pollick, R (2005)**, 'AFRL RED TEAM COOKBOOK', Vol. 1: Red Teaming 101. Plans and Programs Directorate (AFRL/XP), Air Force Research Laboratory. Air Force Materiel Command.
- Microsoft (2005)**, 'Security design by threat modeling', available at [http://msdn.microsoft.com/en-us/library/ee810542\(v=cs.20\).aspx](http://msdn.microsoft.com/en-us/library/ee810542(v=cs.20).aspx). Microsoft Corporation.
- Mirkovic, J and Reiher, P (2004)**, 'A taxonomy of DDoS attack and DDoS defense mechanisms'. *ACM SIGCOMM Computer Communication Review*, Vol. 34, Issue 2, pages 39–53, available at <http://dx.doi.org/10.1145/997150.997156>.
- MIT (2008)**, 'MIT CASCON system for analyzing international conflict', available at <http://web.mit.edu/cascon/>. Massachusetts Institute of Technology.
- Mitre (2000)**, 'Defense-information assurance red team'. The MITRE Corporation.
- Mitre (2013)**, 'Mapping the cyber terrain', available at www.mitre.org/sites/default/files/publications/mapping-cyber-terrain-13-4175.pdf. The MITRE Corporation.
- Mitre (2014)**, 'CAPEC: Common Attack Pattern Enumeration and Characterization', available at <http://capec.mitre.org/>. The MITRE Corporation.
- Moore, A, Cappelli, D, Joseph, H and Trzeciak, R (2007)**, 'An experience using system dynamics to facilitate an insider threat workshop'. In proceedings of the International System Dynamics Conference, Boston, MA.
- OMG (2014)**, 'Threat modeling and sharing', available at www.omg.org/hot-topics/documents/threat/Threat-Modeling-and-Sharing-Presentation.pdf. Object Management Group, Inc.
- Parks, R (2010)**, 'Cross domain red teaming'. *Red Team Journal*. Retrieved 20 November, available at <http://redteamjournal.com/2010/07/cross-domain-red-teaming/#more-2470>. Sandia Corporation.
- Penn State University SIIS Laboratory (2010)**, 'Advanced metering infrastructure security', available at <http://siis.cse.psu.edu/smartgrid.html>. Penn State University Systems and Internet Infrastructure Security Laboratory.
- PricewaterhouseCoopers (2010)**, 'Cyber attacks: is your critical infrastructure safe?', available at www.pwc.com/en_US/us/industry/utilities/assets/cyber-attacks.pdf.
- Rich, E, Martinez-Moyano, I, Conrad, S, Cappelli, D, Moore, A, Shimeall T, Andersen D, Gonzalez, J., Ellison, R, Lipson, H, Mundie, D, Sarriegui, J, Sawicka, A, Stewart, T, Torres, J, Weaver, E and Wiik, J (2005)**, 'Simulating insider cyber-threat risks: a model-based case and a case-based model'. In proceedings of the International System Dynamics Conference, Boston, MA.
- Rogers, M (2006)**, 'A two-dimensional circumplex approach to the development of a hacker taxonomy', *Digital Investigation*, Vol. 3, Issue 2, pages 97–102, available at <http://dx.doi.org/10.1016/j.diin.2006.03.001>.
- Sakhardande, R (2008)**, 'The use of modeling and simulation to examine network performance under denial of service attacks'. Unpublished manuscript. Department of Telecommunications, SUNY Institute of Technology.
- SANS (2009)**, 'Twenty critical controls for effective cyber defense: consensus audit guidelines', available at www.sans.org/critical-security-controls/cag.pdf. The SANS Institute.
- Schneier, B (1999)**, 'Modeling security threats: attack trees', available at www.schneier.com/paper-attacktrees-ddj-ft.html. Counterpane Internet Security.

- Schudel, G and Wood, B (2000)**, 'Modeling the behavior of a cyber terrorist', available at www.csl.sri.com/users/bjwood/cyber_terrorist_model_v4a.pdf. RAND National Security Research Division/GTE/BBN Technologies/Sandia National Laboratories.
- Silverman, B (2014)**, 'Creating human behavior models able to enhance synthetic agents: research results to date', available at www.seas.upenn.edu/~barryg/HBMR.html. Penn Engineering.
- Simmons, C, Ellis, C, Shiva, S, Dasgupta, D and Wu, Q (2009)**, 'AVOIDIT: a cyber attack taxonomy', available at http://issrl.cs.memphis.edu/files/papers/CyberAttackTaxonomy_IEEE_Mag.pdf.
- Skroch, M (2009)**, 'Modeling and simulation of red teaming Part 1: why red team M&S?', available at <http://umbra.sandia.gov/pdfs/resources/redteam.pdf>. Sandia Corporation.
- Sou Park, J, Lee, J, Kuk Kim, H, Jeong, J, Yeom, D and Chi, S (2001)**, 'SECUSIM: a tool for the cyber-attack simulation'. In proceedings of the Third International Conference on Information and Communications Security, pages 471–75.
- Steele, R (2002)**, 'The new craft of intelligence: achieving asymmetric advantage in the face of non-traditional threats'. US Army War College Strategic Studies Institute.
- Stoneburner, G (2001)**, 'Computer security: underlying technical models for information technology security'. Recommendations of the National Institute of Standards and Technology. NIST Special Publication 800–33. US Government Printing Office.
- US Department of Defense (2003)**, 'The role and status of DoD red teaming activities', available at www.fas.org/irp/agency/dod/dsb/redteam.pdf. Defense Science Board.
- University of Foreign Military and Cultural Studies (2011)**, 'Red Team Handbook version 5', available at www.au.af.mil/au/awc/awcgate/army/ufmcs_red_team_handbook_apr2011.pdf. Department of the Army.
- VERIS (2014)**, 'VERIS Community: a resource for learning about the VERIS framework', available at www.veriscommunity.net/doku.php. Verizon.
- WASC (2013)**, 'The Web Application Security Consortium threat classification version 2.0', available at <http://projects.webappsec.org/w/page/13246978/Threat-Classification>.
- Zhou, M and Lang, S (2003)**, 'A frequency-based approach to intrusion detection'. Systemics, Cybernetics and Informatics, Vol. 2, Issue 3, pages 52–56.
- Zhu, B, Joseph, A and Sastry, S (2011)**, 'A taxonomy of cyber attacks on SCADA systems'. IEEE International Conferences on Internet of Things and Cyber, Physical and Social Computing.