

(12) **United States Patent**
Agarwal

(10) **Patent No.:** **US 10,699,008 B2**
(45) **Date of Patent:** ***Jun. 30, 2020**

(54) **THREAT MODEL CHAINING AND ATTACK SIMULATION SYSTEMS AND RELATED METHODS**

(56) **References Cited**

U.S. PATENT DOCUMENTS

(71) Applicant: **ThreatModeler Software Inc.**, Jersey City, NJ (US)

4,959,015 A 9/1990 Rasinski et al.
6,952,779 B1 * 10/2005 Cohen G06F 21/577
726/22

(72) Inventor: **Anurag Agarwal**, Jersey City, NJ (US)

(Continued)

(73) Assignee: **ThreatModeler Software Inc.**, Jersey City, NJ (US)

OTHER PUBLICATIONS

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

Threat Risk Modeling, published online at least as early as Feb. 3, 2018 by OWASP, available at https://www.owasp.org/index.php/Threat_Risk_Modeling.

(Continued)

This patent is subject to a terminal disclaimer.

Primary Examiner — Hosuk Song

(21) Appl. No.: **16/228,738**

(74) *Attorney, Agent, or Firm* — Law Office of Paul B. Johnson; Paul B. Johnson

(22) Filed: **Dec. 20, 2018**

(65) **Prior Publication Data**

US 2019/0147161 A1 May 16, 2019

Related U.S. Application Data

(63) Continuation-in-part of application No. 15/922,856, filed on Mar. 15, 2018, now Pat. No. 10,200,399, (Continued)

(51) **Int. Cl.**
H04L 29/06 (2006.01)
G06F 21/55 (2013.01)
(Continued)

(52) **U.S. Cl.**
CPC **G06F 21/552** (2013.01); **G06F 16/27** (2019.01); **G06F 21/554** (2013.01);
(Continued)


(58) **Field of Classification Search**
CPC G06F 21/577; G06F 21/552; G06F 21/554;
G06F 17/30283; H04L 63/1433; H04L
63/20

See application file for complete search history.

(57) **ABSTRACT**

Threat model chaining methods include providing one or more databases including a threat model components, threats, each threat associated with at least one of the threat model components, and compensating controls, each compensating control associate with one of the threats, providing a diagram interface configured to display a relational diagram defining a first threat model, and configuring the diagram interface to add a component group to the first threat model include in it a second threat model. Attack simulation methods include providing the one or more databases and diagram interface and configuring the diagram interface to visually display attack paths of threats associated with diagrammed threat model components which compromise a selected threat model component. Attack simulation systems include one or more computing devices coupled with one or more databases configured to store and interrelate threats, threat model components, and compensating controls, and allow diagramming and defining of threat models.

20 Claims, 22 Drawing Sheets



Related U.S. Application Data

which is a continuation-in-part of application No. 15/888,021, filed on Feb. 3, 2018, now Pat. No. 10,255,439.

- (60) Provisional application No. 62/507,691, filed on May 17, 2017, provisional application No. 62/527,671, filed on Jun. 30, 2017, provisional application No. 62/530,295, filed on Jul. 10, 2017, provisional application No. 62/520,954, filed on Jun. 16, 2017.

(51) **Int. Cl.**

G06F 16/27 (2019.01)

G06F 21/57 (2013.01)

G06F 30/20 (2020.01)

(52) **U.S. Cl.**

CPC **G06F 21/577** (2013.01); **G06F 30/20** (2020.01); **G06F 2221/034** (2013.01)

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,096,502 B1	8/2006	Fox et al.	
7,260,844 B1	8/2007	Tidwell et al.	
7,433,829 B2	10/2008	Borgia et al.	
7,624,448 B2	11/2009	Coffman	
7,891,003 B2	2/2011	Mir	
7,900,259 B2	3/2011	Jeschke et al.	
8,191,139 B2	5/2012	Heimerdinger et al.	
8,255,995 B2	8/2012	Kraemer et al.	
8,407,801 B2 *	3/2013	Ikegami	G06F 21/577 726/25
8,413,237 B2	4/2013	O'Rourke et al.	
8,413,249 B1 *	4/2013	Chou	G06F 11/00 726/25
9,043,924 B2 *	5/2015	Maor	G06F 11/3688 726/25
9,497,203 B2	11/2016	Honig et al.	
9,774,613 B2	9/2017	Thomas et al.	
10,200,399 B2 *	2/2019	Agarwal	G06F 21/577
10,255,439 B2 *	4/2019	Agarwal	G06F 16/904
10,262,132 B2	4/2019	Reinecke et al.	
2001/0027388 A1	10/2001	Beverina et al.	
2006/0015941 A1	1/2006	McKenna	
2006/0241991 A1	10/2006	Pudhukottai et al.	
2011/0126111 A1	5/2011	Gill et al.	
2011/0178942 A1	7/2011	Watters et al.	
2011/0277034 A1	11/2011	Hanson	
2014/0137257 A1	5/2014	Martinez et al.	

2014/0157417 A1	6/2014	Grubel et al.
2015/0033346 A1	1/2015	Hebert et al.
2016/0162690 A1	6/2016	Reith et al.
2016/0248798 A1	8/2016	Cabrera et al.
2017/0213037 A1	7/2017	Toledano et al.
2018/0255084 A1	9/2018	Kotinas et al.
2018/0324207 A1	11/2018	Reybok, Jr. et al.

OTHER PUBLICATIONS

Threat Model, Wikipedia, published online at least as early as Feb. 3, 2018, available at https://en.wikipedia.org/wiki/Threat_model (note that some of the material in this Wikipedia article appears to be posted by one or more inventors of the present application or its parent provisional applications—some of the material appears to be posted by Brian Beyst, for example, who is an inventor on two of the parent provisional applications).

Comparisons of Threat Model Modeling Methodologies, published online by ThreatModeler at least as early as Apr. 15, 2016, available at <http://threatmodeler.com/comparison-threat-modeling-methodologies/> (note that this article appears to be affiliated with a business organization of, and may have been published under the direction of, one or more inventors of the present application or its parent provisional applications).

Microsoft Threat Modeling Tool, published online by Microsoft, different sections published at different times from Aug. 22, 2017 to Jan. 24, 2018, available online at <https://docs.microsoft.com/en-us/azure/opbuildpdf/security/TOC.pdf?branch=live>.

A description of an on-sale version of systems and methods, on sale in the U.S. by applicants at least as early as Dec. 31, 2011, which on-sale version disclosed some of the elements disclosed in the present application. The description includes a detailed discussion of which claim elements of the claims as originally filed in the present application were disclosed in the 2011 version and which were not, and includes screenshots of user interfaces used in the system/methods.

Lockheed Martin Corporation, "Seven Ways to Apply the Cyber Kill Chain with a Threat Intelligence Platform," published at least as early as 2015.

Roy, Maurer. "Top Database Security Threats and How to Mitigate Them." Jul. 30, 2015. Retrieved from "<https://www.shrm.org/resourcesandtools/hr-topics/risk-management/pages/top-database-security-threats.aspx>" (Year: 2015).

A. Amini, N. Jamil, A.R. Ahmad and M.R. Z'aba, 2015. Threat Modeling Approaches for Securing Cloud Computing. Journal of Applied Sciences, 15: 953-967. Retrieved from "<https://scialert.net/fulltextmobile?doi=jas.2015.953.967>" (Year: 2015).

* cited by examiner








FIG. 1

**FIG. 2**

**FIG. 3**

**FIG. 4**




SELECT ALL NEW DELETE COPY TO LIBRARY HOME KNOWLEDGE BASE TEMPLATES PROFILE CHANGE PASSWORD SIGN OUT		
EDIT SELECT LIBRARY THREAT MODELER ▼		
COMPONENTS		
NAME	TYPE	LABELS
WIFI TR	DEVICE	
MODEM	DEVICE	EMBED. SYS.
WIFI PORT	IOT DEVICE	IOT
HTTPS	PROTOCOLS	
COOKIES	GEN DT STORE	TMT
DATABASE	DB	HARDWARE
LOGIN	APP. COMPON.	WEB, APP, FEAT.
COMMENT	APP. COMPON.	WEB, APP, FEAT.
SHOP CART	APP. COMPON.	WEB, APP, FEAT.
ACH TRANS.	APP. COMPON.	WEB, APP, FEAT.
USB PORT	IOT DEVICE	IOT
FILE SVR	DEPLOY. COMP.	
EMAIL SVR	DEPLOY. COMP.	
FIREWALL	APP. COMPON.	

LOGIN	
TYPE: APPLICATION COMPONENT	
WEB APPLICATION	
PASSWORD	
AUTHENTICATION	
FEATURE	
DESCRIPTION	

PROPERTIES
THREATS +
INDUCING X
ACCT LOCKOUT
SESSION X
HIDJACKING
CLICK JACKING X
SECURITY +
REQUIREMENTS
RANDOMIZING X
SECRET
ENCRYPTION X
ALGORITHM
PARAM. X
DYNAMIC SQL

FIG. 6

600

**FIG. 7**

≡
✓
+
/
✕
□

SELECT LIBRARY

THREAT MODELER ▼

🏠
📁
📄
▶

THREATS		SENSITIVE DATA EXPOSURE		PROPERTIES
NAME	RISK	RISK: VERY HIGH		THREAT AGENTS +
		OWASP-A6		
BUFFER OVERFLOW	VERY HIGH	CAPEC-610		
SESSION HIJACK	VERY HIGH	OWASP-A5-SECURITY MIS-CONFIGURATION		TEST CASES +
PHYSICAL THEFT	VERY HIGH			TEST CASE 01 ✕
HTTP SPLITTING	HIGH			TEST CASE 02 ✕
BRUTE FORCE	HIGH			TEST CASE 03 ✕
DBL ENCODING	MEDIUM	DESCRIPTION		SECURITY REQUIREMENTS +
SENS. DATA EXP.	VERY HIGH			CLICKJACK CHEAT SHEET ✕
INTEGER ATTACKS	MEDIUM	WWW.OWASP.ORG		TEST COMP CONTROL ✕
POINTER ATTACK	MEDIUM			SS-COMPENS. CONTROL ✕
SYMLINK ATTACK	MEDIUM			
EMAIL INJECT.	MEDIUM			
FORMAT STR. INJ.	HIGH			
AUTOMATION ATK	VERY HIGH			
SPAM	VERY HIGH			

800


FIG. 8

The interface displays a 'THREAT MODELER' window with a toolbar at the top containing icons for home, document, and navigation. Below the toolbar is a 'SELECT LIBRARY' dropdown menu. The main content area is divided into two panes. The left pane, titled 'SECURITY REQUIREMENTS', contains a list of requirements with columns for 'NAME' and 'LABELS'. The right pane, titled 'PROPERTIES', shows details for the selected requirement, including 'NAME', 'DESCRIPTION', 'IS COMPENSATING CONTROL', and 'LABELS'. At the bottom right of the right pane are 'SAVE' and 'CANCEL' buttons.

SECURITY REQUIREMENTS		PROPERTIES	
NAME	LABELS	NAME	CODE SNIPPETS
ENCIPHERMENT: USE ONLY STR...	AZURE LIB...	DESCRIPTION	CODE REVIEWS
ENCIPHERMENT: ENSURE RAND...	DOM,...		
INPUT VALIDATION WITH...	HARDWARE,...		
TM - INPUT VALIDATION	JAVA,INI,...	IS COMPENSATING CONTROL ☐	
USING SSL	CERTS,...		
KEY MANAGEMENT: ENSU...	OWASP,...		
TLS: PREVENT CACHING OF...			
CERTIFICATES: USE AN APP...	CERTS,...		
ADD FIREWALL RULES	EVENT,TO,...		
HOST INTEGRITY CHECKING	TO,SSE,...		
SECURE BROWSERS	SCRIPT,LOG,...		
OUTPUT ENCODING			
APPLY PATCHES	XSS,REL,...		
DISABLE HTTP TRACE	INI,HTTP,...		

SAVE CANCEL

FIG. 9



✕
MITIGATIONS

DRAG A COLUMN HEADER AND DROP IT HERE TO GROUP BY THAT COLUMN

THREAT	▼	SOURCE	▼	RISK	▼	STATUS	▼
FILE MANIPULATION		FILE SYSTEM		MEDIUM		OPEN	
ACCESSING EXECUTABLE FILES		FILE SYSTEM		VERY HIGH		OPEN	
CREATE FILES SAME NAME AS...		FILE SYSTEM		VERY HIGH		OPEN	
MAN IN THE BROWSER		IE 11		VERY HIGH		OPEN	
MANIPULATING CONFIG FILE		FILE SYSTEM		VERY HIGH		OPEN	
FORCE USE OF CORRUPT FILE		FILE SYSTEM		VERY HIGH		OPEN	
MANIPULATING INPUT TO CALL		FILE SYSTEM		HIGH		OPEN	
FILE MANIPULATION		FILE SYSTEM		MEDIUM		OPEN	
ACCOUNT FOOTPRINTING		IE 11		VERY HIGH		OPEN	
AUTOMATION ATTACK		IE 11		VERY HIGH		OPEN	
PHYSICAL THEFT		LAPTOP		VERY HIGH		OPEN	
PHYSICAL THEFT		LAPTOP		VERY HIGH		OPEN	
BLUEJACKING		BLUETOOTH PORT		VERY HIGH		OPEN	
PHISHING		OUTLOOK 2010		VERY HIGH		OPEN	
SPAM		OUTLOOK 2010		VERY HIGH		OPEN	
DNS CACHE POISONING		EMAIL SERVER		VERY HIGH		OPEN	

MITIGATE
CLOSE

1100

FIG. 11

COMPENSATING CONTROLS					
DRAG A COLUMN HEADER AND DROP IT HERE TO GROUP BY THAT COLUMN					
CONTROL	THREAT	SOURCE	RISK	STATUS	
FIREWALL	TCP WINDOW SCAN	WINDOWS 7	LOW	MITIGATED	
FIREWALL	TCP SYN SCAN	TCP	LOW	OPEN	
FIREWALL	TCP ISN CTR RATE PROBE	TCP	LOW	OPEN	
FIREEYE	MAN IN THE BROWSER	IE 11	VERY HIGH	MITIGATED	
CYLANCE	MALWARE PROP VIA USB	USB PORT	VERY HIGH	MITIGATED	
AV DLP	FORCE USE CORRUPT FILE	FILE SYSTEM	VERY HIGH	MITIGATED	
MCAFEES DLP	MAN IN THE MIDDLE ATK	WIFI PORT	VERY HIGH	MITIGATED	
BIT9	FILE MANIPULATION	FILE SYSTEM	MEDIUM	MITIGATED	
BIT9	ACCOUNT FOOTPRINTING	IE 11	VERY HIGH	OPEN	
BIT9	AUTOMATION ATTACK	IE 11	VERY HIGH	MITIGATED	
AV HIPS	PHYSICAL THEFT	LAPTOP	VERY HIGH	MITIGATED	
BITLOCKER	PHYSICAL THEFT	LAPTOP	VERY HIGH	MITIGATED	
IPC	BLUEJACKING	BLUETOOTH PORT	VERY HIGH	OPEN	
IRONPORT	PHISHING	OUTLOOK 2010	VERY HIGH	MITIGATED	
IRONPORT	SPAM	OUTLOOK 2010	VERY HIGH	MITIGATED	
AKAMAI DNS	DNS CACHE POISONING	EMAIL SERVER	VERY HIGH	OPEN	

1200

FIG. 12





FIG 13

**FIG. 14**






FIG. 15

**FIG. 16**

**FIG. 17**




FIG. 18







FIG. 19

**FIG. 20**





1

THREAT MODEL CHAINING AND ATTACK SIMULATION SYSTEMS AND RELATED METHODS

CROSS REFERENCE TO RELATED APPLICATIONS

This document is a continuation-in-part application of U.S. patent application Ser. No. 15/922,856 titled "Threat Model Chaining and Attack Simulation Systems and Methods," naming as first inventor Anurag Agarwal, filed Mar. 15, 2018, now pending, which in turn is a continuation-in-part application of U.S. patent application Ser. No. 15/888,021 titled "Threat Modeling Systems and Related Methods Including Compensating Controls," naming as first inventor Anurag Agarwal, filed Feb. 3, 2018, (hereinafter the '021 application), which in turn claims the benefit of the filing date of U.S. Provisional Pat. App. Ser. No. 62/507,691 titled "System and Method of Including Compensating Controls in a Threat Modeling Process," naming as first inventor Anurag Agarwal, filed May 17, 2017, now expired, and which '021 application also claims the benefit of the filing date of U.S. Provisional Pat. App. Ser. No. 62/527,671 titled "System and Method for Identifying and Analyzing the Potential Attack Surface of a Complex System," naming first inventor Anurag Agarwal, filed Jun. 30, 2017, and which '021 application also claims the benefit of the filing date of U.S. Provisional Pat. App. Ser. No. 62/530,295 titled "Method and Apparatus for Early Implementation of Enterprise DevSecOps," naming first inventor Anurag Agarwal, filed Jul. 10, 2017, now expired, and which '021 application also claims the benefit of the filing date of U.S. Provisional Pat. App. Ser. No. 62/520,954 titled "System and Method for Identifying Potential Threats to a Complex System," naming first inventor Anurag Agarwal, filed Jun. 16, 2017, now expired, the disclosures of each of which are hereby incorporated entirely herein by reference.

BACKGROUND

1. Technical Field

Aspects of this document relate generally to threat modeling processes and systems.

2. Background Art

Threat modeling is a process by which vulnerabilities of a system or process may be detailed and prioritized. One example of an existing threat modeler is a modeler marketed under the name THREAT MODELING TOOL by Microsoft Corporation of Redmond, Wash. Threat modeling allows a user to analyze potential attack vectors and prioritize vulnerabilities. While some threat modeling involves threats related to computing networks and systems, threat modeling in general encompasses a broader scope and may involve modeling threats in non-computer-related systems and processes. Some commercial threat modeling tools utilize a dynamic question and answer user interface which generally relies on an underlying data flow diagram (DFD) or process flow diagram (PFD) conceptual basis.

With regards to modeling computing networks, traditional threat modeling systems and methods exist for assessing risk from potential cyber threats at an individual application level without fully accounting for application-application interactions, thereby miscalculating organizational risk. Other systems and methods exist for discovering vulner-

2

abilities to applications deployed in cyber environments by surveying applications. These rely on recognizing previously identified and cataloged vulnerability signatures and are not designed to survey non-application items included in an organization's cyber environment, nor can these systems recognize which threats newly introduced to a cyber system are relevant to security, nor can these systems perform "what-if" scenarios as part of an organization's risk management analysis. Existing methodologies also do not provide for communicating the risks associated with identified threats in a way that non-security experts easily appreciate or understand.

Accordingly, traditional threat modeling methodologies have the following weaknesses: (1) they are effective in analyzing only single applications operating in isolation on a predetermined infrastructure; (2) they require security subject-matter experts for their creation, use, and maintenance; (3) they are resource-intensive to build and maintain, and; (4) they cannot be effectively used to scale a threat modeling practice to meet the needs of enterprises generating tens of software applications per year. Traditional threat modeling methodologies also cannot be easily integrated into existing agile software development approaches or with DevOps practices, and so are often rejected by agile developers and operational teams as non-productive.

Existing threat modeling methodologies also do not help organizations understand: (1) the nature of application interactions; (2) the user of shared components; (3) the effect of including third-party elements within the IT ecosystem; (4) the downstream impact should potential threats be realized; (5) the nature of the organization's comprehensive attack surface; (6) the explicit nature of the organization's attacker population; (7) the effectiveness of deployed or contemplated compensating controls; or (8) a means to communicate the potential impact should threats be realized without communicating confidential information or security details which could further expose the organization to risk from potential threats.

SUMMARY

Embodiments of threat model chaining methods may include: providing one or more databases, the one or more databases including: a plurality of threat model components stored therein; a plurality of threats stored therein, each threat associated with at least one of the threat model components through the one or more databases; and a plurality of compensating controls stored therein, each compensating control associated with at least one of the threats through the one or more databases; providing one or more interfaces, including a diagram interface, configured to be displayed on one or more end user computing devices communicatively coupled with the one or more databases; configuring the diagram interface to display a relational diagram of one of a system, an application, and a process, using visual representations of one or more of the threat model components and visual representations of one or more of the compensating controls, the relational diagram defining a first threat model, and; configuring the diagram interface to, in response to receiving one or more user inputs, add a component group to the first threat model and thereby redefine the first threat model by including in it a second threat model associated with the component group, wherein the component group includes a predefined interrelated group of two or more of the threat model components.

Embodiments of threat model chaining methods may include one or more or all of the following:

3

Configuring the diagram interface to, in response to receiving a user selection of one of the diagrammed threat model components of the first threat model, visually display attack paths of all threats associated with the diagrammed threat model components which compromise the selected threat model component.

Configuring the one or more interfaces to, in response to receiving one or more user inputs, store one or more additional threat model components in the one or more databases, store one or more additional threats in the one or more databases, associate each additional threat with at least one of the additional threat model components through the one or more databases, store one or more additional compensating controls in the one or more databases, and associate each additional compensating control with at least one of the additional threats through the one or more databases.

Embodiments of attack simulation methods may include: providing one or more databases, the one or more databases including: a plurality of threat model components stored therein; a plurality of threats stored therein, each threat associated with at least one of the threat model components through the one or more databases; and a plurality of compensating controls stored therein, each compensating control associated with at least one of the threats through the one or more databases; providing one or more interfaces, including a diagram interface, configured to be displayed on one or more end user computing devices communicatively coupled with the one or more databases; configuring the diagram interface to display a relational diagram of one of a system, an application, and a process, using visual representations of one or more of the threat model components and visual representations of one or more of the compensating controls, the relational diagram defining a first threat model; and configuring the diagram interface to, in response to receiving a user selection of one of the diagrammed threat model components of the first threat model through the one or more interfaces, visually display attack paths of all threats associated with the diagrammed threat model components which compromise the selected threat model component.

Embodiments of attack simulation methods may include one or more or all of the following:

Configuring the diagram interface to, in response to receiving one or more user inputs through the one or more interfaces, visually display one or more attack paths which pass through multiple diagrammed threat model components of the first threat model.

Configuring the diagram interface to, in response to receiving a user toggle of one of the diagrammed compensating controls (hereinafter "first control") of the first threat model to an on state through the one or more interfaces, toggle all attack paths mitigatable by the first control to a mitigated display state.

Configuring the diagram interface to, in response to receiving a user toggle of the first control to an off state through the one or more interfaces, toggle all attack paths mitigatable by the first control to an unmitigated display state.

Configuring the diagram interface to, in response to receiving one or more user inputs through the one or more interfaces to remove the first control from the first threat model, toggle all attack paths mitigatable by the first control to an unmitigated display state.

Configuring the diagram interface to, in response to receiving one or more user inputs through the one or more interfaces, visually display only attack paths of all threats associated with a subset of the diagrammed threat model components.

4

Configuring the diagram interface to, in response to receiving one or more user inputs through the one or more interfaces to alter a profile of an attacking component, alter the visually displayed attack paths.

Configuring the diagram interface to receive one or more user inputs to add a component group to the first threat model and thereby redefine the first threat model by including in it a second threat model associated with the component group, wherein the component group includes a predefined interrelated group of two or more of the threat model components.

Embodiments of attack simulation systems may include: one or more computing devices communicatively coupled with one or more databases, the one or more computing devices displaying, on one or more displays of the one or more computing devices: one or more input interfaces configured to, in response to receiving one or more user inputs, store a plurality of user-defined threat model components in the one or more databases, store a plurality of threats in the one or more databases, associate each of the threats with at least one of the threat model components through the one or more databases, store a plurality of compensating controls in the one or more databases, and associate each compensating control with at least one of the threats through the one or more databases, and; a diagram interface configured to, in response to receiving one or more user inputs, diagram one of a system, an application, and a process, the diagram including one or more of the threat model components and one or more of the compensating controls, to define a first threat model, the first threat model including all threats associated through the one or more databases with the diagrammed threat model components; wherein the diagram interface is configured to, in response to receiving a selection of one of the diagrammed threat model components of the first threat model, visually display attack paths of all threats associated with the diagrammed threat model components which compromise the selected threat model component.

Embodiments of attack simulation systems may include one or more or all of the following:

The diagram interface may be configured to, in response to receiving a user input to toggle one of the diagrammed compensating controls (hereinafter "first control") to an on state, toggle all attack paths mitigatable by the first control to a mitigated display state.

The diagram interface may be configured to, in response to receiving a user input to toggle the first control to an off state, toggle all attack paths mitigatable by the first control to an unmitigated display state.

The diagram interface may be configured to, in response to receiving one or more user inputs removing the first control from the first threat model, toggle all attack paths mitigatable by the first control to an unmitigated display state.

The diagram interface may be configured to, in response to receiving one or more user inputs, visually display only attack paths of all threats associated with a subset of the diagrammed threat model components.

The diagram interface may be configured to, in response to receiving one or more user inputs, alter a profile of an attacking component, wherein the alteration alters the visually displayed attack paths.

The one of a system, an application, and a process may be a computing network.

The diagram interface may be further configured to, in response to receiving one or more user inputs, diagram a communication protocol between two of the diagrammed

5

threat model components, the diagrammed communication protocol including an alphanumeric indicator.

The diagram interface may be configured to, in response to receiving one or more user inputs, add a component group to the first threat model and thereby redefine the first threat model by including in it a second threat model associated with the component group, wherein the component group includes a predefined interrelated group of two or more of the threat model components.

General details of the above-described embodiments, and other embodiments, are given below in the DESCRIPTION, the DRAWINGS, and the CLAIMS.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments will be discussed hereafter using reference to the included drawings, briefly described below, wherein like designations refer to like elements:

FIG. 1 is a diagram representatively illustrating an implementation of a threat modeling system (threat model chaining system) (attack simulation system) (system);

FIG. 2 is a block diagram representatively illustrating an implementation of a threat modeling method;

FIG. 3 is a block diagram representatively illustrating an implementation of a threat modeling method;

FIG. 4 is a flowchart representatively illustrating an implementation of a threat modeling method;

FIG. 5 is a flowchart representatively illustrating an implementation of a threat modeling method;

FIG. 6 is an implementation of an interface of the system of FIG. 1;

FIG. 7 is an implementation of an interface of the system of FIG. 1;

FIG. 8 is an implementation of an interface of the system of FIG. 1;

FIG. 9 is an implementation of an interface of the system of FIG. 1;

FIG. 10 is an implementation of an interface of the system of FIG. 1;

FIG. 11 is an implementation of an interface of the system of FIG. 1;

FIG. 12 is an implementation of an interface of the system of FIG. 1;

FIG. 13 is an implementation of an interface of the system of FIG. 1;

FIG. 14 is a block diagram representatively illustrating an implementation of an attack simulation method;

FIG. 15 is an implementation of an interface of the system of FIG. 1;

FIG. 16 is a flowchart representatively illustrating an implementation of a threat modeling method;

FIG. 17 is an attack tree diagram representatively illustrating an attack path of a threat of a threat model;

FIG. 18 is a block diagram representatively illustrating attack paths of multiple threats of a threat model;

FIG. 19 is the block diagram of FIG. 18 with compensated controls added to the diagram;

FIG. 20 is a block diagram of an asset classification method of an attack simulation method;

FIG. 21 is an implementation of an interface of the system of FIG. 1, and;

FIG. 22 is an implementation of an interface of the system of FIG. 1.

DESCRIPTION

Implementations/embodiments disclosed herein (including those not expressly discussed in detail) are not limited to

6

the particular components or procedures described herein. Additional or alternative components, assembly procedures, and/or methods of use consistent with the intended system and method of including compensating controls in a threat modeling process may be utilized in any implementation. This may include any materials, components, sub-components, methods, sub-methods, steps, and so forth.

As used herein, the term “input field” includes a “selector.” For example, a button or space on a user interface in which a user may move a cursor to and click to make a selection, and a checkbox field, and other similar fields, as well as alphanumeric input fields, are all “input fields” as used herein.

The term “compensating control” in implementations herein may be an alternative mechanism to a security requirement or standard issued by a standards-issuing body that is allowed by the standards-issuing body when the security requirement or standard as stated cannot be met by a party due to legitimate technical or documented business constraints.

In the payment card industry (PCI), as a non-limiting example, compensating controls were introduced in Payment Card Industry Data Security Standard (PCI DSS) 1.0 to give organizations an alternative to security requirements that could not be met due to legitimate technological or business constraints. According to the PCI Council, which was the standards-issuing body issuing the standard (jointly created by the four major credit-card companies VISA, MASTERCARD, DISCOVER, and AMERICAN EXPRESS), compensatory controls in that industry must: (1) meet the intent and rigor of the original stated requirement; (2) provide a similar level of defense as the original stated requirement; (3) be “above and beyond” other PCI DSS requirements (not simply in compliance with other PCI DSS requirements); and (4) be commensurate with the additional risk imposed by not adhering to the original stated requirement. Examples of compensating controls for information technology may include: using audit trails and logs for payroll management instead of segregation of duties (having two individuals responsible for separate parts of payroll management); the use of database security applications and services, network access control (NAC), data leak prevention strategies, and e-mail encryption in lieu of comprehensive encryption (i.e., in lieu of converting all electronic data into ciphertext and changing cryptographic keys periodically); two-factor authentication with a change of password every 60 days in lieu of long complex passwords; and so forth.

Notwithstanding the above, in implementations herein “compensating control” may have a more general definition. For example, in implementations a “security requirement” may be defined as a mechanism for stopping or closing a threat at the source of the threat, and a “compensating control” may be defined as a mechanism for stopping or closing a threat not at the source, but between the source and some protected component (i.e., preventing a threat emanating from a source from reaching a protected component). In other implementations a “security requirement” may be defined as a hardware solution and a “compensating control” may be defined as a software solution, such as a change in the code or software added to a system. In still other implementations a “security requirement” may be generally defined as any solution which is costlier or more time consuming and a “compensating control” may be defined as a solution which is not as secure or sure-proof as the relevant security requirement but which is less cost-prohibitive or time-prohibitive.

The threat modeling system and related methods discussed herein are implemented using computing devices and/or networks. Referring to FIG. 1, an implementation of a threat modeling system (threat model chaining system) (attack simulation system) (system) **100** is shown. FIG. 1 only shows a representative example, and there are many other contemplated systems that could be used to implement the threat modeling processes. System **100** includes a computing device **102** having a display **104**. While the computing device is drawn as a desktop computer it could be a laptop, a mobile phone or tablet, or any other type of computing device. The same goes for all other computing devices shown in the drawings.

Device **102** is shown communicatively coupled with server **106** which is communicatively coupled with a database (DB) **108**. The coupling may be direct, such as through a wired connection, or through a local area network (LAN), or remotely through telecommunication network **110** (which may be the Internet). In some systems the server and database could be housed on the same machine as the computing device **102** using virtualization. In implementations device **102** could be accessed by an administrator of the system to choose settings, add or remove users, add or remove items from the database, and so forth. System **100** only shows one computing device **102**, though in implementations the number of computing devices **102** may be scaled up to any number. Likewise, only one server and database are shown, but these also may be scaled up to any number as needed.

Other computing devices may be included in system **100**. Computing device **112** includes display **114** and is an example of a computing device which is communicatively coupled with device **102** both directly (such as through a hardwired or wireless LAN), and coupled directly with the server (such as through a hardwired or wireless LAN), and also may be coupled with the server and/or the device **102** through telecommunication network **110**. System **100** is shown with only one device **112** but in implementations it could be scaled up to any number of devices **112**.

Computing device (device) **116** is an example of a computing device that is not directly coupled with either device **102** or the server but is only coupled thereto through the telecommunications network **110**. Nevertheless, device **116** may access the server and database through the telecommunications network. Although only one device **116** is shown, this may be scaled up to any number. Device **116** has a display **118**, as shown.

Also shown are a web server **120** and a remote server (server) **122**, each of which may be included in implementations of system **100**. By non-limiting example, device **116** may access the server **106** and database through the web server **120**, such as by navigating to a uniform resource locator (URL) and providing login credentials. Computing devices **102** and **112** could do the same. Although only one web server is shown, this may be scaled up to any number as needed.

None of the computing devices shown in FIG. 1 are directly coupled with remote server **122**, which may by non-limiting example be a third-party server, or multiple servers (such as a portion of a server rack) or any portion thereof. System **100** could, for example, exclude server **106**, and utilize only remote servers **122** which have access to the database **108** (which may be stored on the remote servers), and each of the computing devices may access the database through the remote servers and through one or more described web servers such as through one or more user

interfaces displayed on the displays of the computing devices when accessing correlated URLs.

In other implementations one or more application servers could be included in the system, the application server(s) positioned relationship-wise between an end-user device and the database(s) to facilitate operation of the methods that will be later described.

As indicated, these are only examples of how to implement a threat modeling system, and many other layouts are possible. System **100** may also include many other elements which are not shown for brevity. In a small business or organization wherein only one computing device may be needed to do threat modeling, system **100** could be implemented using a single computing device **102** with a database **108** stored thereon, or with one computing device **102** coupled with a server **106** and database **108** through a local connection (wired or hardwired), or using a cloud-stored database that the users access through user interfaces through remote servers **122** and/or web servers **120**. In an organization in which multiple computing devices may need to do threat modeling the system **100** could be implemented using a computing device **102** having the database stored thereon, or coupled through a local or Internet connection to a database **108** stored elsewhere within the organization's computing devices such as on a server **106** or remotely on remote servers **122** accessed via a web server **120**, with other computing devices **112** and/or **116** coupled either directly with device **102** and/or **106** and/or through the telecommunication network **110**. In implementations in which remote servers are utilized these may be scaled up to any needed number.

The threat modeling system and methods include the modeling of threats utilizing software which users access and interact with through a variety of user interfaces, some examples of which will be described hereafter, but a brief description of the processes facilitated by the software will now be discussed.

Referring to FIG. 2, a representative example of a threat modeling process (process) (method) **200** includes generating a threat model (model) **208** for any application, process, or system under consideration. By non-limiting example, this could include modeling the possible threats to commuting to work safely, modeling the possible threats to preventing the spread of an infectious disease, or modeling the possible attacks on a computer network (cybersecurity). Model **208** is used to generate an original threat report (report) **214** which in implementations includes identified threats, the status of identified threats (threat status), and the source(s) of identified threats, among other things.

As illustrated in FIG. 2, process **200** may include storing a plurality of threats **202** and threat model components (components) **204** in a data store **206**. This may include, by non-limiting example, storing titles, definitions or descriptions, and/or associated images in the database **108** for each component and/or threat. The threats and components are used to create threat model **208**.

Process **200** in implementations includes a user selecting from among the available components those components which are relevant to any given system, process or application. This is represented by the arrow between "components" and "relevant sources." The relevant sources **212** are relevant sources of threats that are determined by the system **100** either after, or while, the user is selecting the components that are involved in the specific application, system or process and defining relationships between and among the relevant components. The system also retrieves from the database relevant threats **210** that were previously correlated

with the chosen components or combinations of components through the database, to form the threat model **208** (this is representatively illustrated by the arrow between “threats” and “relevant threats”). The threat model thus includes relevant threats and the relevant sources of those threats. The threat model is used to generate a threat report **214**.

There may be some components with which no threat is associated (and they are therefore not relevant sources of threats), and there may be some threats that are dependent on one or more relationships between components. For example, when modeling a computing network some components may communicate with one another using a hyper-text transfer protocol secure (HTTPS) protocol or instead with a transmission control protocol (TCP), and this relationship may determine whether there is a relevant threat to include in the model (or which relevant threat to include in the model). Although these relationships between components are communicative couplings and/or protocols in some instances, the relationships themselves may be considered “components” in a broad sense so that, in FIG. 2, the potential relationship types between components would themselves be considered components **204** that are stored in the data store and which the user utilizes to build a diagram of the system, application or process. Communication protocols may use conductive wires (electric signals), optic fibers (optic signals), wireless technologies (electromagnetic signals), and so forth.

FIG. 2 is a simplified diagram. Each component and each plurality of components is a potential relevant source for one or more threats. For example, one threat may be “Bluejacking” and one component, which would be correlated to this threat through the database, could be “BLUETOOTH port.” Accordingly, if a user includes a BLUETOOTH port in a diagram of a computing system, the system **100** will identify that port as a relevant source for bluejacking in the associated threat model and threat report. The component in this example is a physical component of a computing device or system/network. In other implementations components/sources may not be physical components. For example, if one is modeling the threats involved in commuting to work safely one threat may be “freeway collision” and one component, which would be correlated with this threat through the database, would be “merging onto freeway.” Thus, in this latter example “merging onto freeway” would be a relevant source for “freeway collision.” In this example the component (and relevant source) is defined as an action or step, and not as a physical component.

Regardless of whether the threats and components/sources are physical elements or steps, actions, etc., the database allows for the correlation of any threat to any number of components and likewise, the correlation of any component to any number of threats. Accordingly, using the above example, the BLUETOOTH port may be defined as a relevant source for a Bluejacking threat, but it may not be the only relevant source (another source may be a BLUETOOTH device wirelessly connected with the system being analyzed). Similarly, Bluejacking may be defined as one threat that may occur due to vulnerabilities of the BLUETOOTH port, but there may be other vulnerabilities or threats that are present in the system due to the BLUETOOTH port.

Referring to FIG. 3, another implementation of a threat modeling process (process) (method) **300** is shown. Method **300** includes the steps described above with respect to FIG. 2 but also includes additional steps. A modified data store **302** includes data store **206** but also includes compensating controls **304** stored in the database. The stored compensating

controls include, by non-limiting example, a title, definition, image, and/or other items for each compensating control. Each compensating control may be associated with one or more threats and/or with one or more components and/or with one or more security requirements through the database (security requirements may in turn be associated with one or more components and/or one or more threats through the database). Method **300** includes user selection of one or more compensating controls (relevant compensating controls **308**) from among all compensating controls **304** stored in the database, and the relevant compensating controls together with the threat model **208** previously discussed (in other words the relevant threats **210** and relevant sources **212**) are included in the modified threat model **306**. Modified threat model is used to generate modified threat report **310**.

Referring now to FIG. 4, method **300** may include other steps and may be shown by threat modeling process (process) (method) **400** which includes creating a threat model of the application, system or process of interest (step **402**), generating a threat report based on the threat model identifying relevant threats (step **404**), generating a modified threat model with the status of one or more threats changed to “mitigated” using one or more of the relevant compensating controls (step **406**), and generating a modified threat report and/or a change log including the threats designated as “mitigated” by the one or more relevant compensating controls (step **408**).

Referring now to FIG. 5, another implementation of a threat modeling process (process) (method) **500** is representatively illustrated. Method **500** includes cycling through the process of selecting/deselecting a compensating control (i.e., adding it to the threat model, or removing it) as often as the user chooses. After the threat model is initially built and the initial threat report is generated (step **502**) the user may define a compensating control (step **504**) (this step may also be done before the threat model has been built or the threat report generated). The user may select the threats mapped to the compensating control and/or deselect threats mapped to the compensating control (step **506**).

The user may add the compensating control to the threat model and thereby modify the model (step **508**). System **100** matches threats identified in the original threat model with those mapped to the selected compensating control (step **510**). The user at this point may be able to further modify the threats mapped to the compensating control. The status of every matched threat may be changed to “mitigated” or some other status by the user (step **512**). A modified threat report may then be generated (step **514**). The user is then brought to a point at which he/she may select whether to revert the model to its previous state (step **516**). If the user selects to not revert the model, the process skips to step **522**, otherwise the compensating control is removed from the model (step **518**) and the status of matched threats are reverted to their pre-matched values (step **520**). The user then determines whether to add another compensating control (step **522**) and if so the process cycles to step **504** again, otherwise the process ends.

Reference will now be made to several example user interfaces which may be utilized to accomplish the above general processes and other processes as will be described. It should be stressed that these are only examples, and that other user interfaces could be used to accomplish the methods. Similarly, although specific user interfaces are described with respect to specific functionalities (dropdown menus, buttons, fields, tags, text prediction, etc.), the practitioner of ordinary skill in the art will be able to mix and

11

match these functionalities and/or use other functionalities with the user interfaces to make the user experience intuitive and easy. For example, in instances where a dropdown menu is present this could be replaced by a search field, or a radio button selection, and so forth. Nevertheless, the user interface implementations as depicted in the drawings are useful for a variety of reasons, as will be understood.

Referring now to FIG. 6, a representative example of a user interface (interface) 600 is shown. This interface is titled the “knowledge base” page and may be displayed on any of the displays of system 100 shown in FIG. 1, depending on the particular system setup. Interface 600 has a number of selectors or menu items near the top of the screen such as, by non-limiting example, a three-bar menu selector, a “select all” selector, a “new” selector, an “edit” selector, a “delete” selector, a “copy to library” selector, a “select library” dropdown selector, a “home” selector, a “knowledge base” selector, a “templates” selector, a “profile” selector, a “change password” selector, and a “sign out” selector. In implementations some of these selectors may be further organized into dropdowns so as to take up less space on the interface. Additionally, in implementations of system 100 some or all of the menu items may be present on other interfaces, such that they are permanent menu items at the top of those other interfaces as the user navigates from one interface/page to another.

This “knowledge base” interface is an interface where the user may view and edit “components” that are stored in the database. The title “components” is thus present at the left near the top of the screen, and below this are three columns for “name,” “type,” and “labels.” Below the column titles are search boxes where a user may begin typing and predictive text will bring up a list of items that begin with the letters typed so far. Each row shows one component type, name, and labels (if any). The list is scrollable as indicated by the scroll bar, and only a small number of components are listed here. For example, the topmost item is a WiFi router (named WIFI RTR in shorthand, and this and any other text used in the system may be in some type of shorthand or may be spelled out in its entirety in various implementations), its type is “Device,” and no labels are applied. Shown in FIG. 6 are the following types: device, IoT device (which represents “internet of things” device), communication protocols (such as HTTPS), generic data store, database, application component, and deployment component. Other available types not shown may include, by non-limiting example: cloud, external system, FTP server, generic external interactor, generic process, generic trust boundary, Modbus, perimeter control, third party software, third party library, transaction, WINDOWS application, WINDOWS OS, and so forth. These are only representative examples, and the user may define new types as desired and apply those types to a new component (or existing component), which will associate the type with the component name through the database. The rows may be organized in alphabetical order by any of the columns (name, type, labels) by clicking on the title of the column and may be organized in reverse alphabetical order by again clicking on the same column header (and toggled with continued clicking). The labels column may be useful for, among other things, searching for components that have a label applied. For example, if the user is looking for embedded system components, the user may search in the labels column for “embedded system” (or a shortened version if the system uses a shortened version) to find all components that have the embedded system tag applied.

12

A user may select any row by clicking anywhere in that row. The selected row will be highlighted, as representatively illustrated in FIG. 6 wherein the “login” row is highlighted. Once a row is highlighted any data associated with the component through the database is shown to the right. For example, in the middle of the screen the component name “login” and its type “application component” are shown. Below this the tags applied to this component are shown, which in this case include: web, application, feature, authentication, password. A description can be associated with the component through the database—in this case no description has been provided. If a user desires to edit the name, type, labels, or description the above “edit” selector may be selected and the user will be able to input new information for any of these fields, which upon pressing a “cancel” button will revert to the prior data or upon pressing a “save” button will write the new data to the database.

Interface 600 also shows a “properties” section to the right, which shows the threats and security requirements associated with this component through the database. The user may press the plus icons to add new threats or security requirements, as desired, or may use the “X” icons to remove threats and/or security requirements associated with the component through the database. In this example the login component has the following threats associated with it: inducing account lockout; session hijacking; and clickjacking. The security requirements associated with the component in this example are: parameterized queries—dynamic SQL; breach—randomizing secrets per request; and encryption algorithm (the first two only show a short description on the interface screen for ease of viewing).

The threats and security requirements associated with the component through the database will at least partially facilitate the later threat modeling when a user is depicting a specific system layout. For example, if a user is depicting a computer system layout and adds the “login” component to the system layout then one or more or all of the associated threats may become “relevant threats” to the overall system as previously described. In some systems 100 all the threats associated with this element would become “relevant threats” to the specific application, process, or system being modeled, and in other systems 100 some of these threats may be automatically dropped from “relevant threats” for the specific application, process or system being modeled if other components added to the system would inherently mitigate those threats (such as an HTTPS protocol instead of HTTP protocol being used between two components).

With regards to the other selectors, a user may select the “select all” option to select all components then listed (if the user has done a search this would limit to selecting all those components populated by the search), and the user may then press the delete selector to delete all of these components and their relationships (associated threats, security requirements, tags, descriptions, etc.) from the database. The user may instead, once all are selected, press the “copy to library” selector to copy the selected items to a default library set up by the user previously through another interface. The user may similarly select only one or a few components to delete or copy to the default library. The “select library” selector allows a user to display various libraries, here the “Threat-Modeler” library is displayed but the user may create his/her own libraries, may view other built-in libraries or third-party libraries accessible through the interface through a web server or other server which communicates with another third-party database, and so forth.

The user may select the “new” icon to bring up a popup interface in which the user may define a new component,

13

giving it a name, selecting the component type from a dropdown menu, adding an image to be displayed to represent the component (such as a lock for a security feature, a USB indicator for a USB device, etc.), labels, and a description, and may then select a “cancel” selector to cancel the new component or may select a “save” selector to write the new component information and relationships to the database, and selecting either save or cancel will also remove the popup interface and revert back to interface 600.

The home menu item at the top right of the screen will navigate the user to a home screen. The knowledge base icon will navigate to interface 600. The templates selector navigates to an interface similar to interface 1000 of FIG. 10 (hereinafter described) but with a blank canvas 1002 for creating a new template which may be used later for creating other relational diagrams (diagrams) 1004 (also described hereafter). The profile selector brings up a popup window which displays the current user’s name, email, department, role (admin, etc.), and last login date/time as stored in the database, and a close button to close the window. The change password selector brings up a popup window to change the user’s password with a “close” selector to cancel the change and a “save” selector to write the new data to the database. In implementations the home, templates, and knowledge base icons are permanently displayed at the top of all interfaces of system 100 and the profile, change password, and sign out selectors are available from a dropdown icon which is also displayed at the top of all interfaces of system 100. The sign out selector allows the user to log out of the system. Interface 600, as well as other interfaces of the system, may also include a refresh selector to refresh the page and a help icon to bring up a help menu. An icon displaying a title and/or logo of the software may also be displayed at the top of the screen (such as the top left) which when selected may navigate to the home screen.

When the triple-bar menu icon is selected it brings up a popup expandable menu which, when all items are fully expanded, appears as interface 700 of FIG. 7. The “threat library” menu item includes the sub-menu items “threats,” “test cases,” and “threat agents,” the “security requirements” menu item includes the sub-menu items shown, the “threat patterns” menu item includes the sub-menu items shown, and the “enterprise” menu item includes the sub-menu item “users.” When any of the sub-menu items are selected an interface very similar to interface 600 is displayed.

For example, FIG. 8 shows interface 800 which is accessed by selecting the “threats” sub-menu item. The title “threats” is shown near the top of a list of threats which are organized into rows and columns showing each threat’s name, risk level, and attached labels. Several threats are shown, and any column may be searched using the search box directly below the name, risk, or labels headers, which behaves similarly as described for FIG. 6, and the data may be organized alphabetically (or reverse alphabetically) by any column as described with respect to interface 600. The selected threat is highlighted, and its information is displayed to the right including the name, risk level, labels, and a description, which in this case includes a URL. The properties section includes threat agents which may be added, though none are added to this threat—but examples include things like “insider,” “black hat,” “hactivist,” “cyber terrorists,” “authorized external user,” and so forth which are threat agents that had previously been identified by the user. Threat agents may be added and removed. Test cases may also be added and removed, this threat has three test cases added which were previously entered into the database by the user. Security requirements may also be

14

added by the user and, in this way, compensating controls may be associated with specific threats through the database—this is one of the ways the system facilitates step 506 of FIG. 5. Security requirements may be added to the database originally through another interface, described hereafter, and identified there as a compensating control, then added to a specific threat from interface 800. Such an association through the database facilitates the system displaying various compensating controls in a diagrammed system, method or application and the threats those compensating controls can mitigate.

The menu items at the top of FIG. 8 are the same commands/links as those shown in FIG. 6 but are shown here in icon format to give another representative example, where the select all selector is represented by a checkmark, the new selector is represented by a plus icon, the edit selector is represented by a pencil icon, the delete selector is represented by an “X” icon, the copy to library selector is represented by an icon of two documents, the home selector is represented by a home icon, the knowledge base selector is represented by a file folder icon, the templates selector is represented by a document icon, and the profile, change password, and sign out selectors are available by selecting a dropdown menu represented by a downward facing triangle icon.

Referring back to FIG. 7, if the “test cases” sub-menu item is selected an interface similar to interfaces 600/800 is displayed, but showing test cases, and allowing the user to add, edit, delete, copy, add labels, and so forth, similarly as previously described with respect to the “threats” sub-menu item. The threat agents, security requirements, code snippets, code reviews, components (already described with respect to interface 600 since the components interface is set as the default interface when the “knowledge base” selector is selected), data elements, roles, widgets, component types, attributes, and user sub-menu items all have similar functionality and bring up similar interfaces when selected. Data elements, which represent data elements that may be captured by any diagrammed system, application or process (such as credit card numbers, billing addresses, pins, phone numbers, email addresses, order history, birth date, medical history, insurance history, and so forth) may be associated with specific threats and with specific security requirements in the same way as those associations may be made for components.

If a user selects the security requirements sub-menu item an interface similar to interfaces 600/800 will be shown similar to interface 900 of FIG. 9, which allows the user to order security requirements alphabetically (or reverse) by name or label and search for specific security requirements. Selecting any security requirement will highlight it and display its name, description, labels, an indicator to indicate whether the security requirement is a compensating control, and any code snippets or code reviews (under a “properties” header) that have been associated with the security requirement (which may be added or removed from that interface similar to what is described with respect to the “properties” elements of interfaces 600/800).

Interface 900 is actually the interface visible when the user selects the “new” selector to create a new security requirement. From this window the user may add a name, add a description, check a box (or leave it blank) to indicate whether the security requirement is a compensating control, add any labels (which are previously input into the database by the user), and press cancel to abort the addition or press save to write the data to the database.

15

From any of the previously mentioned interfaces, if the user presses the home selector an interface similar to interface **1300** of FIG. **13** will be displayed, which lists all previously stored threat models by name and showing their version. Either column may be organized alphabetically or in reverse, and the PDF button will export the list to PDF. If any specific threat model is selected it will be highlighted and its associated threat report (threat report interface) **1302** will be displayed, which will be discussed hereafter. From the top menu items the user may select the new selector to create a new threat model, the edit selector to edit the name, version, risk level, an "internal" toggle, and labels associated with the selected threat model, a delete selector to delete the selected threat model, a diagram selector to view the diagram for the selected threat model, a report selector to export to PDF the threat report (which shows for each threat the threat name, source, risk level, status, and creation date), a threat tree selector to view a diagrammed threat tree, showing threats of the threat model, and other selectors already described.

If the threat tree selector is selected a threat tree is displayed in which threats are organized as sub-elements of data elements and/or components, and also displayed are mitigating security requirements or compensating controls that may be implemented to mitigate the threats and an indicator of whether each threat has been mitigated. Portions of the tree may be collapsed, expanded, or viewed in vertical tree or horizontal tree format. The interface showing the tree diagram has selectors to zoom in, zoom out, revert to 100% zoom, toggle on/off a thumbnail overview image in a corner of the display, save an image of the tree, open a legends window which shows the colors in which various elements are diagrammed (threat agent, widget, component, role, data element, security requirement, threat, protocol, node, and project), and a filter selector which allows the user to remove and add back in any of the aforementioned items from the tree. The tree may be dragged and moved in any direction for viewing any portion of it easier.

Continuing with FIG. **13**, if the new selector is selected (which, as with all other selectors, could be implemented as an icon or logo absent any wording), a popup window appears allowing the user to enter a name, version, toggle an "internal" indicator, select a risk level from a dropdown (from among previously determined risk levels stored in the database), apply labels (which allow the user to type and populate a list of already entered labels stored in the database to select one, or to add a new label by typing a new item and selecting "add new"), or the user may select a "new from template" button to begin a new diagram from an existing diagram template, or a "new from features" button to begin a new diagram from user-selected features, or the user may select "import" to begin a diagram from a previously stored diagram (such as from an older version of the software), or the user may select "empty" to begin a diagram with a blank canvas.

If the user selects "new from features" the diagram interface of FIG. **15** displays with a blank canvas **1002**, and a popup selector (not shown) allowing the user to select "business requirement" features which are desired to be included in the model such as, by non-limiting example, authentication, funds transfer, credit verification, bill pay, or other features already stored in the database previously by the user, and after selecting the desired features, the blank canvas will populate with a node for each feature. The user can then modify and/or add components to the diagram as desired. In implementations each of these added "business

16

requirements" will add a "component group" to the diagram and threat model, as will be described hereafter.

If the user selects "empty" the diagram interface (interface) **1000** of FIG. **15** will appear, showing the blank canvas **1002** as shown. To the left of the canvas is a toolbox module which may be minimized using the left facing double arrows and re-expanded by selecting the same arrows which will then be right facing. To the right of the canvas is a properties module which also may be minimized and expanded in similar fashion, along with sub-menu items "general," "components," "data elements," "roles," and "widgets" each of which may be expanded downward or minimized upward using the double arrows (in FIG. **15** they are all minimized).

The toolbox module is searchable, with the user able to begin typing and the system suggesting through a list populated just below the search box components which begin with the letters (or include the letters) typed by the user. The dropdown to the right of the search box may be used to allow the user to search from among all components or a subset such as one or more specific component libraries as previously discussed, with ThreatModeler being an example component library) and/or one or more specific component types as stored in the database (and previously discussed with respect to FIG. **6**). The user may also scroll down through the components using the scrollbar function.

In the example of FIG. **10** the components in the toolbox are of varying sizes, but in implementations they may all be of similar sizes and shapes (or identical sizes and shapes). In implementations in which the system is deployed within a single organization with only a local library of components the components listed in the toolbox will be those that have been input by the user or that were previously loaded into the database during software installation. In implementations in which the system includes some remote communication with other libraries the toolbox may display components available from other parties, such as through a cloud computing services, e.g., MICROSOFT AZURE or the like (in implementations the entire system and methods could be implemented using cloud computing in instances where a local software installation or local database are not desired).

The toolbox, as can be seen, includes components that were previously entered into the database through interface **600** as previously described. From interface **1000** the user may select a desired component and drag it to a desired location on the canvas, and by dragging multiple items the user may begin to add the components of a system, application or process to be modeled. By non-limiting example, FIG. **10** shows interface **10** with a relational diagram (diagram) **1004** already fully created by the user. As can be seen, the user has dragged multiple components onto the canvas **1002** and has defined various relationships between them. The user dragging components to the canvas is one way in which the user selects relevant components or relevant sources as described with respect to FIG. **2**, and it is from these relevant components (relevant sources) that the system identifies relevant threats by identifying the threats that are associated through the database with those components.

In implementations a communication protocol (protocol) can be defined between components by clicking on a component and dragging the cursor to a second component. This will create an arrow, such as those shown on FIG. **10** between components on the canvas, and will display a default protocol (here the default protocol is HTTPS), but a user may right click the arrow to display a popup selector allowing the user to deselect HTTPS and/or select one or more other protocols (such as TCP, SMB, WiFi, 3G, 4G, AJAX, binary, BLUETOOTH, FTP, FTPS, HTTP, IMAP,

17

MAPIRPC, SMTP, SMTPS, USB, etc.), and when any protocol is selected that protocol name will be displayed (or if multiple are selected then “multiple” or some term/icon meaning “multiple” will be displayed—in the image this is shown as “MULT” for short)—these protocol arrows will also be color coded by type (such as green for HTTPS, red for multiple, and so forth). Protocol arrows may also be deleted as desired. The alphanumeric descriptions are shown here next to the link arrows, but in implementations the wording may be superimposed over the link arrows. They are shown here next to the arrows for easier readability.

The components displayed in the toolbox of FIG. 10 are generally shown without icons. The USB Port, SSO (single sign on), BLUETOOTH PORT, and BLUETOOTH PROXY components are shown with icons as an example that all components may be shown in the toolbox with an icon if desired. It may also be seen that the user may input components that constitute branded software elements, such as PAY PAL, AKAMAI DNS, SKYPE, etc. In implementations all compensating controls will be displayed in the toolbox with a similar icon, such as the padlock icon shown for the SSO component of FIG. 10.

As a user adds components and links them together, the threat model includes the threats that are associated with the components/protocols. As previously described, each component, each protocol, and each data element may be directly associated with one or more threats through the database. Accordingly, each time a component is added to the diagram, or a new link drawn between components and protocol selected, if there are any threats associated through the database with the component and/or protocol those threats are then included in the threat model. A user may at any time right click on a blank area of the canvas when no component/link is selected and a “threat report” item will be selectable which, when clicked, will generate a popup threat which will include a list of all threats. In implementations this will look much like the threat report 1302 shown in FIG. 13, and will include a first “threats” tab list of all threats (and the number of threats), identifying the risk level of each threat, showing the threat status (open/mitigated), identifying the source of the threat (which is a component on the diagram), and having an actions column from which the user may display a description of the threat (previously associated with the threat through the database) (the description selector here shown with the letter “D” though the selector could instead show “description” or the like) and leave a comment to store to the database (the comment selector having the phrase “note” though in other implementations it could say “notes” or “comment” or the like). The user may organize the table by any column (alphabetically or reverse) and may also organize by column headers by dragging the column headers to the location indicated. For example, the user could drag the “status” header to the identified location and then drag the “risk” header to the right of it and the table would then be organized first by status and then by risk level. When this “dragging” is done a “ghost header” is dragged up to the identified location in the sense that the existing header remains in place. The organization may be reverted to an earlier state by deleting the ghost headers using an “x” icon.

The threat report displayed also includes a security requirement tab (showing number of security requirements) and displaying a table similar to the threats table. The security requirements table lists the security requirements that are associated with one or more of the identified threats in one column, lists a source in another column (which is the component associated with the threat), includes an “implemented” column indicating whether the security require-

18

ment has been implemented, an “optional” column indicating whether the security requirement is optional, and a similar “actions” column to the previously described actions column, this time the description giving a description of the security requirement and allowing the user to record a written note. This table may be organized similar to the previously described table.

With respect to the “implemented” column, in implementations this is a checkbox for each security requirement indicating whether the security requirement has been implemented. The user may manually check those security requirements that have been implemented and manually uncheck those that have not. This column may also have been previously populated with some “implemented” checkmarks based on mitigations that the user has selected for specific compensating controls, which will be discussed later with respect to FIG. 11. Further, the “optional” column may also include checkboxes which a user may manually toggle to indicate which security requirements are optional or not. This column also may be populated by the user selecting certain mitigations as will later be discussed with respect to FIG. 11—for example if there are two security requirements associated through the database with a specific threat, but only one needs to be implemented, then if the user selects one of those security measures as mitigating the threat the other security measure may populate on this list as being checked “optional.”

The threat report further includes a “test cases” tab (showing the number of test cases) and displays a table listing test cases associated with the threats through the database. As previously described with respect to FIG. 8, each threat may have test cases associated with it. An example test case would be, for example, a “dictionary-based password attack” for a password-related threat. The test case tab would list this test case and have an “actions” column which, when an icon is selected, pops up a description, which in this case states “use a password cracking tool what will leverage the dictionary to feed passwords to the system and see if they work” and further list techniques like setting the test to try all words in the dictionary, common misspellings, and combinations of words and common misspellings of combinations. The test case tab/table thus offers tests that the user may try to test against the threats both before and after security requirements are implemented. This table may be organized and sorted similar to the previously described tables.

The threat report further includes a “code reviews” tab (showing the number of code reviews) and displays a table listing the code reviews that are associated with the security requirements through the database (as previously shown with respect to FIG. 9). An example code review would be, for example, an “Authentication: Weak Password” code review. The code review table would list this code review and have an “actions” column which, when an icon is selected, pops up a description, which in this case states “Password strength should be enforced upon a user setting/selecting one’s password” and gives examples of code that may be used to set up such an enforcement. The code review tab thus offers sample code that the user may use to implement measures against threats which may complement or work together with implemented security requirements. This table may be organized and sorted similar to the previously described tables.

Multiple selectors are shown at the top of interface 1000 in addition to the permanent selectors that are available on several interfaces. The “select all” selector selects all components on the canvas (the user may then deselect some, if

desired, for example for grouping some or all the components as later described). The “copy” selector copies selected components and the “paste” selector pastes the copy onto the canvas (this may also be done with keyboard shortcuts, and shortcuts may in implementations be used for all other selectors described herein for system 100). The “delete” selector deletes the selected components from the canvas and the “clear canvas” selector deletes all components from the canvas. The “zoom in” and “zoom out” and “100%” zoom in and out of the canvas and revert to a 100% zoom, respectively. The “comment” selector populates an empty comment box on the canvas in which the user may type a comment which will remain on the canvas until later removed (though it may appear in minimized or icon format until hovered over or opened). The “overview” selector displays a thumbnail overview of the entire diagram in a corner of the canvas.

The “undo” selector undoes the last action (or multiple last actions if selected multiple times) and the “redo” selector does the opposite. The leftmost “templates” selector pops up a menu from which one or more previously saved business requirement “features” may be selected to add to the canvas (these may be nodes and/or component groups, for example, and may be the same nodes/component groups that are described above when the user uses the “new from features” function when generating a new model). The “save as image” selector has a dropdown and allows the user to save an image of the diagram as then displayed in one of various formats such as PNG, SVG, etc., or to take a “snapshot” which saves the diagram at the present state to the database. The “PDF” selector generates a PDF of the threat report based on the current configuration of the diagram components (the threat report described to some extent above). The “settings” dropdown selector allows the user to toggle between straight or curved link lines (protocol lines/arrows), showing or hiding the link wording (in FIG. 10 the lines are curved and the link wording is shown), making the diagram public or non-public, and showing or hiding comments (the comment function described above).

The “save as template” selector allows a user to save the entire diagram as either a threat model or a business requirement “feature” (in other words saving it as a component group) and in either case the user may give the feature a name, add labels, and select the type as either threat model or business requirement, then press cancel to cancel or save to store the new template to the database. The “snapshots” selector populates a list of previously saved snapshots, any one of which may be opened from the list or deleted.

The “group” and “ungroup” selectors allow the user to create a group containing multiple components or to delete a group (but not the included components). As seen in FIG. 10 for example, there is a “WINDOWS 7” group which includes file system, PDF client, SKYPE, OUTLOOK 2010, MCAFEE HIPS, MCAFEE AV VSE, IE11, BIT LOCKER, MCAFEE DLP, and OFFICE 2010 components. There is also a Laptop group containing the WINDOWS 7 group and further containing WiFi port, USB port, ethernet port, HDMI port, and BLUETHOOTH port components. Finally, there is an Office Network group which includes the Laptop group and also includes file server, SHAREPOINT, printer, WiFi access point, IRONPORT, email server, BLUECOAT proxy, and SSO (single sign on) components. Then there are other components (external email gateway, AKAMAI DNS) which are not part of any group. A group may be formed from any one or more components, and the AZURE group is seen containing only a single component: ONEDRIVE.

When a group is formed the user may, using the right-side “general” dropdown, add a custom display name, select a container type (from container, trust boundary, collection, or some other option stored in the database previously by the user), select a component (for example for the WINDOWS 7 the component “WINDOWS machine” is selected, which shows that some components stored in the database may include other components), select other display choices such as title and background/border color, and select a Common Platform Enumeration (CPE) identification from among a list previously stored in the database (for example in this case ID a version of WINDOWS 7 is selected) (the list of CPE IDs may be searched using filters to easily find the appropriate one), and the user may also add notes.

As further examples, the WINDOWS 7 group is identified as a container group, a WINDOWS machine, and a specific WINDOWS 7 operating system version is identified. The Laptop group shown in FIG. 10 is listed as a container group and no CPE ID is selected. The Office Network group is identified as a trust boundary group type and no CPE ID is selectable. The AZURE group is identified as a trust boundary group type and no CPE ID is selectable.

Each grouping of components, however, could be diagrammed separately as an independent threat model and then saved as a component so that it may be imported into another threat model/diagram. When a user adds any of these component groups to a blank or existing diagram/threat model the threat model of the component group is added to (and/or nested within) the threat model of the existing diagram/threat model. In this way the user can modify a threat model by incorporating previously defined threat models. This ability is generally termed “threat model chaining” herein and is a useful mechanism for allowing a user to diagram complex systems/processes without having to repeatedly build common elements among the systems/processes.

Each component group may thus be redefined as a discrete “component” and may then be included as a single icon in the toolbox menu. By non-limiting example, referring to FIG. 10, the WINDOWS 7 component group could be defined as a component, then the user could, in another diagram (or the same diagram), select and add a WINDOWS 7 component to the diagram to import into the diagram and associated threat model the threats associated with the WINDOWS 7 threat model. That same could be done for the LAPTOP component group. Accordingly, a component group and associated threats added to a diagram may in turn already include other nested/chained threat models therein, so for example if a user defined a laptop component group such as that in FIG. 10 as a “WINDOWS LAPTOP” component then, when a user later adds a WINDOWS LAPTOP element to a diagram/threat model by selecting a WINDOWS LAPTOP component from the toolbox and dragging it onto the diagram, the threats associated with the laptop itself, as well as the nested/chained threats associated with the included WINDOWS 7 threat model, are automatically included in the threat model for the then displayed diagram.

Referring back to FIG. 10, the “compensating controls” selector may be selected to display a popup similar to compensating control report (report) 1200 of FIG. 12. This report is populated from previous relations stored in the relational database. As indicated previously with respect to FIG. 6, each component may be associated with one or more threats and one or more security requirements through the database using interface 600, and as described with respect to FIG. 8 each threat may be associated with one or more

21

security requirements through the database. As further indicated with respect to FIG. 9, each security requirement may be identified through the database as a compensating control. Accordingly, based on these relationships and selections the list populated in FIG. 12 shows all possible compensating controls that could be implemented to mitigate threats that are present in the diagrammed system, application or process.

By non-limiting example, the threat of “physical theft” is associated with the component “laptop” through the database, but no security requirements are directly associated with the laptop component. Nevertheless, the “physical theft” threat is associated with the security requirements “MCAFEЕ FRP” and “BITLOCKER” through the database, both of which security requirements are identified through the database as compensating controls (MCAFEЕ FRP being a USB encryption tool and BITLOCKER being a disk encryption tool). Accordingly, when a user adds the laptop component to a diagram, if the user selects the “compensating controls” selector this list will then include both MCAFEЕ FRP and BITLOCKER as compensating controls for the physical theft threat. If the user removes the laptop component from the diagram (deletes it) then the MCAFEЕ FRP and BITLOCKER compensating controls will no longer appear in the list (unless they are also associated as compensating controls for some other threat present in the diagram). In implementations the list populated will also show compensating controls which are directly associated with a component that is added to the diagram. In other words, in the above example there are no security requirements directly associated with the laptop component, but if a third security requirement were directly associated with the component through interface 600, and if the third security requirement was identified through the database as a compensating control, then that third compensating control would also populate in the list of interface 1200.

Referring again to FIG. 10, it may be seen that this diagram includes the MCAFEЕ DLP, BITLOCKER, MCAFEЕ HIPS, and MCAFEЕ AV VSE components grouped in the WINDOWS 7 group. The user may, for example have modeled the WINDOWS 7 group, then used the compensating controls selector at the top of interface 1000 to identify that there are certain threats that these components would mitigate. The user may then add those components to the WINDOWS 7 group by dragging them from the toolbox to a location within the group outline to add them to that group. Then, upon right clicking on any specific security requirement component, the user may select a “mitigations” selector which pops up mitigations interface (interface) 1100 as shown in FIG. 11. The user could add one of these components at a time and model the mitigations, generate a new threat report, then either remove that component or add another compensating control, and in this way test out various compensating controls. This is one way in which the system facilitates steps 512, 514, 516, 518, 520, and 522, among other steps, of FIG. 5.

Interface 1100 displays a table which lists all threats and their associated sources (component associated with the threat), risk levels, and status, and highlights the listed threats which the security requirement is configured to mitigate as identified previously through the database through interface 800 (these threats are “mitigatable” by the selected security requirement). The user may deselect highlighted threats and/or may highlight other threats, then may select “close” to cancel or may select “mitigate” to toggle those threats to mitigated through the database. Once this is done, for example, the threat report as seen in FIG. 13 (or

22

popup shown from the diagram interface) will list the mitigated threats as “mitigated” and the compensating controls list shown in FIG. 12 will also show the mitigated risks as mitigated.

The threat report of FIG. 13 allows a user to manually change the risk level and threat status directly from the home menu using the “change risk” and “change threat status” drop downs once a threat is selected by clicking on any row. Changing the risk level of a threat from this interface will only change the risk level for this specific threat model. Threat statuses that may be used include, by non-limiting examples: open, closed, mitigated, fixed, not applicable, needs more details, not tested, secure usage practice, and so forth.

The systems and methods described herein may also be used for analyzing an attack surface of a complex system or process, which will be described now. In implementations the attack surface of a system or process is comprised of the sum of all open and unmitigated potential threats to an “asset” identified through threat modeling.

As described herein, a modeled system or process may include a modeled deployed or contemplated computing network, which could include by non-limiting example: one or more applications; one or more on-premises infrastructures; one or more cloud-based infrastructures; one or more hybrid infrastructures; serverless architectures; microservices; one or more embedded devices; one or more IoT devices; one or more mobile devices; one or more Industrial Control Systems (ICS); one or more cyber-physical systems (CPS); one or more third party systems; one or more organizational networks or intranets; non-cyber elements interacting with a cyber system such as, by non-limiting example tangible assets, intangible assets, property(ies), plant(s), equipment, liquid assets, brands, reputation, residential structures, real property, utility services, unattached removable items and assets; an infrastructure system, such as a transportation infrastructure, an electrical grid, a telecommunications network, and so forth, all in implementations accessible and utilized through a highly interconnected network of intranets, wireless networks, and the Internet.

All such components of a modeled process or system may include potential threats which, if discovered and exploited by an adversary, may yield one or more attack vectors to one or more assets of the system or process. The “asset” may be any item selected by a user to be identified as an asset. The modeled “attackers” included in any threat model may be actually modeled as a person in implementations, though in other implementations may be simply modeled as an entry point or device included in the process or system which an attacking person may use to cause harm.

It may be pointed out that increased interconnectivity of a computing system with other systems (such as the Internet, third party systems, end user systems, etc.) may increase economic value and efficiency though these may also increase organizational risk due to the increase in adversarial actors and a constantly evolving threat landscape. The threat modeling chaining and attack simulation systems and methods described herein allow organizations to manage threats at a comprehensive organizational level notwithstanding an ever-changing threat landscape.

A modeled “attack” as used herein is a modeled attempt by an adversarial entity to traverse the modeled system or process from an attack surface to one or more assets within the system or process. The routes from all potential attackers to a selected asset are the “attack vectors” or attack paths to that asset.

23

Referring now to FIG. 14, a block flow diagram illustrates steps that may be included in a general attack surface analysis. Step 1402 includes user generation of a diagram of a system or process, as has been described previously with respect to FIG. 10. Step 1404 includes system generation of a threat report for each component or component group and step 1406 includes system combination of the individual threat reports into a comprehensive threat report for the overall diagrammed system/process (this would include, for example, including threat report elements for nested or chained threat models, as has been explained above), and then steps 1408-1414 include steps which may occur in any order. In step 1408, once a user has selected an asset to analyze, the system summarizes data to show all attack vectors associated with threats which may compromise that asset. At step 1410 the user analyzes the various attack vectors to determine what compensating controls may be included to protect the asset. At step 1412 the user adds or removes compensating controls to the diagram and/or toggles compensating controls between ON/OFF states. At step 1414 the user determines the effectiveness of the compensating controls or other risk management methods (such as changing communication protocols, changing the relative location of the asset within the modeled environment, adding non-compensating control elements between the asset and attack locations, and so forth).

FIG. 16 further breaks down step 1402. At step 1602 the user identifies the components of the system or process and chooses their relative placement within the diagram, at step 1604 the user identifies the type and placement of any communication protocols between components, and at step 1606 any other attributes of components (including protocols) may be further designated (this step may be excluded in some methods, and these steps may be done in any order in implementations).

FIG. 17 shows a tree diagram that representatively illustrates a method of determining an attack vector for a selected asset. After a system or process is modeled using the diagram interface, described above, then in the diagram interface or in another interface either an asset is selected or an associated component is selected and relevant threat attack paths illustrated. For example, in the topmost node of FIG. 17 an asset 1702 is shown, and the attack tree or threat tree shows that the asset is associated through the database with four components. Components 1704, 1706, 1708, and 1710 are each associated with threats (i.e., the threats that are associated with each component through the database). Accordingly, threats 1712, 1714, 1716, 1718, 1720, 1722, and 1724 are shown. It may be seen that the system then determines that threat 1722, through component 1708, is a threat to asset 1702, and thus attack vector 1726 (shown by lines of heavier weight) is shown between threat 1722 to asset 1702. The diagram of FIG. 17 may be shown on an interface of the system, though in implementations FIG. 17 simply models the method the system is utilizing behind the scenes to show the attack vector(s) in the diagram interface, as in some implementations (such as those shown in the drawings) the threats are not explicitly shown on the diagram interface(s). FIG. 17 shows an attack vector by which an adversarial entity may target the asset, i.e., by exploiting threat 1722 of component 1708. It should be noted that the asset may be some sub-component within a diagrammed component, or it may be some element that is somehow associated with the component (for example the component may be a database, and the asset may be sensitive client information stored within the database—in that instance the asset could be defined broadly as the database, or it could be

24

defined narrowly as the sensitive information within the database, other information in the database not being included in the definition).

It is noted that each asset could have any number of threats that could threaten to compromise it. Accordingly, the system identifies all threats which may compromise the asset in order to visually diagram attack vectors for the user. Some assets may be threatened by only a single threat, some may be threatened by two or three threats, or more, and so forth. Along these same lines, it is pointed out (as has been discussed to some extent above) that each individual component (or grouped set of components) of a threat model could, itself, be associated with its own threat model through the database. Because of this, the overall threat model that is shown (for instance in FIG. 13) in implementations could be called a threat model portfolio as it includes all sub-threat models and nested threat models. For example, an overall threat model A could include components B, C, and D. Component B could be a single component, component C could be a previously modeled group of components having its own threat model, and component D could be a previously modeled group of components having its own threat model that also includes therein a nested threat model for a component group E. Accordingly, the threat model A would include all threat models associated with components and component groups B, C, D, and E, including all nested threat models. As described above, this “threat model chaining” may allow for quick and simple building of process/system models without having to recreate commonly included system/process elements.

In implementations one or more interfaces of system 100 may be utilized to list the top potential threats (such as the top ten potential threats), the top most vulnerable components, the composition of threats by risk type, the composition of threats by status (mitigated, unmitigated, open, etc.), and so forth. The data may of course be organized and displayed in many other ways. In implementations an interface could list each threat, list each threat risk level, and list the source(s) for each threat. In implementations each threat source listing may include a component and an underlying source, so for example one component may be associated with multiple sources of a threat. An example would be an online banking interface which is shown on a diagrammed threat model as a component, and a threat may be a virus being introduced to the bank’s environment through a breach in the banking interface, and one source of the threat may be a legitimate user who unwarily exposes the ATM to a virus on his/her computer, and another source of the threat may be an individual hacker seeking to introduce the virus to the banking environment, and another source of the threat may be a state-sponsored entity with high end mechanisms to seek to introduce the virus to the banking environment.

In implementations an interface of the system may list the most vulnerable components in ascending/descending order, along with the risk level for each listed component, the number of total threats to that component, and the number of open or unmitigated threats for that component. In implementations an interface of the system may list multiple data elements associated with components (for example user names, social security numbers, financial info, credit card numbers, employee ages, etc.) (in this case the component likely being a database housing the data), a data classification for each data element (such as confidential, restricted, public, etc.) and an exposure level for each data element (such as very high, high, low, etc.). In implementations an interface of the system may break down the risk levels by percentage, such as 45% very high risks, 25% high risks,

25

15% medium risks, and so forth. In implementations an interface of the system may break down risk status by percentage, such as 45% open threats, 25% closed threats, 30% mitigated threats, etc. In implementations an interface of the system may visually illustrate the number of threat introductions plotted versus time, such as the number of threat introductions by month.

Any of the interfaces discussed herein may, as would be expected, change (and in implementations may change dynamically) as edits are made to the diagram (for example when a compensating control is toggled between ON/OFF states, this may change many or all of the interface displays accordingly by varying percentages, mitigating threats, etc. Thus when a compensating control is introduced and/or turned on or off, the threat model, threat report, mitigations interface, compensating control report, and so forth may all change accordingly to reflect the updated configuration.

It is pointed out that any threat model of the system may be subjected to an attack vector analysis. So, for example, an attack vector analysis may be performed on a simple two-component diagram, or an attack vector analysis may be performed on a hundred-component diagram that includes the two-component diagram as a nested sub-threat model.

In implementations an interface of the system may show a display similar to FIG. 17 except mapping all components/assets vulnerable from a single threat. In other implementations an interface of the system may show a display of all threats affecting a single component or asset. In the latter case the visual display may show attack vectors branching outwards from the asset, tracing communication protocols and extending through multiple components as far back as the locations/sources of the underlying threats. This is representatively illustrated in FIG. 18, which shows a type of display that may be shown on an interface of the system.

FIG. 18 shows components A-G. Component A has been selected as an asset, and the attack vectors of all threats threatening the asset (component A) are diagrammed using arrows. The arrows in FIG. 18 are shown in different formats so that they can be distinguished from one another. In other implementations they may all have the same appearance. The arrows are also shown pointing away from the asset, but in other implementations this may be reversed so that arrows point towards the asset from the underlying threats. It may be seen that a threat emanating from component G threatens to compromise component A (or the asset within or associated with component A), the attack vector 1802 passing through component F and component B to reach component A. A threat emanating from component E also threatens to compromise component A, the attack vector 1806 passing through component B to reach component A. A threat emanating from component D also threatens to compromise the asset, the attack vector 1804 passing through component C and component B to reach component A.

The attack vector displays may be modified based on contemplated compensating controls. FIG. 19, for example, shows the diagram of FIG. 18 but modified to include six compensating controls (each labeled "CC"). These compensating controls may mitigate different threats emanating from different components. Compensating control 3 (CC3) may mitigate one or more threats from component D. CC2 may mitigate one or more threats from components C and/or D. CC5 and CC6 may each mitigate one or more threats from component E. CC4 may mitigate one or more threats from components F and/or G. CC1 may mitigate one or more threats from components B, C, D, E, F, and/or G. Using such a modeled attack vector diagram, the user may be able to alter compensating controls and make decisions about

26

where in the system to deploy compensating controls. The user may determine, for example, that CC1 may mitigate many threats more efficiently than several compensating controls deployed elsewhere in the system. Or it may be determined that the asset is better protected by some other combination of compensating controls.

The elimination of any threat of course protects downstream components. Accordingly, for example, if CC3 eliminates or mitigates one threat emanating from component D, then the downstream effects of that threat to components C, B, and A are all mitigated. On the other hand, if CC1 eliminates or mitigates the same threat, it would only eliminate or mitigate the threat for component A, not for components B or C. Accordingly, there are security advantages to having compensating controls further upstream.

Changing the compensating controls in such a diagram could also change the other displays on other interfaces, for example the listing of top threats, listing of top ten most vulnerable components, other top ten listings, composition of threats by risk, composition of threats by status, composition of threats by source, composition of threats by other criteria, the overall threat diagram model, the mitigations interface, the compensating control report, and so forth. This dynamic change of all interfaces allows the user to do "what if" analyses by adding/removing any compensating control and/or toggling any compensating control between ON/OFF states and then seeing how all the aforementioned displays/reports etc. change accordingly. The change may, for example, illustrate one or more threats as mitigated or closed instead of open or unmitigated, reduce the total number of open or unmitigated threats, reduce exposure levels of data elements, alter the percentages of threats by risk level (very high, high, medium, low, very low, etc.), alter the percentages of threats by status (open, closed, unmitigated, mitigated), or the reverse of all these, on any of the interfaces or reports of the system.

The term "asset" as it is used herein may refer to anything that has value to a user or organization, and therefore requires security measures to protect it from theft, manipulation, destruction, compromise, or other forms of abuse. Attackers are an active population of computer environment users and vary in the skills, toolsets, opportunities, financial backing, and other attributes required for a successful attack on an asset. Most security measures to protect a computer system-related asset may be categorized into one of two categories: (1) defensive and/or monitoring mechanisms to inhibit attacks in progress or mitigate the effects thereafter, and; (2) reducing the number of potential threats through which attacks may be initiated to begin with. The latter can in many cases be less costly overall, and the ability of the systems and methods described herein to test and retest various configurations, including modeling the use of compensating controls, allows users to focus efforts on this method.

In implementations threat modeling and/or attack simulation methods may include the following steps: (1) generally defining a threat model (including model type, outputs to be generated, and overall objectives for different stakeholders); (2) visually diagramming a system/process using database-stored components (including communication protocols and compensating controls); (3) classifying one or more data elements within a component (or associated with a component), and/or one or more components themselves, as assets; (4) identifying and enumerating potential threats to the assets based on the diagrammed components; (5) analyzing the identified threats (threat analysis); (6) identifying and analyzing potential attackers based on the threat analy-

sis; (7) analyze the effectiveness of deployed or contemplated compensating controls to mitigate the potential threats; (8) analyze a threat model and attack surface as modified by compensating controls; and (9) measure and quantify the effectiveness of the model and/or method relative to the desired outputs and objectives previously determined. This is just one example of steps included in such methods, and other methods may exclude some of these steps, or include other steps, and so forth. In an implementation of the above steps, step (6) interrupts step (5), and in an alternative implementation step (5) is allowed to be completed without interruption. Furthermore, in implementations steps (7) and (8) may be iteratively bounced between during any given implementation, so that a user adds a compensating control then checks the outcome, toggles the configuration of the compensating control then checks the outcome, removes a compensating control then checks the outcome, moves a compensating control to a new location then checks the outcome, etc., before moving to step (9).

Referring to step (1), different stakeholders may have different threat model output needs, so different model types can be built using system 100 and the methods described herein. For example, software application developers, operations teams, embedded system teams, IoT device teams, cloud deployment and architecture teams, ICS system or cyber-related system teams, and so forth, may create different threat model types. Architects may model threats to a residential structure, structural engineers may model threats to occupants, neighborhood planners may model supporting infrastructures needed for developments, and so forth. Home decorators may model occupant flow and so determine the most secure places for expensive home décor. In non-computing systems/process that are modeled, the protocols may represent things other than communication for example in the home décor example the protocols may represent methods of moving an occupant between rooms. In implementations a user may select one of many model types to begin with (to populate the proper predefined components with which to diagram the model). This could be selecting, by example, from among examples such as the following: application threat model; operational threat model; cloud threat model; embedded threat model; IoT threat model; cyber-physical system threat model; and so forth. A user may still access all components in the diagram interface by searching for them but selecting the proper model may allow the most relevant components to populate at the top of the list.

Still referring to step (1), a user may in implementations also select role-based threat model outputs (such as with yes/no selectors) such as: secure coding requirements; operational checklists; executive metrics; executive reporting; security metrics; security reporting, and so forth, and these may also affect the position of components within the toolbox list on the diagram interface. In other implementations these selections (and the selections in the above paragraph) may only be used for record-keeping purposes. The input may also include the name of the model, the version, the name of the model creator, the date created, the update person, the update date, comments, desired objectives by development/operations/security/executives/business leaders, and others, and notes or toggle elements to indicate whether those objectives are met, in progress, or in some other state (and any or all of these features may be included on one or more of the already-described interfaces).

A representative example of the aforementioned step (3) is illustrated in FIG. 20, which may involve the classification of assets which may be associated with, located nearby or

relevant to, or utilized by diagrammed components. In implementations the asset classification for each asset includes: an asset class; an asset, and; an associated component. In implementations the asset classification may be used as an input for analyzing potential attackers by providing insight to possible purposes attackers may have in targeting the asset. As an example, in implementations of cybersecurity modeling the assets may be classified as: non-confidential data; confidential consumer data; confidential corporate data; system capabilities (such as electronic funds transfers); non-digital assets (such as infrastructure or IT-system controlled equipment, and so forth. In a residential security model the classification could be: lightweight valuables; heavy weight valuables; personal information; life and health of occupants; residential structure/utilities; home security (entry keys/codes), etc.

Step (5) has already been described to some extent previously with respect to FIG. 17 with the creation of an attack tree (a partial attack tree is shown in FIG. 17—a full attack tree would include all assets as topmost nodes, and in implementations the attack tree may be similar or identical to the “threat tree” shown in FIG. 14 of the ‘021 Application). Referring to FIG. 17, the attack vector connects the asset with the component from which the threat emanates, and this association was determined by the previous classification step. In implementations the attack/threat tree may be simplified by only showing one topmost node (one asset) and only its associated components and threats. In implementations the attack tree only shows three levels—the topmost level being the assets, the middle level being components, and the bottom level being threats. This model can be simplified further to have only two levels—the topmost level includes asset/component pairings, and the bottom level shows the threats. In this simplified model the attack tree may show the asset identified as being nested within a component or otherwise attached or coupled with it (in implementations the topmost level would then show, for example, component 1 including asset 1 in level 1, component 2 including asset 1 in level 1, component 3 including asset 1 in level 1, etc., with level two (the bottom level) showing the threats coupled with the components). As a representative example, the asset could be jewelry, the components could be a front door, a back door, and a window, and a threat to the front door could be a bump-key, and since an attacker may use a bump-key to enter the front door, the bump-key is a threat to the asset.

Step (5) includes a second part, which is analyzing the identified threats. In implementations this involves various stages. One stage includes creating a profile for each threat. This may include utilizing real-world threat intelligence to provide attributes to said threats, and the attributes may include defined technical means to exploit the threat. A numerical value may be assigned to this means attribute (a high value indicating a higher level of technical means required to exploit the threat. This means attribute may also include the requisite skill, toolset, social network, financial backing, etc. the attacking entity would need to utilize the threat and traverse the corresponding attack vector. As examples, the technical means for the aforementioned bump-key against the front door is relatively low, whereas the technical means to alternatively burrow through the home’s foundation to bypass other security controls is relatively high.

Another threat attribute is the access required to carry out the attack. A number is assigned to the access attribute, with a high number indicating a higher required level of access. A low access attribute may for example be attached to a

29

lower level employee's alphanumeric password (which may be known by an immediate supervisor), while a high access attribute may be attached to biometric signatures for key personnel (which may require high access levels such as a privileged system administrator).

Another threat attribute is a defensive distance, or number of security controls and their effectiveness, between a threat and the asset. A high numerical value indicates a greater number of effective security measures in between. Another threat attribute is the potential impact if the threat is realized—a high value indicates a high technological or business impact in the event the threat is realized.

The second stage involves analyzing the attacker population. Two methods are disclosed herein. In one method (the preferred method), potential attackers are analyzed based on the identified threats and identified associated assets. The other method allows users to adopt existing threat intelligence about known attackers who have utilized existing threats (for example data of a 30% chance of a home invasion in any given year compared with a 1 in 20 million chance of a terrorist attack at the home). In the preferred method objective attributes and characteristics of attackers are provided even where such intelligence is not known.

The third stage involves determining a likelihood of a threat being realized. When the above preferred method is utilized this includes three sub steps of determining a subset of the attacker population with the sufficient means, sufficient motivation, and sufficient opportunity to exploit the threat, to determine the probability of an attacker subset coming from the attacker population (this discards attackers who would be unable to exploit the threat for whatever reason) (if the other method is used the statistically-provided probability is used). A determination is then made of whether the asset is attractive to the attacker, if not the likelihood of threat exploit is set to a minimum value, but if it is attractive the likelihood of threat exploit is the product of an attacker coming from the attacker population and the level of attraction. The next step calculates the product of the impact should the identified threats be realized and the likelihood of the threats being realized for each threat. Each threat is then prioritized according to some scheme (e.g., critical, high, moderate, etc.).

Step (6) involves analyzing attackers using means, motive, and opportunity (as identified above, in the preferred embodiment this occurs partway through step (5)—but if real world intelligence is used this step is obviated and skipped). If the attacker lacks either sufficient means, or sufficient motivation, or sufficient opportunity to conduct an attack, the attacker is disregarded, whereas if the attacker has all three the attacker is included in the organizational attacker population. These determinations could be done, for example, by any known or discovered method to assign a meaningful and relevant value to each category, and if the value of the attacker's technical means meets or exceeds the above set technical means value (and similarly for a motivation value compared with the above identified access level value and an opportunity value compared with the above identified defensive distance value) then the attacker is included in the organizational attacker population, otherwise not.

Step (7) has previously been described with respect to FIG. 5.

Step (8) includes reviewing the threat report after all of the above steps have been done. Step (9) involves comparing the outputs obtained from the threat model and process relative to the desired outputs and desired objectives indicated in step (1). In implementations steps (1) and (9) may

30

be done through system interfaces. In other implementations they may simply be done manually. In still other implementations they may both be excluded.

Referring now to FIG. 21, a representative example of attack simulation methods is shown. FIG. 21 shows the methods illustrated using the diagram interface 1000, though in other implementations the methods may be implemented using a separate interface that takes the components diagrammed in any given threat model and displays them in a separate interface just for the purpose of illustrating attack vectors. In FIG. 21 it is to be understood that the user has previously identified, through the database, the ONE DRIVE element as an asset. In other implementations the asset could be linked to the ONE DRIVE component but could be displayed as its own separate component—for example a “client info database file” shown next to or within the ONE DRIVE component (though in such a case the “asset” itself would be a “component” as that term is used herein). In the shown implementation the ONE DRIVE component itself is identified as the asset.

A number of dashed protocol lines are shown in FIG. 21. These are shown in this fashion after the user has selected the ONE DRIVE component and made a selection to show all threats that could compromise the component. For example, this could be done in implementations by right-clicking on the ONE DRIVE component and selecting a selector such as “Show Attack Paths,” though the attack paths may be shown using other methods that the practitioner of ordinary skill in the art may determine. Once the selection is made, all threats reaching the ONE DRIVE component are highlighted in some manner. In the representative example they are highlighted by showing dashed and “moving” arrows, meaning an animation which shows the dashed portions of the lines as moving in one direction or another. In some implementations this animation could show the dashes moving in the direction towards the asset, which may assist the user to trace attack paths back to the outermost threats. In other implementations the arrows could be reoriented when the attack paths are shown so that the arrow tips all point downstream towards the asset, which may also assist the user to trace attack paths back to outermost threats. In implementations the attack paths are shown in a different color and/or using a greater line thickness than other protocol lines. In implementations the attack paths are shown in bright red and using thick lines while non-compromised paths are shown in thin black lines.

It can be seen that there are threats associated with the WIFI ACCESS POINT and with the SHARE POINT components, and that these threats reach the ONE DRIVE component through OFFICE 2010 and IE11 components. There are a number of potential compensating controls in the diagrammed system which may be utilized to mitigate these threats. For example, the laptop component group is shown including the MCAFEE HIPS, MCAFEE AV VSE, BIT LOCKER, and MCAFEE DLP compensating controls, and there is also an SSO compensating control diagrammed between the IE11 and ONE DRIVE components. All compensating controls in this diagram are shown as having an icon of a padlock on them so that the user can easily view the compensating controls to modify the threat model and the attack simulation (the compensating controls in implementations do not introduce any new threats but only mitigate threats).

Referring now to FIG. 22, in this screen the user has toggled the MCAFEE AV VSE and SSO elements to an ON configuration. In implementations this is done by right-clicking each component and selecting “BLOCK” from a

31

menu, which then changes all attack paths that are mitigatable by the selected compensating control to a mitigated display state (in some implementations the user may be able to select which attack paths are mitigatable from this interface). In FIG. 22 these compensating controls are highlighted by outlining them in a thicker line to show that they are turned on (though this display may be excluded in other implementations). The attack paths that are mitigated by the compensating controls that have been toggled to an ON state (or to a blocking state) are now shown as solid lines again, while those that are not mitigated are still shown in dashed format. In implementations the mitigated attack paths are shown in a thick green line. The user may reselect the compensating controls and select unblock to revert the threat model to the previous state. The user may also delete any given compensating control from the diagram altogether, and this will have the effect of changing any attack paths that were mitigated by that compensating control to an unmitigated display state.

While the mitigated display state for attack paths are described herein as green and bold and the unmitigated display state for attack paths are described as dashed, red, bold and animated, the practitioner of ordinary skill in the art could select other visual techniques to differentiate attack paths from non-attack paths and mitigated attack paths from non-mitigated attack paths. When a user toggles a compensating control to an ON state this may have the effect of dynamically changing the threat report, the compensating control report, the mitigations interface, and so forth by toggling relevant threats to mitigated (and toggling the same compensating control to the OFF state or removing it altogether from the diagram may have the opposite effect). In implementations one or more of the user interfaces (such as those of FIGS. 21-22) may allow user selection of certain threat sources and deselection of others, so as to only show attack paths of threats associated with a selected subset of the diagrammed components. The threats are shown in FIGS. 21-22 as emanating from outermost components (and the system/method may default to checking for threats at all outermost nodes), but in implementations they could emanate from internal sources (such as a USB port in an internal system administrator computer being a threat source for a virus). An attacker could be visually displayed at each threatened component, for example an "attacker" component could be placed on the diagram coupled with the WIFI ACCESS POINT and SHAREPOINT components to indicate the presence of an attacker. In other implementations the attacker need not be displayed in this way, but the user may understand that an attacker is reaching for the asset through the end points from where the threats are shown emanating on the diagram. From these end point components, or from an actual "attacker" component, the user may edit an attack profile (such as to indicate along a spectrum from a low funded individual hacker all the way to a state sponsored team of skilled and well-funded hackers) and this may modify the attack paths. For example, if the latter profile is identified then certain compensating controls may not mitigate the threats (for example the SSO element may not be effective, and another compensating control would be needed instead), and accordingly when the profile is changed one or more of the attack paths may change to a mitigated or unmitigated display state in response. The identification of hacker type may be done, for example, through dropdown lists, checkmark items, and so forth that had been previously set through the database by the user, and the ability of a compensating control to mitigate threats by lower-end attackers but not higher-end attackers may also be

32

defined/set through the database in the input interfaces previously described (though using selectors not shown—but which the practitioner of ordinary skill in the art will know how to implement).

In implementations another interface may allow the user to get a list of all threats which may compromise the asset and all attack paths associated with those threats. The asset has been described above as being anything of value to the user, it could for example be a database, a financial asset, a trade secret (or files related to trade secrets), and so forth.

One more comment is made here with respect to threat model chaining. The user may select a sub-component or component group (for example the user could right click the WINDOWS 7 component in FIG. 22) and select an "open in new window" selector to open a new diagram interface showing the diagram just for that component group. The user may then view just the threats associated with the component group, using previously identified methods, and/or may identify assets within the component group (for example OUTLOOK 2010) to show attack paths related just to this component. In implementations a user may right click or otherwise select a component group on an existing threat model by selecting a box outlining the component group and then select a "show threats" or other menu item to treat the entire component group as an asset and to show all attack paths thereto.

In implementations the designation of a component as an asset is done as simply as clicking on any component and selecting a "show threats" or similar menu item from a list. In other implementations a component may be designated as an asset by another mechanism and/or using another interface. The asset may be selected and its features edited, as well, as detailed to some extent above (for example some of the asset classification may be done here, though this may alternatively be done using one or more other interfaces). The attack simulation systems and methods herein allow a user to obtain a thorough understanding of threats and risks to multi-faceted systems/processes, and the impact to relevant assets, should potential threats be realized.

Although the examples of computing system models have been described in detail, system 100 may be utilized to model other threats. For example, referring to a disease epidemic threat model, the threats may be sources or ways the disease may spread, the security requirements may be methods or elements to reduce the effects of those sources, such as sterilizing surgery instruments, washing hands, using air filters, breathing masks, gloves, and the components may be steps or interactions that occur with respect to persons who may be infected, such as admitting them, dealing with infected clothing, performing medical operations on them, etc. In such a scenario the threat model may have no components, threats, or security requirements that relate directly to computing devices or systems, though the threat may be modeled on system 100 which utilizes computing devices. The same may be said for the aforementioned threat model of commuting to work. Numerous other threat models, involving any of myriad types of threats, could be modeled using system 100. Other types of threat models could include: modeling physical and technology threats to airplanes, smart homes, smart TVs, an electric grid, and so forth.

Accordingly, the system and interfaces described herein allow a user to model a variety of systems, applications, and processes to determine threats and to mitigate those threats. The user may, at any time, delete any security requirement component (which may be a compensating control component) from a diagram and then generate the threat report

again to revert to the threat statuses prior to mitigation by any particular component. This is one way in which step 518 of FIG. 5 may be accomplished, since removing the mitigating component will revert the mitigated threats back to the open status (except for threats which are also mitigated by another mitigating component). In implementations there may be other ways to “roll back” the threat report to an earlier status, such as with one or more selectors on the threat report itself to revert to an earlier configuration, or one or more “back” selectors on the diagram interface, and so forth.

Utilizing the methods and systems described herein an organization or individual can quickly and efficiently model deployed or contemplated systems, applications and/or processes and assess relevant threats to determine a mitigation strategy. The user may assess new threats to the entire attack surface of a system, application or process as the new threats arise and accordingly develop modified mitigation strategies with minimal disruption to existing operations. Additionally, the methods and systems described herein allow users to study the ability of a compensating control to block new or existing threats without penetration testing (pen-testing), which can be resource-intensive, slow, and not as thorough as desired. The systems and methods also allow assessment of future compensating controls which have not yet been implemented (and accordingly where penetration testing is not yet an option). Any of the threat modeling described herein may utilize existing threat modeling methodologies including STRIDE, P.A.S.T.A., TRIKE, OCTAVE, and/or VAST methodologies, which may include modeling the system, application or process as a data flow diagram (DFD) or a process flow diagram (PFD). The methods and systems herein integrate well with agile and DevOps development environments, and allow a threat modeling process which is simple enough to allow non-security experts to participate in the threat modeling process and to benefit from the outputs of the system/method.

In implementations a compensating control record in the database may include a definition or details which include: an articulation of the rigor and intent of a prescribed security requirement which the compensating control is meant to satisfy; a defensive level of the compensating control; security and defensive capabilities of the compensating control, and; additional risk imposed by not adhering to the prescribed security requirement.

Systems and methods described herein may have other features that have not been described in detail. Users may be able to access the interfaces of the system through a web page, as described, and may be able to register using user accounts, with login credentials, functionality for retrieving lost passwords etc. Individual users may be able to import custom libraries and other elements into the database including custom components, threats, threat agents, test cases, security requirements, code snippets, code reviews, data elements, roles (developer, customer, user, admin, etc.), widgets, component types, attributes, templates, threat models, and the like. In implementations widgets may be a subset of components and may include reusable modular components that may be included in some components. Some examples of widgets may be, by non-limiting example: forms, cookies, banners, embedded video, RSS feeds, SMS functionality, SILVERLIGHT, secure hardware token, file upload, XML parser, and so forth. Widgets may be associated with threats through the database, using an interface similar to interface 600 (but accessible by selecting “widgets” from interface 700), and may be added to any diagram in the same way that components are added, which may

include defining communication protocols between the widget and other components/widgets. Widgets may in implementations comprise identifying information about a component, user roles which can interact with a component, and the means by which the component receives, transmits, or stores data elements.

In implementations a software installation on one or more computing devices to implement the systems and methods may include storing in the database default libraries, such as over 650 threats, and a number of standard or default definitions for components, default risk levels to choose from, though the user may set up custom threats and associations through the database (which essentially sets up custom risk algorithms), custom risk levels to choose from, and so forth, as described above. The ability of the user to diagram a system, application or process visually allows for non-security experts to analyze threats using a threat model. In implementations, previously generated diagrams may be able to be imported into the system using a PDF or VIZIO drawing.

In implementations the system is configured to interface with third-party software applications and libraries (including software as a service (SAAS)) services, by communicating with the third-party software/library using its application programming interface (API). This may be used, for example, to store standards from standards-issuing bodies in the database and to associate them with security requirements and/or compensating controls through the database. This may assist a user to determine when a component added to a diagrammed threat model would qualify as a compensating control as defined by the standards issuing body. The system may be said to include a “contextual threat engine,” which is the underlying database relations allowing the system to dynamically update the threat model each time a component is mapped onto the canvas, or connected with another component using a communication protocol, or when a protocol is changed or deleted, or a component removed, etc. The systems and methods described herein facilitate comprehensive enterprise-level attack surface analysis. Different users may use different interfaces within an organization in some implementations. Information technology (IT) professionals may define the actors, threats, components, etc. specific to an organization, non-IT professionals may map out systems using the canvas, executives may review the high-level threat details included in the home page or threat model, and so forth (though in other organizations any user may use any interface).

In implementations one or more user inputs may be used using one or more user interfaces to import a threat model from a third-party software or database into the system. When this is done an administrator may map components or elements of the imported threat model to stored components of the system so that the relevant threats and etc. will be included in the imported threat model. Along these lines, it is pointed out here that both end users and administrators are “users” of the system as that term is used herein. Accordingly, when the claims refer to “one or more user inputs” these inputs may be carried out by an end user and/or they may be carried out by an administrator, such as an administrator maintaining a cloud database which several different end users are accessing to implement different deployments of the method/system.

It is further pointed out that systems and methods disclosed herein may use multiple computing devices, interfaces, displays, databases, and so forth. For example, any attack simulation system disclosed herein could include multiple computing devices (each having at least one dis-

35

play) and multiple databases in which data is stored, accessed, and manipulated using multiple interfaces displayed on the displays of the computing devices. The computing devices may be remote from one another and coupled with the database(s) through the Internet and/or locally. Any attack simulation methods and/or threat model chaining methods disclosed herein may include receiving user inputs at multiple interfaces displayed on displays of multiple computing devices coupled with multiple databases to accomplish the method(s) at least in part by storing, accessing and manipulating data in the databases. By non-limiting example, in implementations an administrator could access the one or more databases using an administrator computer to input threat model components, threats, compensating controls, and the like, and associate these elements with one another (and also may create some default threat models which the user may use or manipulate—for example an “ATM” threat model by non-limiting example which an end user could use or include in his/her own diagrammed system, application or process), while an end user could access the one or more databases using a separate remote computer to diagram a system, process or application and thereby generate a threat model to be displayed on the end user computer.

In places where the description above refers to specific embodiments of threat model chaining and attack simulation systems and related methods, one or more or many modifications may be made without departing from the spirit and scope thereof. Details of any specific embodiment/implementation described herein may, wherever possible, be applied to any other specific implementation/embodiment described herein.

What is claimed is:

1. A threat model chaining method, comprising:
providing one or more databases, the one or more databases comprising:
a plurality of threat model components stored therein;
a plurality of threats stored therein, each threat associated with at least one of the threat model components through the one or more databases; and
a plurality of compensating controls stored therein, each compensating control associated with at least one of the threats through the one or more databases;
providing one or more interfaces, including a diagram interface, configured to be displayed on one or more end user computing devices communicatively coupled with the one or more databases;
configuring the diagram interface to display a relational diagram of one of a system, an application, and a process, using visual representations of one or more of the threat model components and visual representations of one or more of the compensating controls, the relational diagram defining a first threat model; and
configuring the diagram interface to, in response to receiving one or more user inputs, add a component group to the first threat model and thereby redefine the first threat model by including in it a second threat model associated with the component group, wherein the component group comprises a predefined interrelated group of two or more of the threat model components.
2. The threat model chaining method of claim 1, further comprising configuring the diagram interface to, in response to receiving a user selection of one of the diagrammed threat model components of the first threat model, visually display

36

attack paths of all threats associated with the diagrammed threat model components which compromise the selected threat model component.

3. The threat model chaining method of claim 1, further comprising configuring the one or more interfaces to, in response to receiving one or more user inputs, store one or more additional threat model components in the one or more databases, store one or more additional threats in the one or more databases, associate each additional threat with at least one of the additional threat model components through the one or more databases, store one or more additional compensating controls in the one or more databases, and associate each additional compensating control with at least one of the additional threats through the one or more databases.

4. An attack simulation method, comprising:

providing one or more databases, the one or more databases comprising:

- a plurality of threat model components stored therein;
- a plurality of threats stored therein, each threat associated with at least one of the threat model components through the one or more databases; and
- a plurality of compensating controls stored therein,

each compensating control associated with at least one of the threats through the one or more databases;

providing one or more interfaces, including a diagram interface, configured to be displayed on one or more end user computing devices communicatively coupled with the one or more databases;

configuring the diagram interface to display a relational diagram of one of a system, an application, and a process, using visual representations of one or more of the threat model components and visual representations of one or more of the compensating controls, the relational diagram defining a first threat model; and

configuring the diagram interface to, in response to receiving a user selection of one of the diagrammed threat model components of the first threat model through the one or more interfaces, visually display attack paths of all threats associated with the diagrammed threat model components which compromise the selected threat model component.

5. The attack simulation method of claim 4, further comprising configuring the diagram interface to, in response to receiving one or more user inputs through the one or more interfaces, visually display one or more attack paths which pass through multiple diagrammed threat model components of the first threat model.

6. The attack simulation method of claim 4, further comprising configuring the diagram interface to, in response to receiving a user toggle of one of the diagrammed compensating controls (hereinafter “first control”) of the first threat model to an on state through the one or more interfaces, toggle all attack paths mitigatable by the first control to a mitigated display state.

7. The attack simulation method of claim 6, further comprising configuring the diagram interface to, in response to receiving a user toggle of the first control to an off state through the one or more interfaces, toggle all attack paths mitigatable by the first control to an unmitigated display state.

8. The attack simulation method of claim 6, further comprising configuring the diagram interface to, in response to receiving one or more user inputs through the one or more interfaces to remove the first control from the first threat model, toggle all attack paths mitigatable by the first control to an unmitigated display state.

37

9. The attack simulation method of claim 4, further comprising configuring the diagram interface to, in response to receiving one or more user inputs through the one or more interfaces, visually display only attack paths of all threats associated with a subset of the diagrammed threat model components.

10. The attack simulation method of claim 4, further comprising configuring the diagram interface to, in response to receiving one or more user inputs through the one or more interfaces to alter a profile of an attacking component, alter the visually displayed attack paths.

11. The attack simulation method of claim 3, further comprising configuring the diagram interface to receive one or more user inputs to add a component group to the first threat model and thereby redefine the first threat model by including in it a second threat model associated with the component group, wherein the component group comprises a predefined interrelated group of two or more of the threat model components.

12. An attack simulation system, comprising:

one or more computing devices communicatively coupled with one or more databases, the one or more computing devices displaying, on one or more displays of the one or more computing devices:

one or more input interfaces configured to, in response to receiving one or more user inputs, store a plurality of user-defined threat model components in the one or more databases, store a plurality of threats in the one or more databases, associate each of the threats with at least one of the threat model components through the one or more databases, store a plurality of compensating controls in the one or more databases, and associate each compensating control with at least one of the threats through the one or more databases, and;

a diagram interface configured to, in response to receiving one or more user inputs, diagram one of a system, an application, and a process, the diagram including one or more of the threat model components and one or more of the compensating controls, to define a first threat model, the first threat model including all threats associated through the one or more databases with the diagrammed threat model components;

wherein the diagram interface is configured to, in response to receiving a selection of one of the diagrammed threat model components of the first threat model, visually display attack paths of all

38

threats associated with the diagrammed threat model components which compromise the selected threat model component.

13. The attack simulation system of claim 12, wherein the diagram interface is configured to, in response to receiving a user input to toggle one of the diagrammed compensating controls (hereinafter "first control") to an on state, toggle all attack paths mitigatable by the first control to a mitigated display state.

14. The attack simulation system of claim 13, wherein the diagram interface is configured to, in response to receiving a user input to toggle the first control to an off state, toggle all attack paths mitigatable by the first control to an unmitigated display state.

15. The attack simulation system of claim 13, wherein the diagram interface is configured to, in response to receiving one or more user inputs removing the first control from the first threat model, toggle all attack paths mitigatable by the first control to an unmitigated display state.

16. The attack simulation system of claim 12, wherein the diagram interface is configured to, in response to receiving one or more user inputs, visually display only attack paths of all threats associated with a subset of the diagrammed threat model components.

17. The attack simulation system of claim 12, wherein the diagram interface is configured to, in response to receiving one or more user inputs, alter a profile of an attacking component, wherein the alteration alters the visually displayed attack paths.

18. The attack simulation system of claim 12, wherein the one of a system, an application, and a process comprises a computing network.

19. The attack simulation system of claim 18, wherein the diagram interface is further configured to, in response to receiving one or more user inputs, diagram a communication protocol between two of the diagrammed threat model components, the diagrammed communication protocol including an alphanumeric indicator.

20. The attack simulation system of claim 12, wherein the diagram interface is configured to, in response to receiving one or more user inputs, add a component group to the first threat model and thereby redefine the first threat model by including in it a second threat model associated with the component group, wherein the component group comprises a predefined interrelated group of two or more of the threat model components.

* * * * *