US012301627B2

(12) **United States Patent**
Crabtree et al.

(10) **Patent No.: US 12,301,627 B2**
(45) **Date of Patent: *May 13, 2025**

(54) **CORRELATING NETWORK EVENT ANOMALIES USING ACTIVE AND PASSIVE EXTERNAL RECONNAISSANCE TO IDENTIFY ATTACK INFORMATION**

(71) Applicant: **QOMPLX LLC**, Reston, VA (US)

(72) Inventors: **Jason Crabtree**, Vienna, VA (US); **Andrew Sellers**, Monument, CO (US); **Richard Kelley**, Woodbridge, VA (US)

(73) Assignee: **QOMPLX LLC**, Reston, VA (US)

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

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **18/892,283**

(22) Filed: **Sep. 20, 2024**

(65) **Prior Publication Data**

US 2025/0030745 A1 Jan. 23, 2025

Related U.S. Application Data

(63) Continuation of application No. 17/237,346, filed on Apr. 22, 2021, now Pat. No. 12,206,708, which is a continuation-in-part of application No. 16/777,270, filed on Jan. 30, 2020, now Pat. No. 11,025,674, which is a continuation-in-part of application No. 16/720,383, filed on Dec. 19, 2019, now Pat. No. 10,944,795, which is a continuation of application

(Continued)

(51) **Int. Cl.**

H04L 9/40	(2022.01)
G06F 16/2458	(2019.01)
G06F 16/951	(2019.01)

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

CPC ***H04L 63/20*** (2013.01); ***G06F 16/2477*** (2019.01); ***G06F 16/951*** (2019.01); ***H04L 63/1425*** (2013.01); ***H04L 63/1433*** (2013.01); ***H04L 63/1441*** (2013.01); *H04L 63/0807* (2013.01); *H04L 63/1466* (2013.01)

(58) **Field of Classification Search**

CPC . H04L 63/20; H04L 63/1425; H04L 63/1441; H04L 63/1433; H04L 63/0807; H04L 63/1466; G06F 16/2477; G06F 16/951

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,477,572 B1 11/2002 Elderton et al.
6,906,709 B1 6/2005 Larkin et al.

(Continued)

OTHER PUBLICATIONS

Eberle, William, Holder, Lawrence; Incremental Anomaly Detection in Graphs, 2013 IEEE 13th International Conference on Data Mining Workshops.

(Continued)

Primary Examiner — Cheng-Feng Huang

(74) *Attorney, Agent, or Firm* — Galvin Patent Law LLC; Brian R. Galvin

(57) **ABSTRACT**

A system and method for correlating network event anomalies to identify attack information, that identifies anomalous events within the network, identifies correlations between anomalies and other network events and resources, generates a behavior graph describing an attack pathway derived from the correlations, and determines an attack point of origin using the behavior graph.

28 Claims, 36 Drawing Sheets

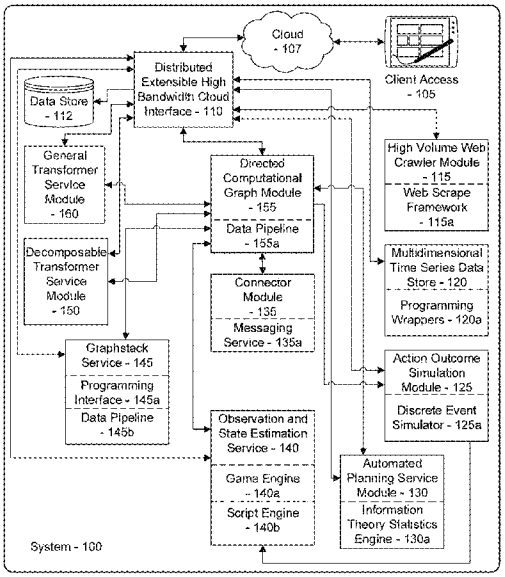

MICROSOFT CORP.
EXHIBIT 1001

Related U.S. Application Data

No. 15/823,363, filed on Nov. 27, 2017, now Pat. No. 10,560,483, which is a continuation-in-part of application No. 15/725,274, filed on Oct. 4, 2017, now Pat. No. 10,609,079, which is a continuation-in-part of application No. 15/655,113, filed on Jul. 20, 2017, now Pat. No. 10,735,456, which is a continuation-in-part of application No. 15/616,427, filed on Jun. 7, 2017, now abandoned, and a continuation-in-part of application No. 15/237,625, filed on Aug. 15, 2016, now Pat. No. 10,248,910, which is a continuation-in-part of application No. 15/206,195, filed on Jul. 8, 2016, now abandoned, which is a continuation-in-part of application No. 15/186,453, filed on Jun. 18, 2016, now abandoned, which is a continuation-in-part of application No. 15/166,158, filed on May 26, 2016, now abandoned, which is a continuation-in-part of application No. 15/141,752, filed on Apr. 28, 2016, now Pat. No. 10,860,962, which is a continuation-in-part of application No. 15/091,563, filed on Apr. 5, 2016, now Pat. No. 10,204,147, and a continuation-in-part of application No. 14/986,536, filed on Dec. 31, 2015, now Pat. No. 10,210,255, and a continuation-in-part of application No. 14/925,974, filed on Oct. 28, 2015, now abandoned, said application No. 15/616,427 is a continuation-in-part of application No. 14/925,974, filed on Oct. 28, 2015, now abandoned.

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,139,747	B1	11/2006	Najork
7,227,948	B2	6/2007	Ohkuma et al.
7,310,632	B2	12/2007	Meek et al.
7,409,721	B2	8/2008	Hernacki et al.
7,530,105	B2	5/2009	Gilbert et al.
7,546,333	B2	6/2009	Alon et al.
7,685,296	B2	3/2010	Brill et al.
7,818,224	B2	10/2010	Boerner
7,818,417	B2	10/2010	Ginis et al.
7,925,561	B2	4/2011	Xu
7,971,252	B2	6/2011	Lippmann et al.
8,069,190	B2	11/2011	McColl et al.
8,281,121	B2	10/2012	Nath et al.
8,346,753	B2	1/2013	Hayes
8,457,996	B2	6/2013	Winkler et al.
8,464,346	B2	6/2013	Barai et al.
8,516,596	B2	8/2013	Sandoval et al.
8,751,867	B2	6/2014	Marvasti et al.
8,752,178	B2	6/2014	Coates et al.
8,782,080	B2	7/2014	Lee et al.
8,881,288	B1	11/2014	Levy et al.
8,914,878	B2	12/2014	Burns et al.
9,100,430	B1	8/2015	Seiver et al.
9,110,706	B2	8/2015	Yu et al.
9,152,727	B1	10/2015	Balducci et al.
9,183,387	B1 *	11/2015	Altman H04L 63/1425
9,231,962	B1	1/2016	Yen et al.
9,235,732	B2	1/2016	Eynon et al.
9,325,732	B1	4/2016	Stickle et al.
9,344,444	B2	5/2016	Lippmann et al.
9,400,962	B2	7/2016	Prasad
9,466,041	B2	10/2016	Simitsis et al.
9,507,944	B2	11/2016	Lotem et al.
9,602,530	B2	3/2017	Ellis et al.
9,628,501	B2	4/2017	Datta Ray et al.
9,652,538	B2	5/2017	Shivaswamy et al.
9,654,495	B2	5/2017	Hubbard et al.
9,679,125	B2	6/2017	Bailor et al.
9,686,308	B1	6/2017	Srivastava
9,699,205	B2	7/2017	Muddu et al.
9,712,553	B2	7/2017	Nguyen et al.
9,736,173	B2	8/2017	Li et al.
9,762,443	B2	9/2017	Dickey
9,774,522	B2	9/2017	Vasseur et al.
9,911,088	B2	3/2018	Nath et al.
9,954,884	B2	4/2018	Hassell et al.
9,967,265	B1	5/2018	Peer et al.
10,021,116	B2	7/2018	Bassett
10,044,726	B2	8/2018	Dulkin et al.
10,083,236	B2	9/2018	Crosby
10,191,768	B2	1/2019	Bishop et al.
10,210,246	B2	2/2019	Stojanovic et al.
10,210,470	B2	2/2019	Datta Ray
10,212,176	B2	2/2019	Wang
10,216,485	B2	2/2019	Misra et al.
10,643,144	B2	5/2020	Bowers et al.
10,846,391	B1	11/2020	Bonney et al.
11,392,875	B2	7/2022	Carstens et al.
11,836,664	B2	12/2023	Ackerman et al.
2005/0165822	A1	7/2005	Yeung et al.
2010/0037293	A1 *	2/2010	StJohns H04L 9/3263 726/2
2010/0299651	A1	11/2010	Fainekos et al.
2011/0307467	A1	12/2011	Severance
2012/0215575	A1	8/2012	Deb et al.
2014/0244612	A1	8/2014	Bhasin et al.
2014/0324521	A1	10/2014	Mun
2015/0020199	A1 *	1/2015	Neil H04L 63/1408 726/23
2015/0347414	A1	12/2015	Xiao et al.
2016/0006629	A1	1/2016	Ianakiev et al.
2016/0012235	A1	1/2016	Lee et al.
2016/0057159	A1	2/2016	Yin et al.
2017/0090893	A1	3/2017	Aditya et al.
2018/0082304	A1	3/2018	Summerlin et al.

OTHER PUBLICATIONS

Huang, Alex, A Comparison of Value at Risk Approaches and a New Method with Extreme Value Theory and Kernel Estimator.

Marozzo, Fabrizio; Talia, Domenico; Trunfio, Paolo; P2P-MapReduce—Parallel Data Processing in Dynamic Cloud Environments, Journal of Computer and System Sciences, 78 (2012) 1382-1402.

Simonian, Joseph, Davis, Josh., Robust Value-At-Risk: An Information Theoretic Approach, Applied Economic Letters, 2010, 17, 1551-1553.

* cited by examiner

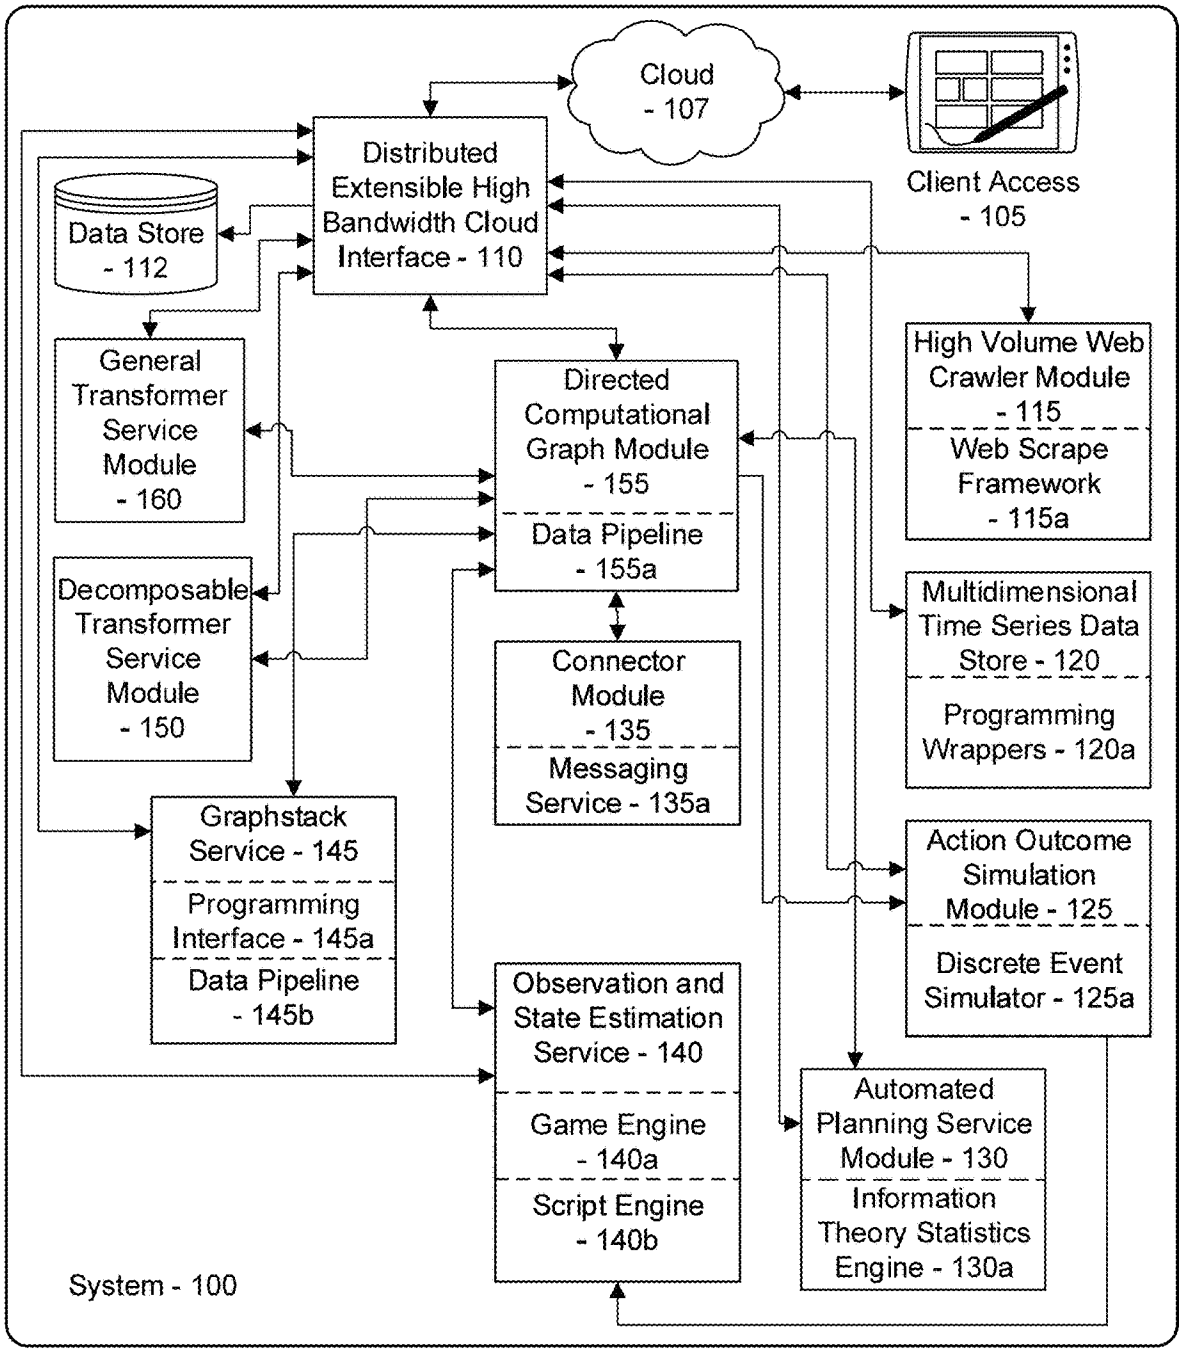
Cloud - 107
Client Access - 105
Distributed Extensible High Bandwidth Cloud Interface - 110
Data Store - 112
General Transformer Service Module - 160
Directed Computational Graph Module - 155
Data Pipeline - 155a
High Volume Web Crawler Module - 115
Web Scrape Framework - 115a
Decomposable Transformer Service Module - 150
Connector Module - 135
Messaging Service - 135a
Multidimensional Time Series Data Store - 120
Programming Wrappers - 120a
Graphstack Service - 145
Programming Interface - 145a
Data Pipeline - 145b
Action Outcome Simulation Module - 125
Discrete Event Simulator - 125a
Observation and State Estimation Service - 140
Game Engine - 140a
Script Engine - 140b
Automated Planning Service Module - 130
Information Theory Statistics Engine - 130a
System - 100

Fig. 1

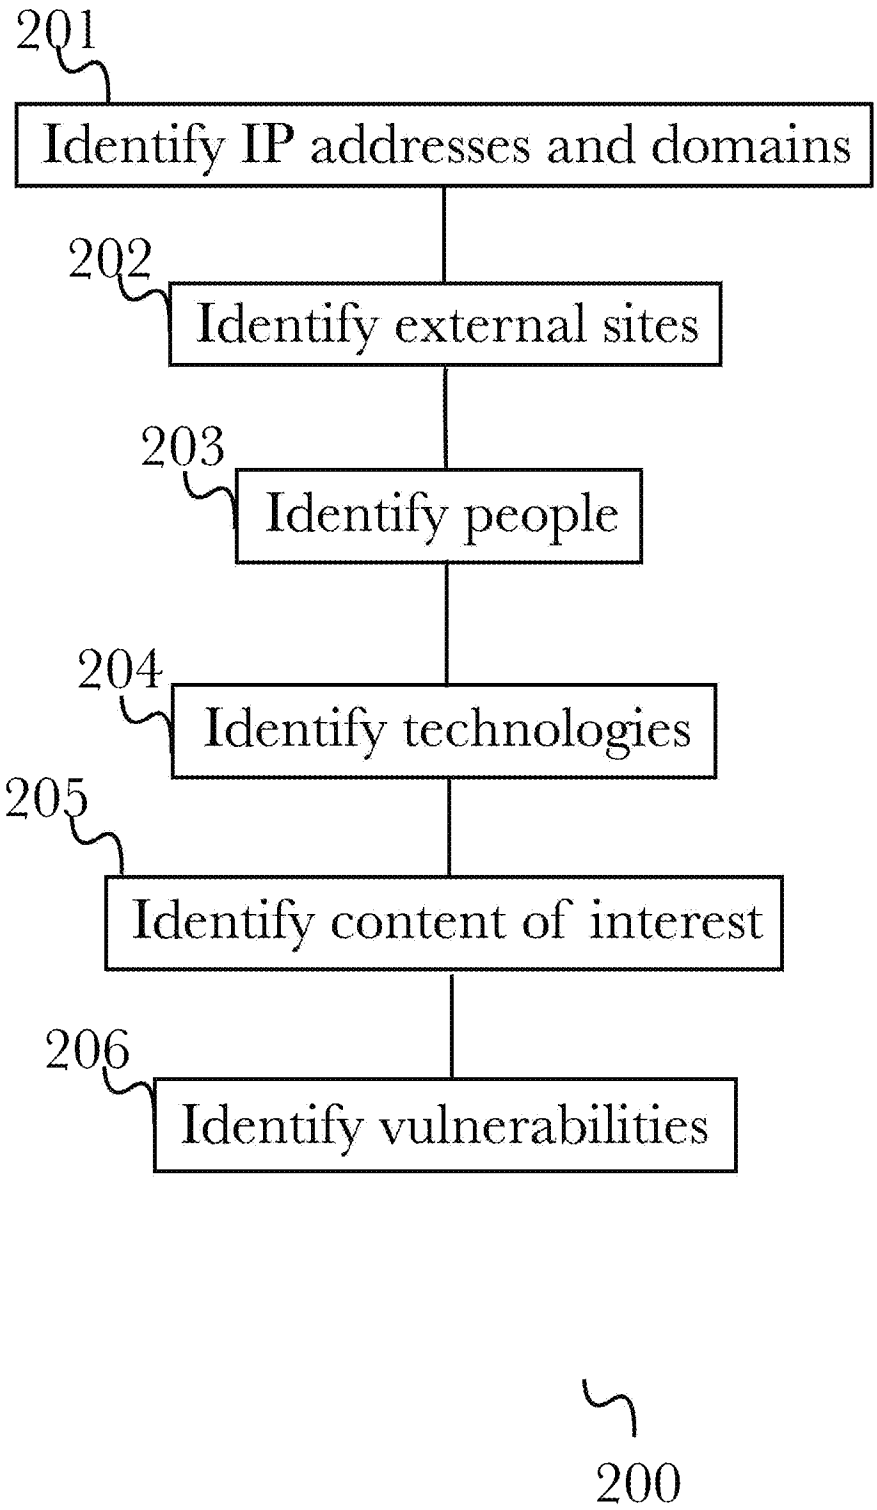
201
Identify IP addresses and domains
202
Identify external sites
203
Identify people
204
Identify technologies
205
Identify content of interest
206
Identify vulnerabilities
200

Fig. 2A

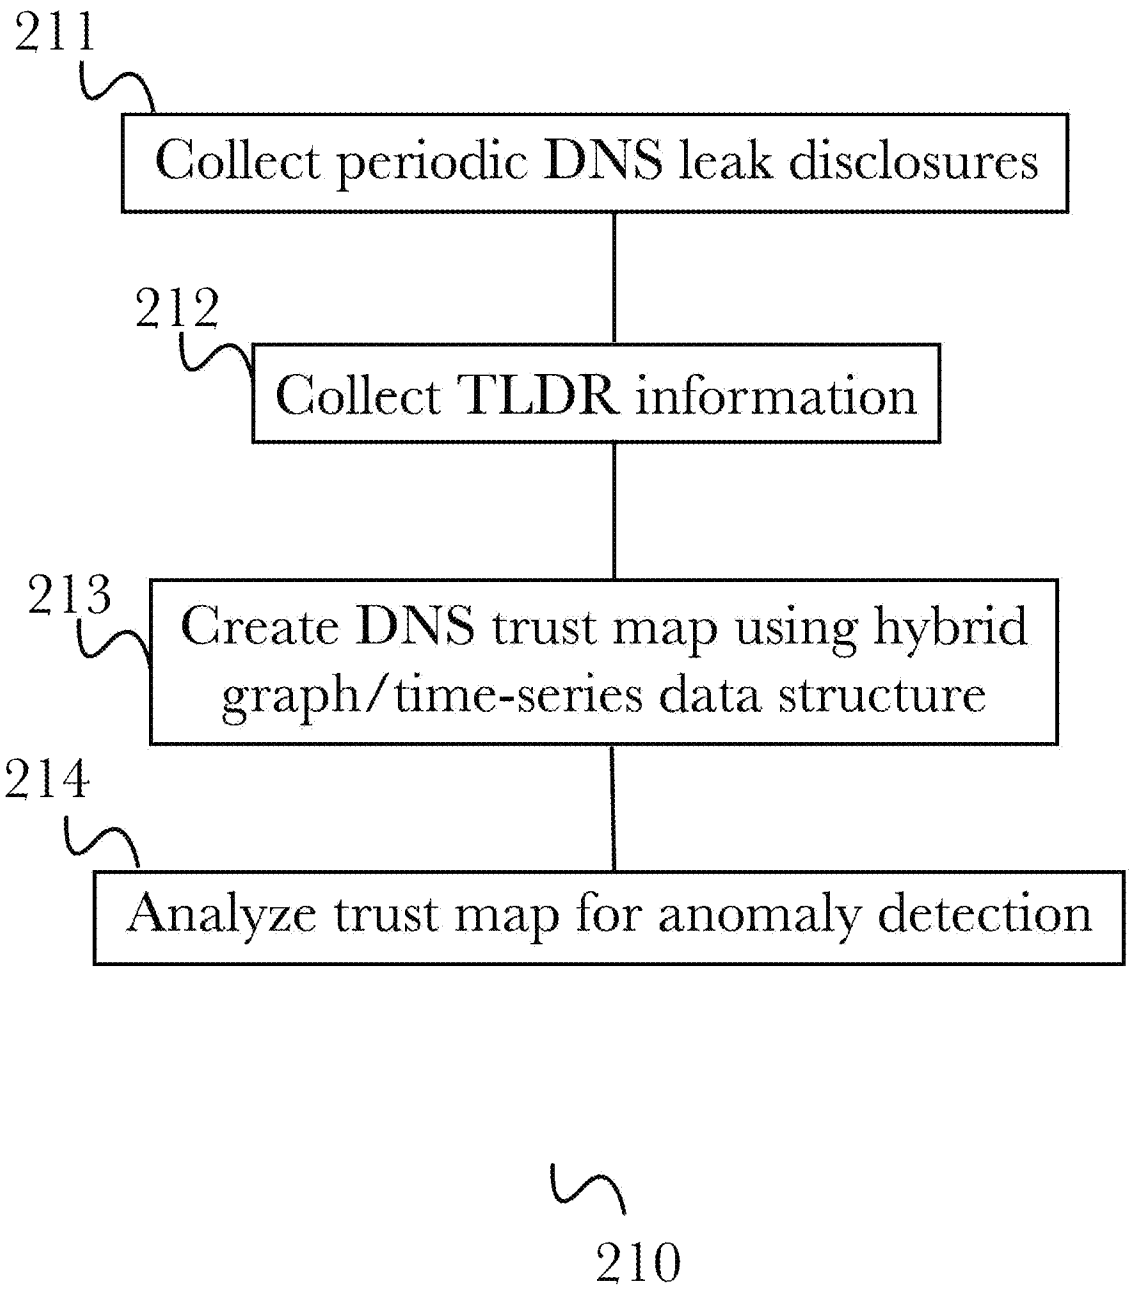
211
Collect periodic DNS leak disclosures
212
Collect TLDR information
213
Create DNS trust map using hybrid graph/time-series data structure
214
Analyze trust map for anomaly detection
210

Fig. 2B

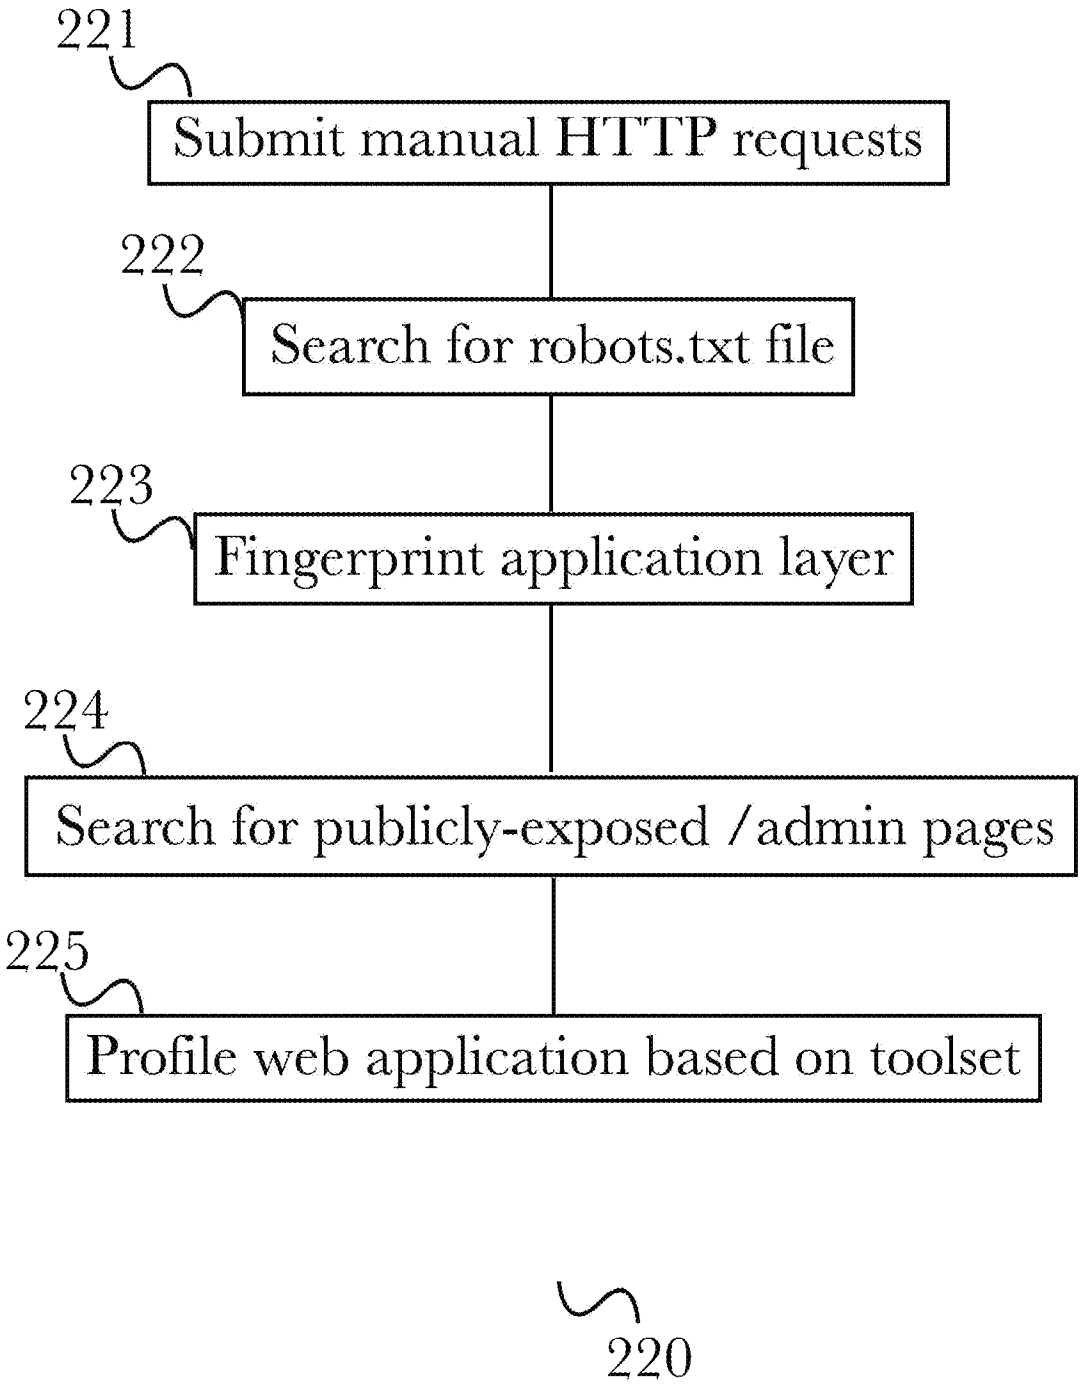
221
Submit manual HTTP requests
222
Search for robots.txt file
223
Fingerprint application layer
224
Search for publicly-exposed /admin pages
225
Profile web application based on toolset
220

Fig. 2C

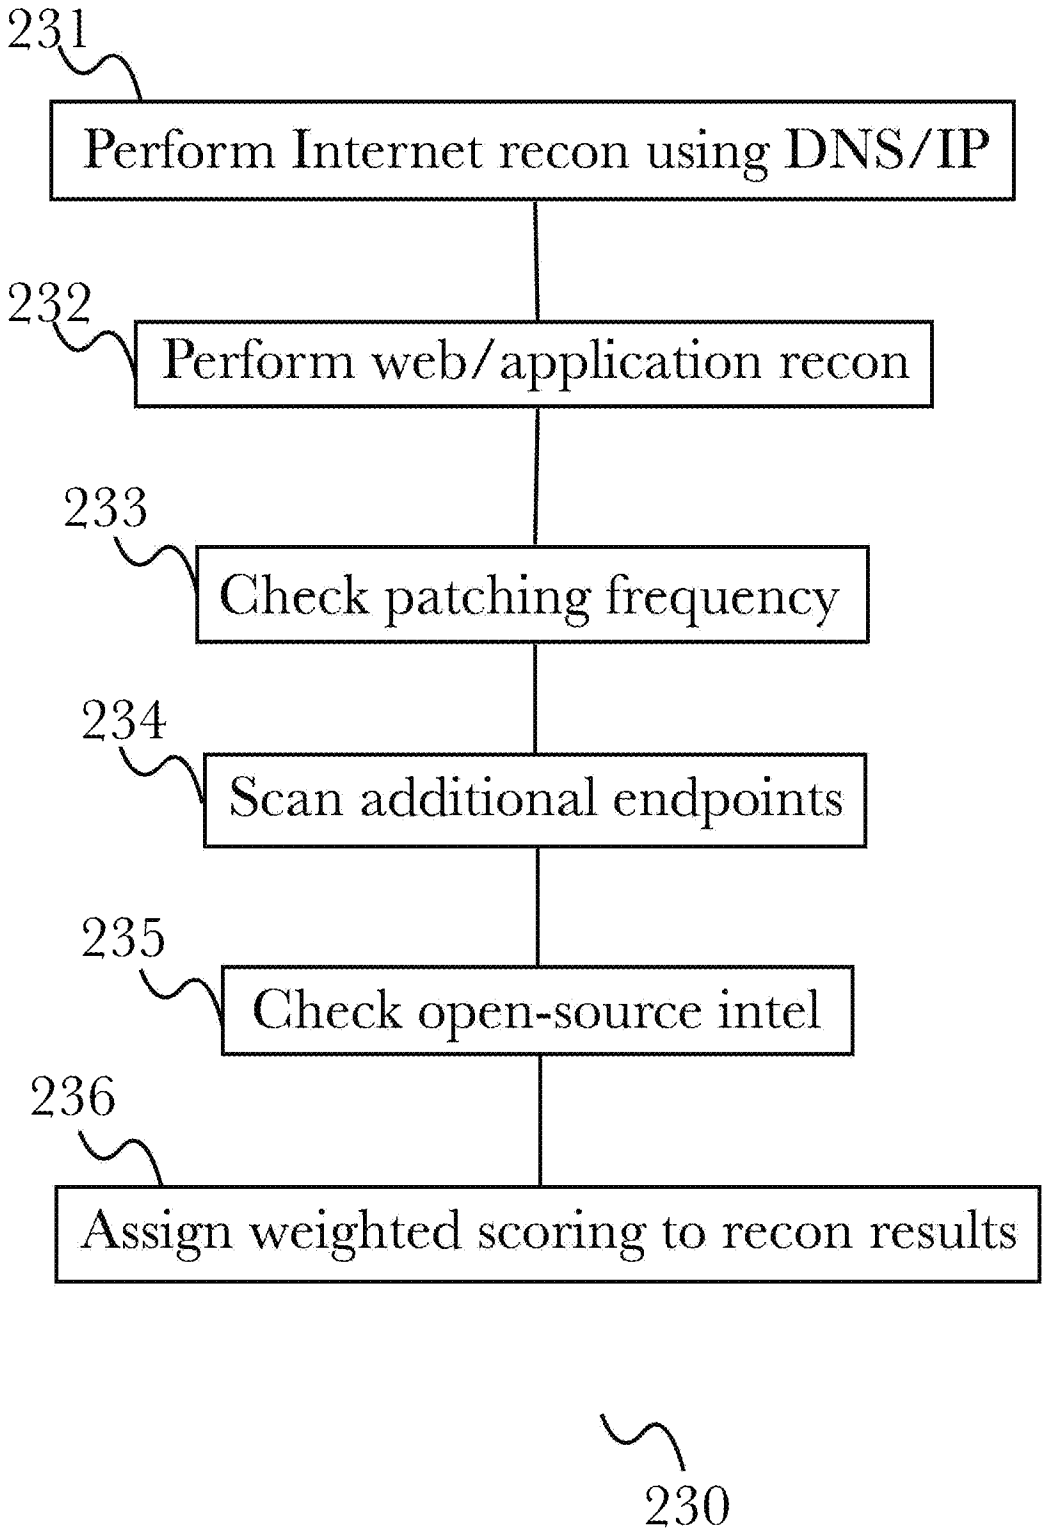
231
Perform Internet recon using DNS/IP
232
Perform web/application recon
233
Check patching frequency
234
Scan additional endpoints
235
Check open-source intel
236
Assign weighted scoring to recon results
230

Fig. 2D

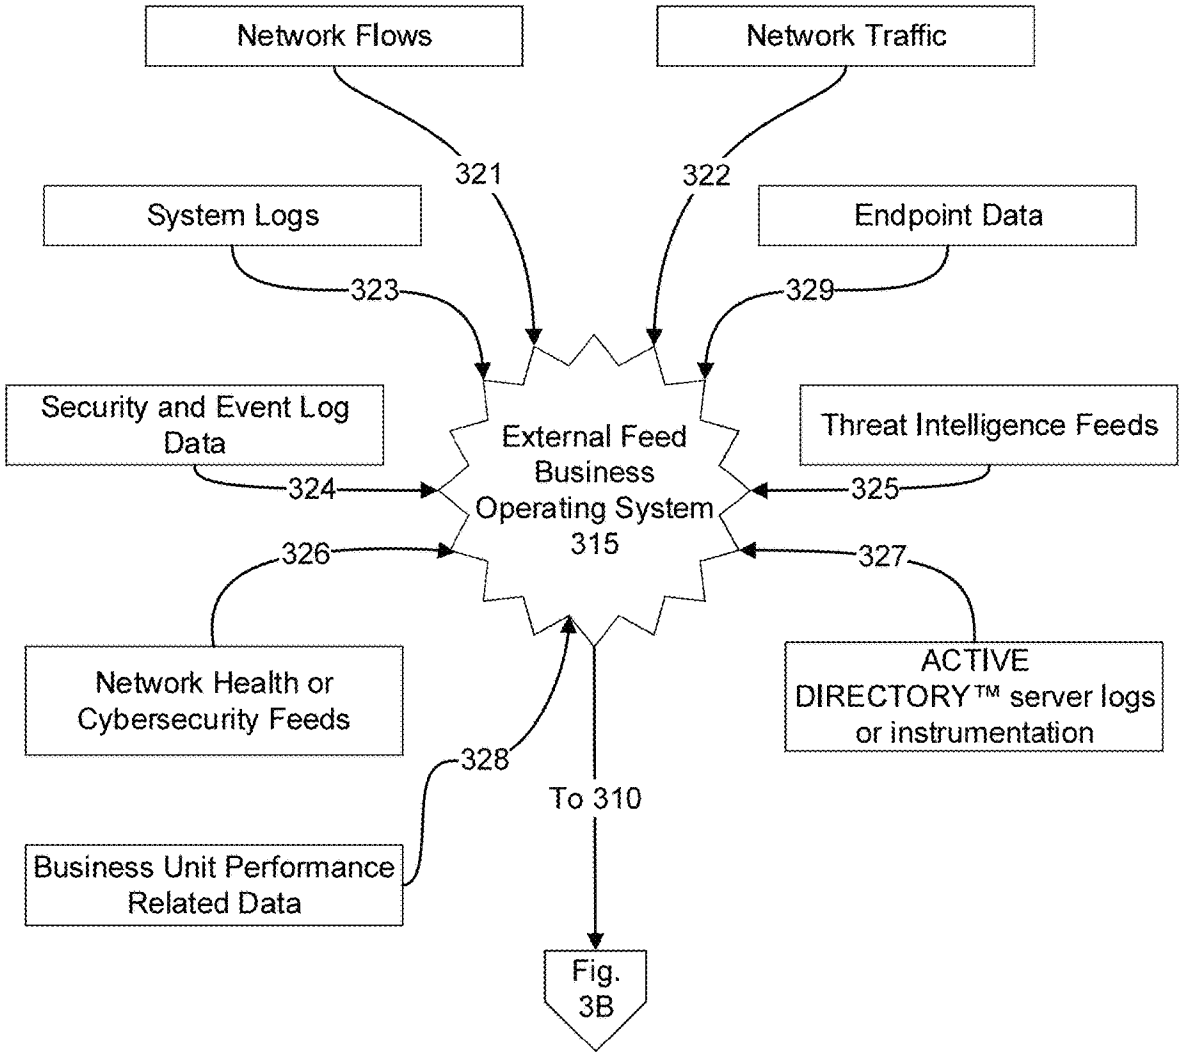
Network Flows
Network Traffic
321
322
System Logs
Endpoint Data
323
329
Security and Event Log Data
External Feed Business Operating System 315
Threat Intelligence Feeds
324
325
326
327
Network Health or Cybersecurity Feeds
ACTIVE DIRECTORY™ server logs or instrumentation
328
To 310
Business Unit Performance Related Data
Fig. 3B

Fig. 3A

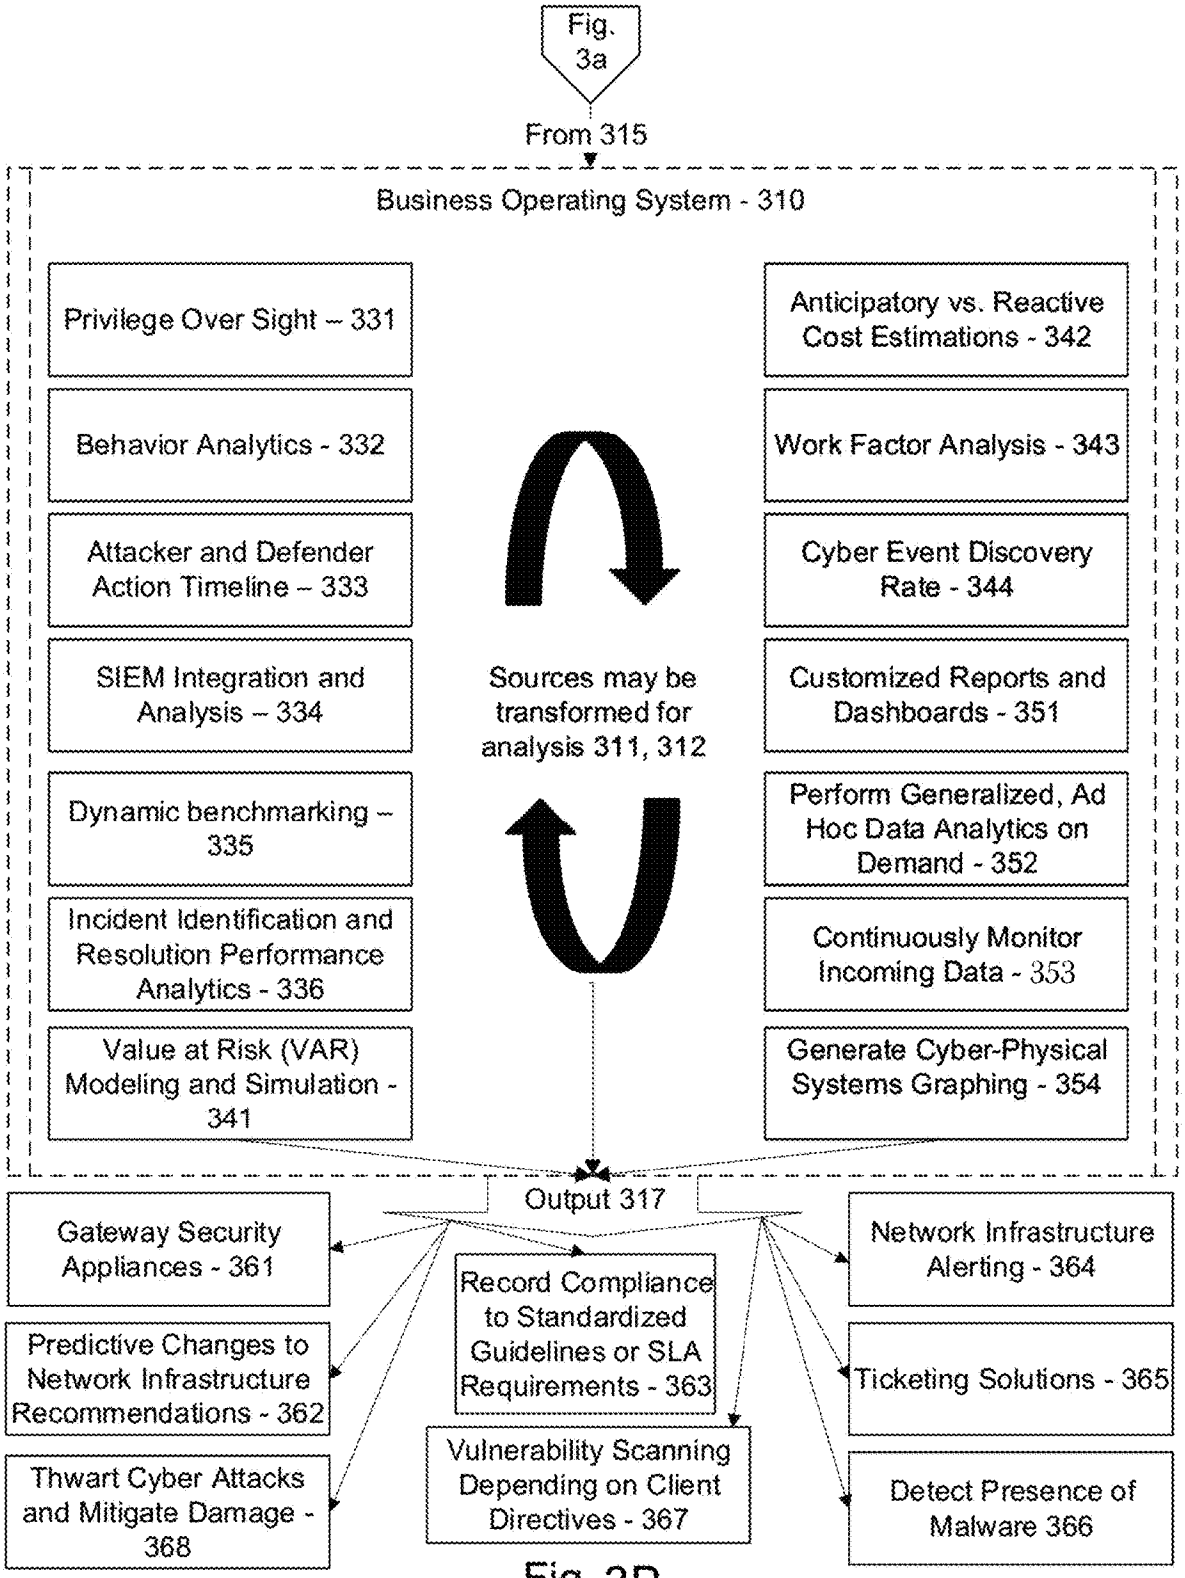
Fig. 3a
From 315
Business Operating System - 310
Privilege Over Sight – 331
Behavior Analytics - 332
Attacker and Defender Action Timeline – 333
SIEM Integration and Analysis – 334
Dynamic benchmarking – 335
Incident Identification and Resolution Performance Analytics - 336
Value at Risk (VAR) Modeling and Simulation - 341
Sources may be transformed for analysis 311, 312
Anticipatory vs. Reactive Cost Estimations - 342
Work Factor Analysis - 343
Cyber Event Discovery Rate - 344
Customized Reports and Dashboards - 351
Perform Generalized, Ad Hoc Data Analytics on Demand - 352
Continuously Monitor Incoming Data - 353
Generate Cyber-Physical Systems Graphing - 354
Output 317
Gateway Security Appliances - 361
Predictive Changes to Network Infrastructure Recommendations - 362
Thwart Cyber Attacks and Mitigate Damage - 368
Record Compliance to Standardized Guidelines or SLA Requirements - 363
Vulnerability Scanning Depending on Client Directives - 367
Network Infrastructure Alerting - 364
Ticketing Solutions - 365
Detect Presence of Malware 366

Fig. 3B

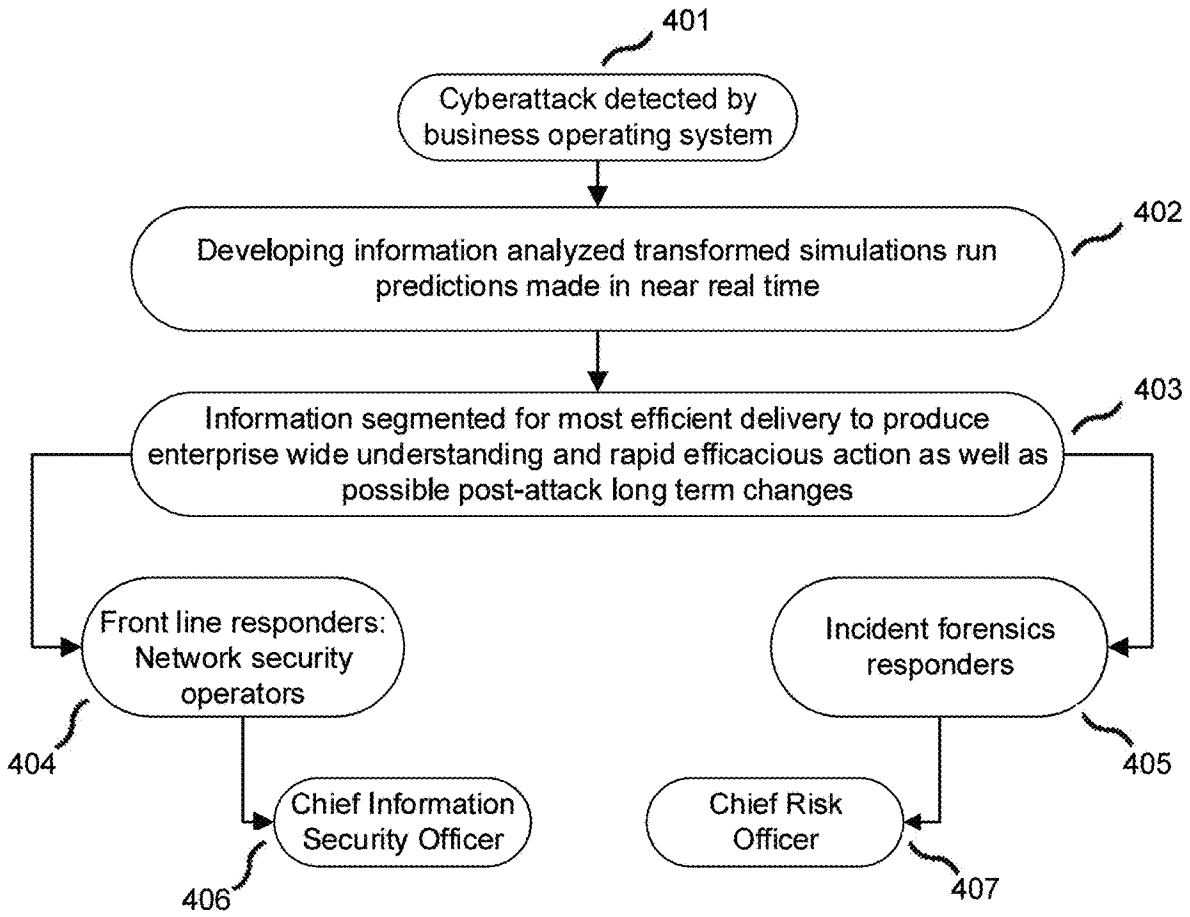
401
Cyberattack detected by business operating system
402
Developing information analyzed transformed simulations run predictions made in near real time
403
Information segmented for most efficient delivery to produce enterprise wide understanding and rapid efficacious action as well as possible post-attack long term changes
Front line responders: Network security operators
404
Incident forensics responders
405
Chief Information Security Officer
406
Chief Risk Officer
407

400

Fig. 4

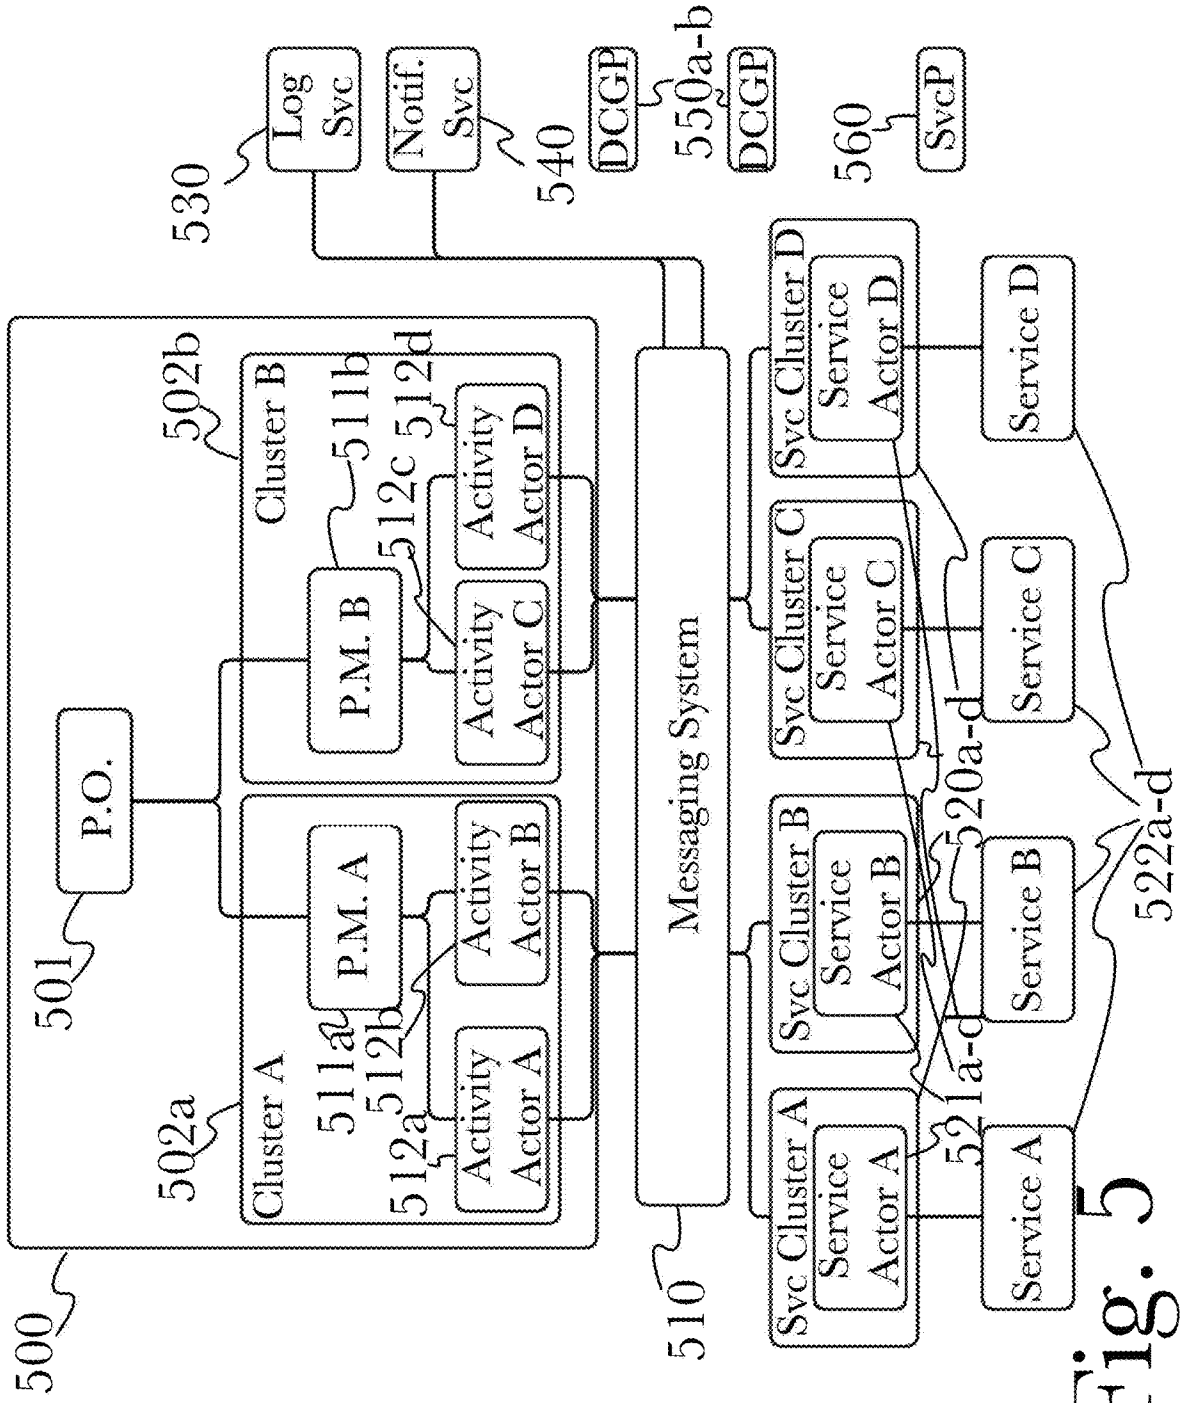
500
501
P.O.
502a
Cluster A
P.M. A
511a
512a
512b
Activity Actor A
Activity Actor B
502b
Cluster B
P.M. B
511b
512c
512d
Activity Actor C
Activity Actor D
530
Log Svc
Notif. Svc
540
DCGP
550a-b
DCGP
560
SvcP
510
Messaging System
Svc Cluster A
Service Actor A
Svc Cluster B
Service Actor B
Svc Cluster C
Service Actor C
Svc Cluster D
Service Actor D
521a-d
520a-d
Service A
Service B
Service C
Service D
522a-d

Fig. 5

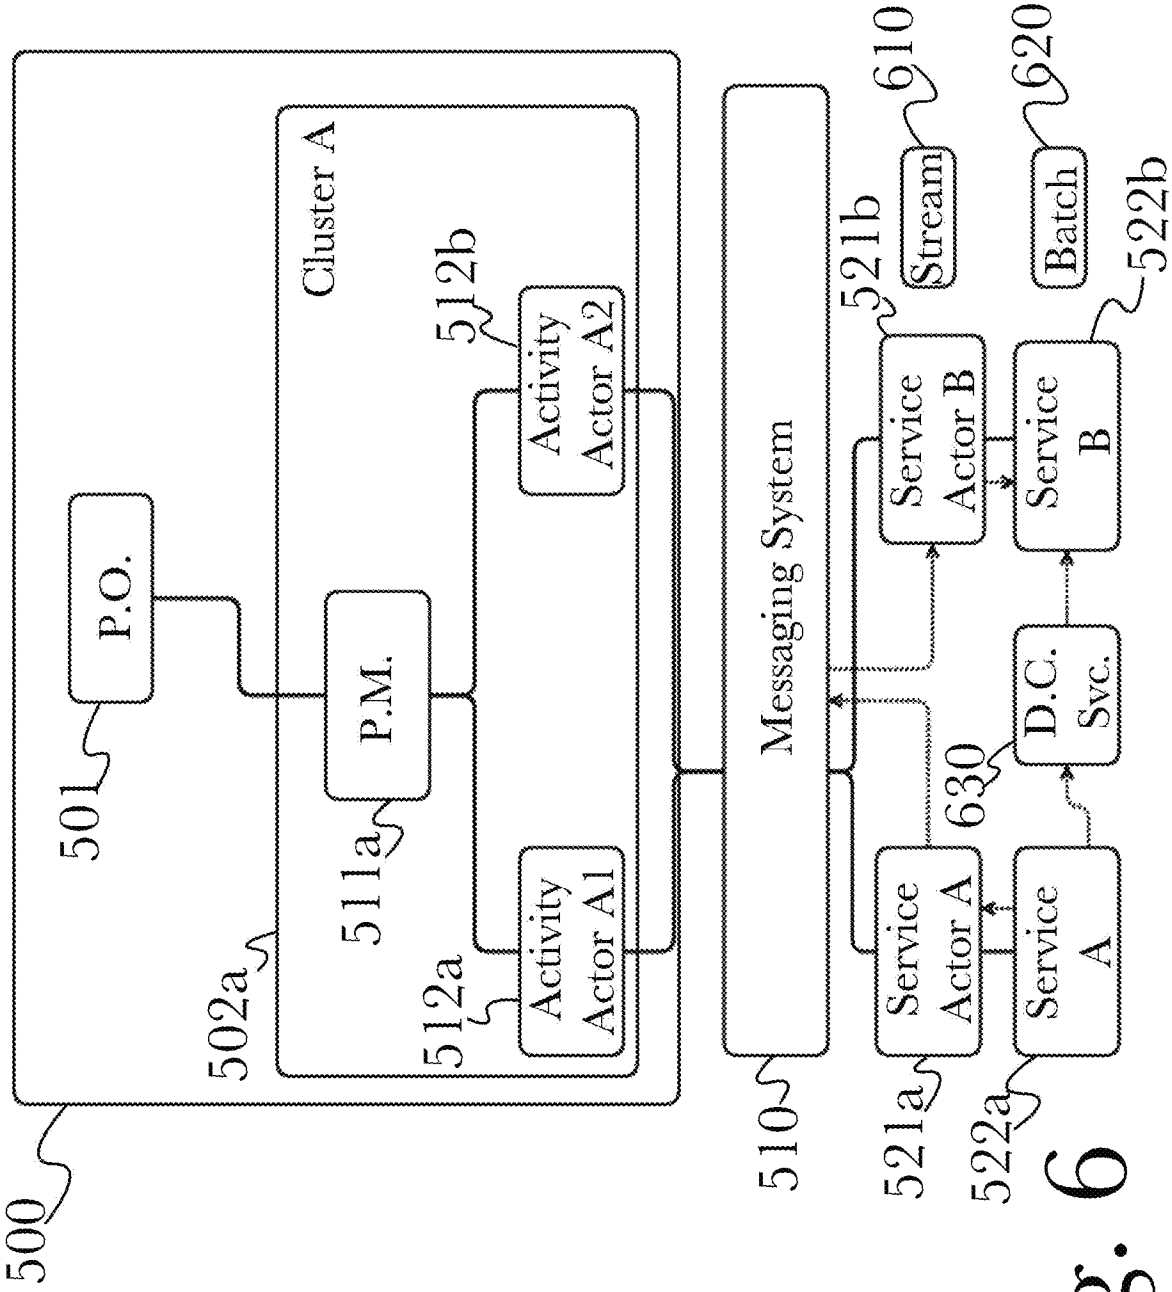
500
501
P.O.
502a
Cluster A
511a
P.M.
512a
Activity Actor A1
512b
Activity Actor A2
510
Messaging System
521a
Service Actor A
521b
Service Actor B
522a
Service A
522b
Service B
630
D.C. Svc.
610
Stream
620
Batch

Fig. 6

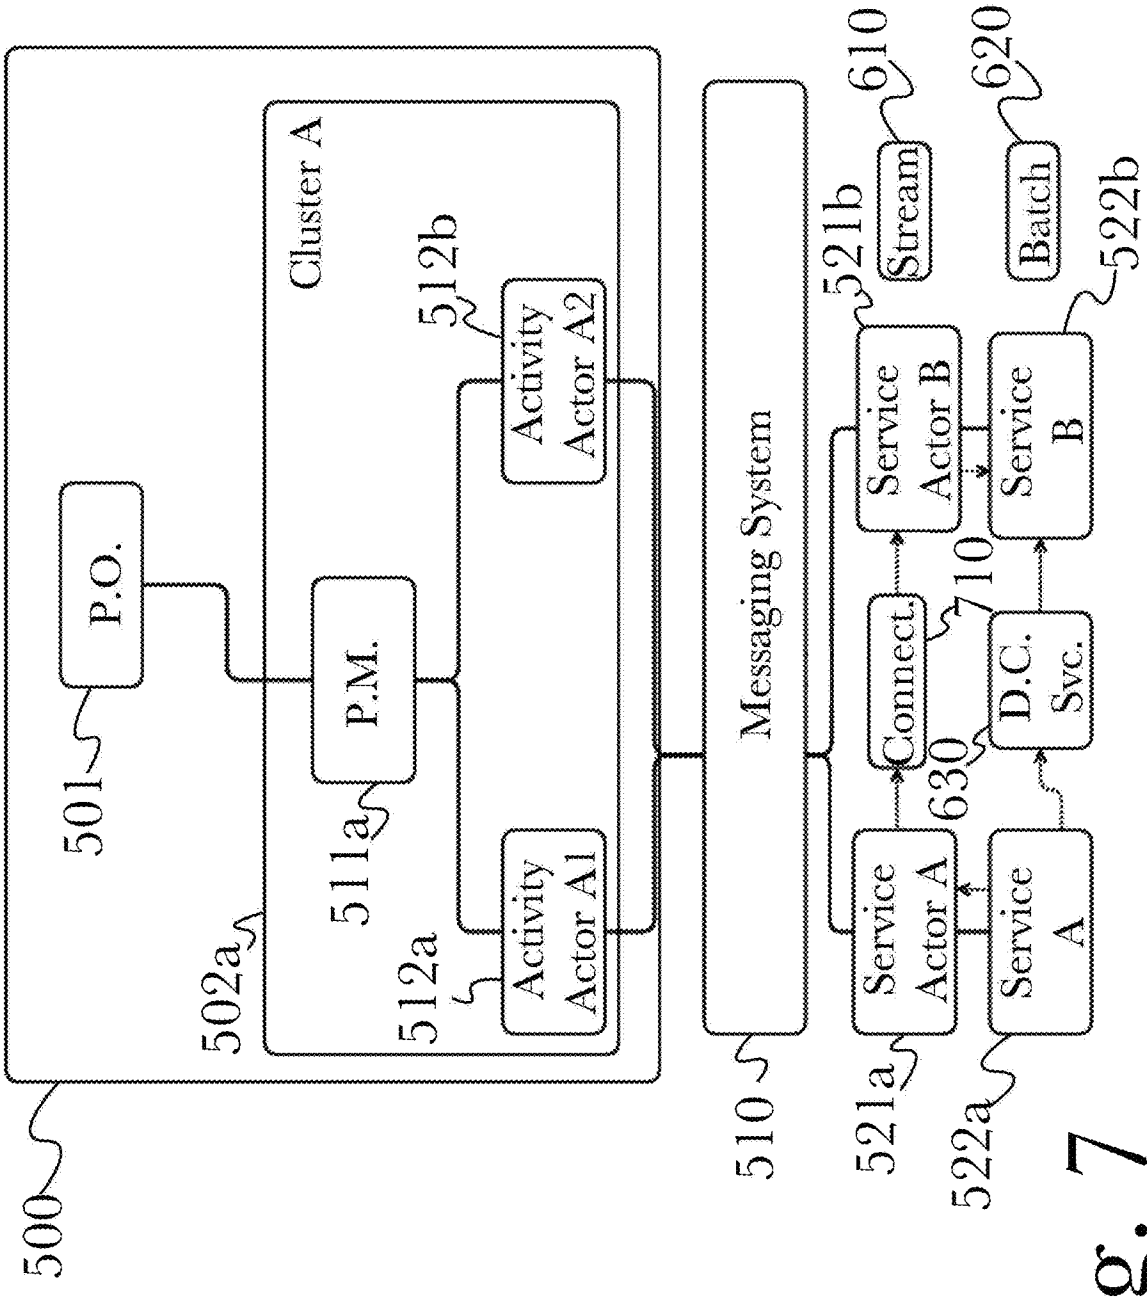
500
501
P.O.
502a
Cluster A
511a
P.M.
512a
Activity Actor A1
512b
Activity Actor A2
510
Messaging System
521a
Service Actor A
521b
Service Actor B
630
Connect.
710
D.C. Svc.
522a
Service A
522b
Service B
610
Stream
620
Batch

Fig. 7

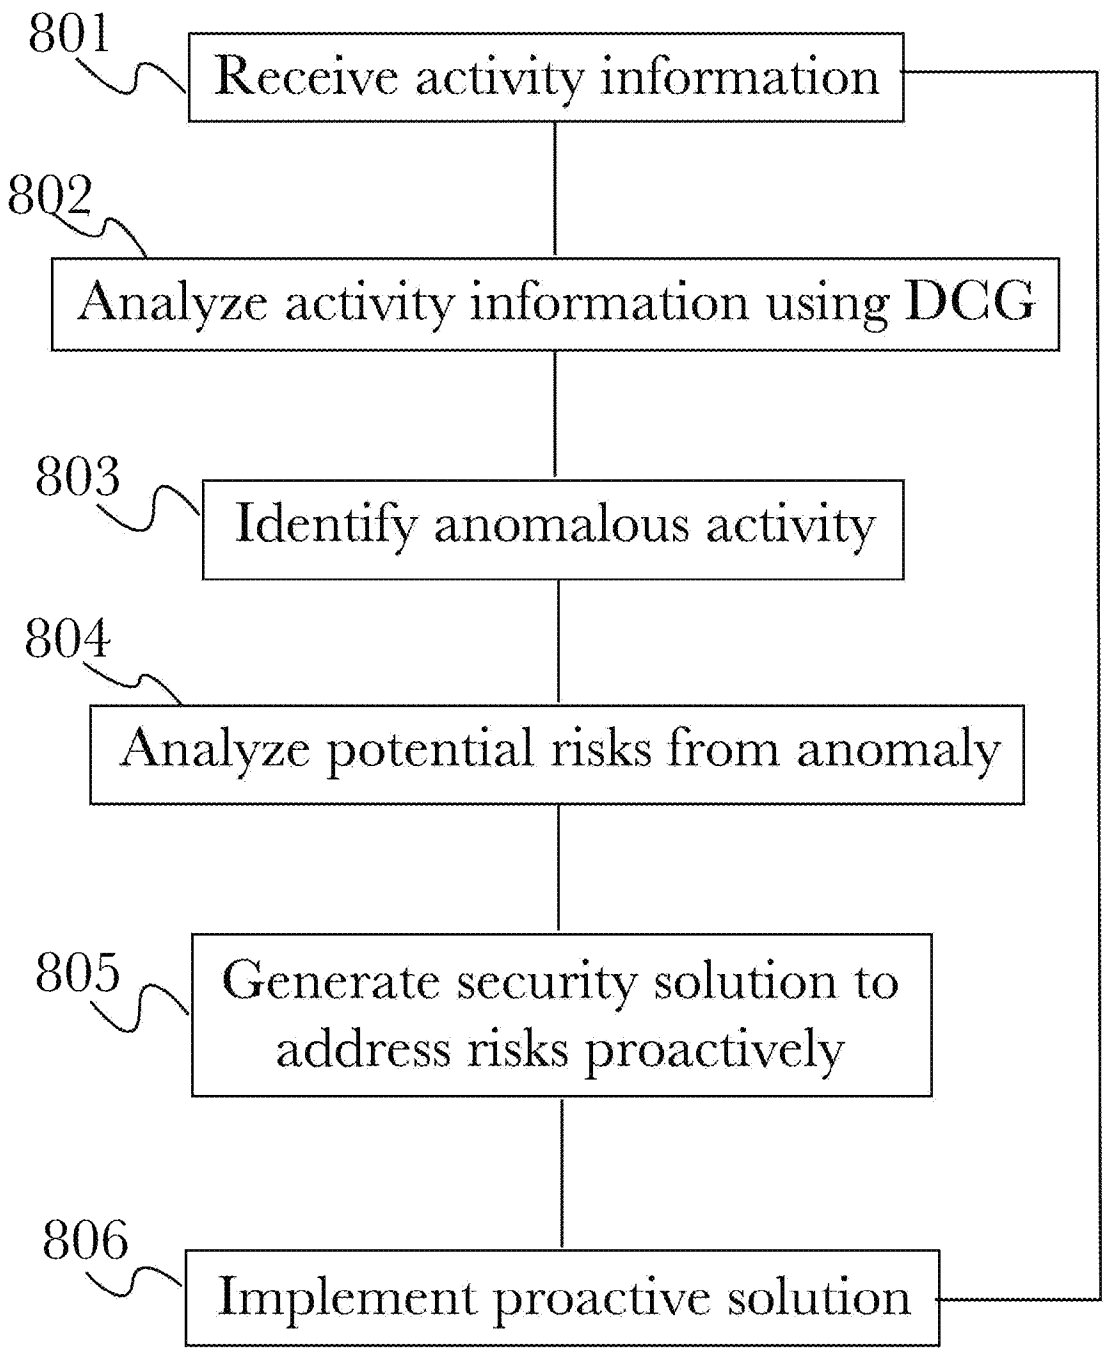
801
Receive activity information
802
Analyze activity information using DCG
803
Identify anomalous activity
804
Analyze potential risks from anomaly
805
Generate security solution to address risks proactively
806
Implement proactive solution

Fig. 8

800

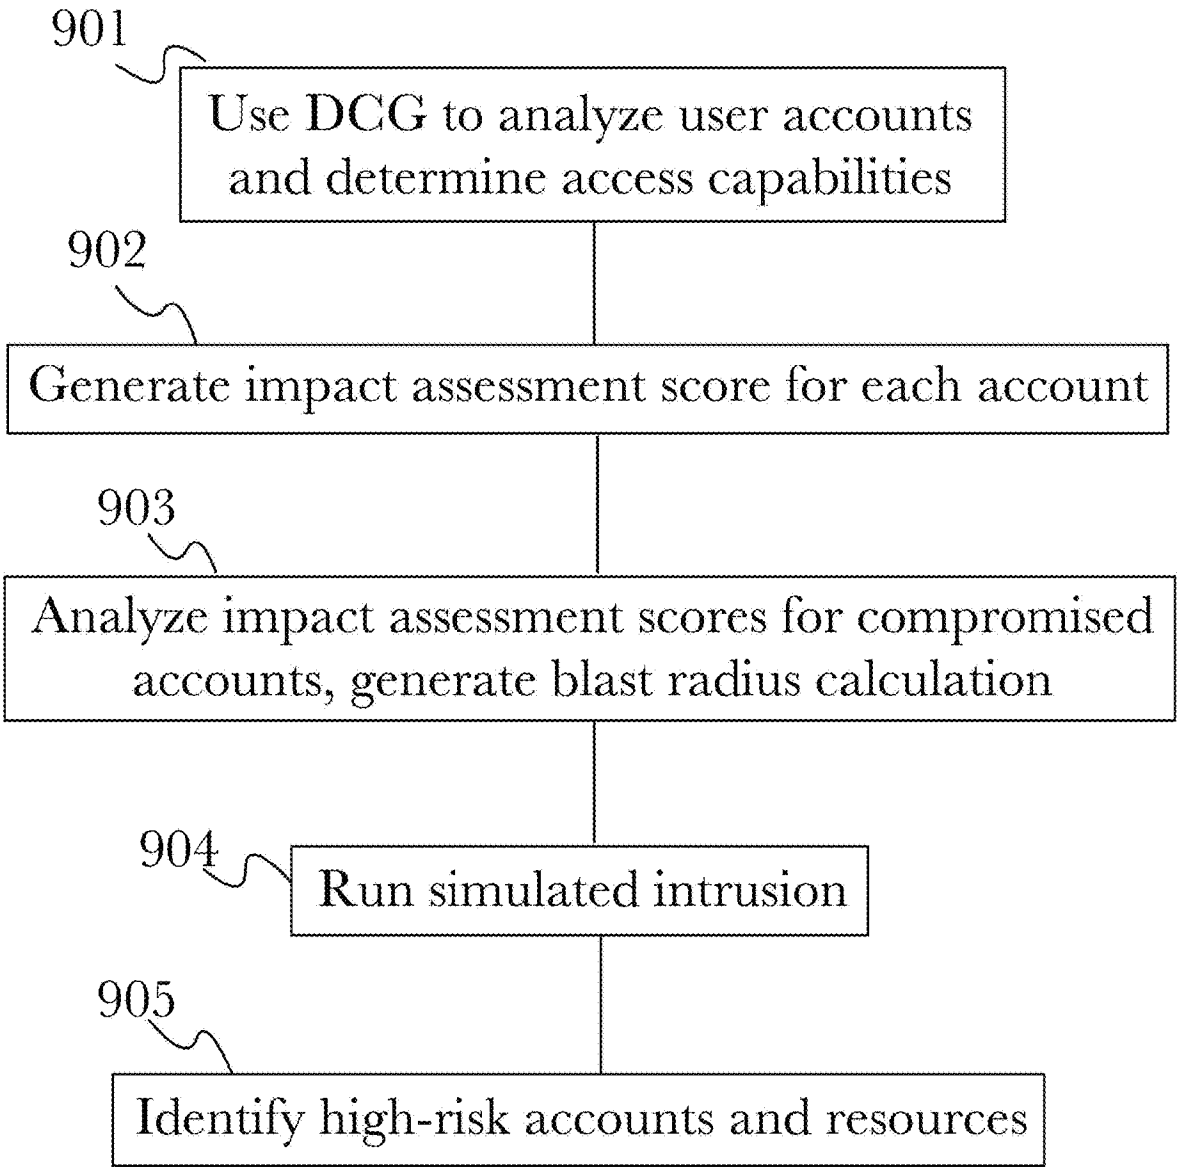
901
Use DCG to analyze user accounts and determine access capabilities
902
Generate impact assessment score for each account
903
Analyze impact assessment scores for compromised accounts, generate blast radius calculation
904
Run simulated intrusion
905
Identify high-risk accounts and resources

Fig. 9

900

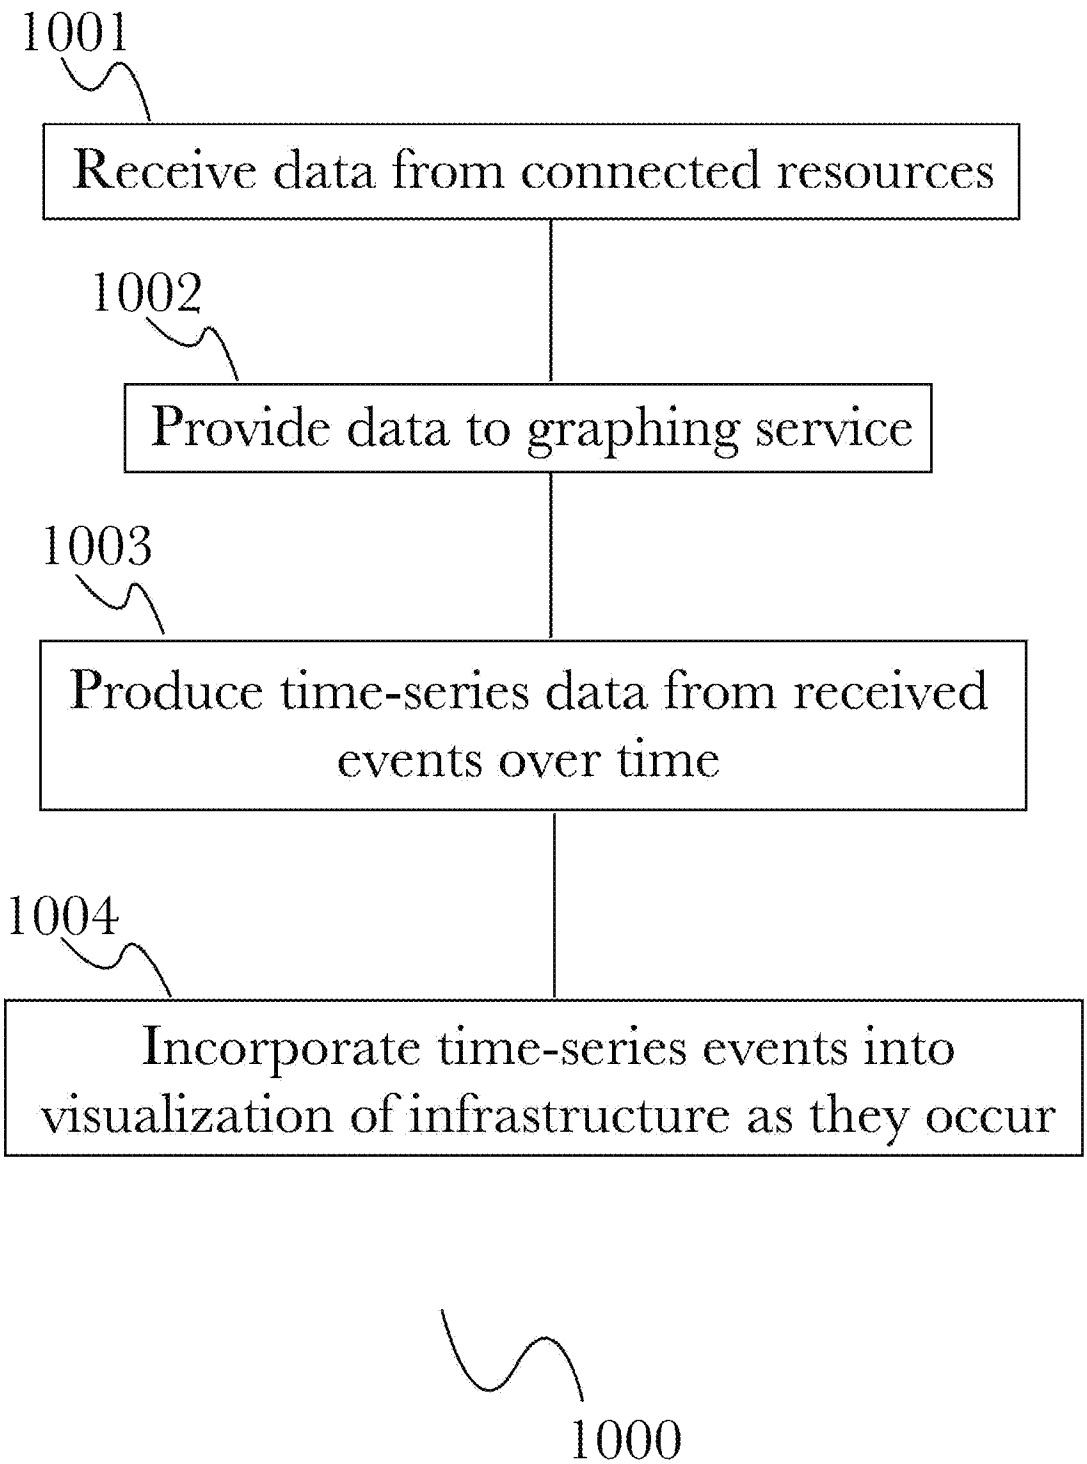
1001
Receive data from connected resources
1002
Provide data to graphing service
1003
Produce time-series data from received events over time
1004
Incorporate time-series events into visualization of infrastructure as they occur
1000

Fig. 10

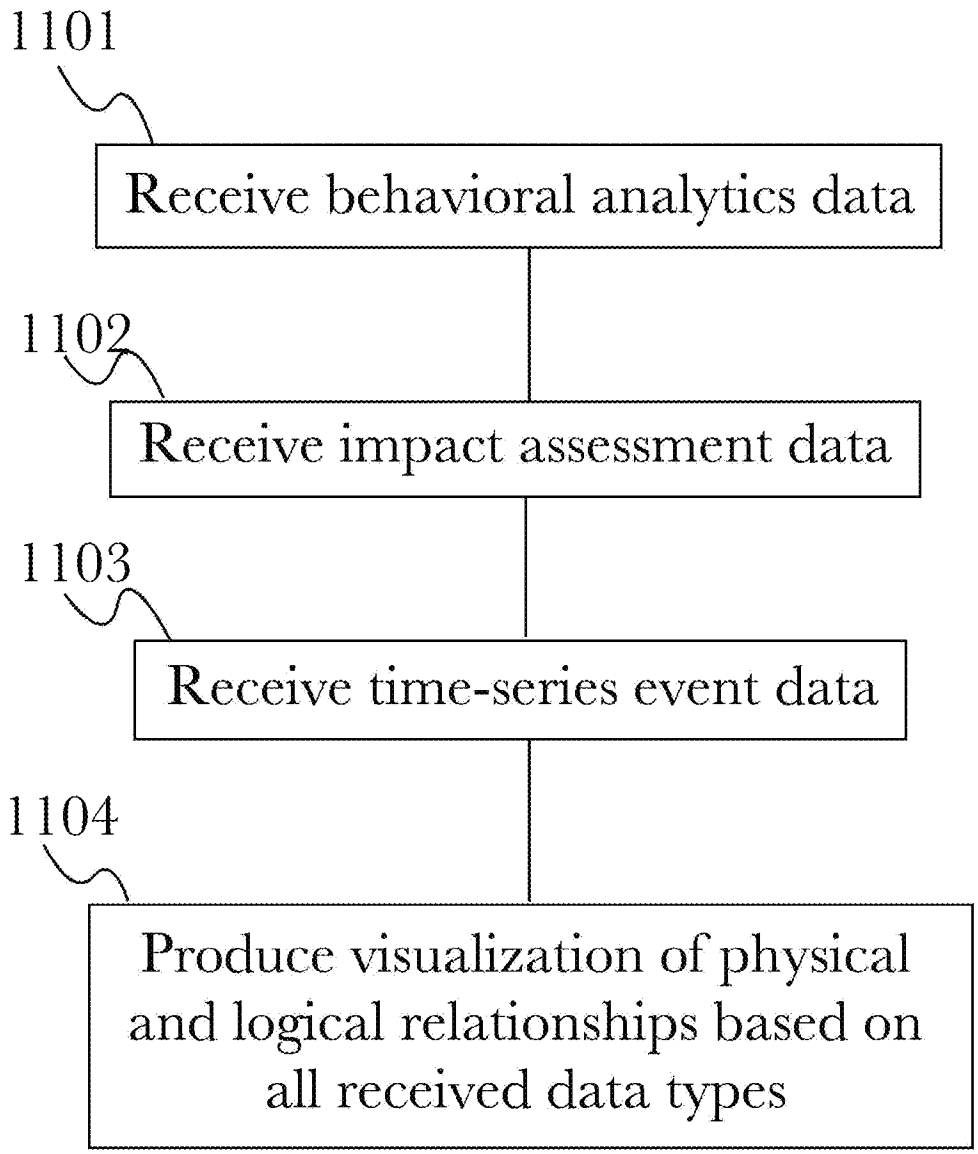
1101
Receive behavioral analytics data
1102
Receive impact assessment data
1103
Receive time-series event data
1104
Produce visualization of physical and logical relationships based on all received data types

Fig. 11

1100

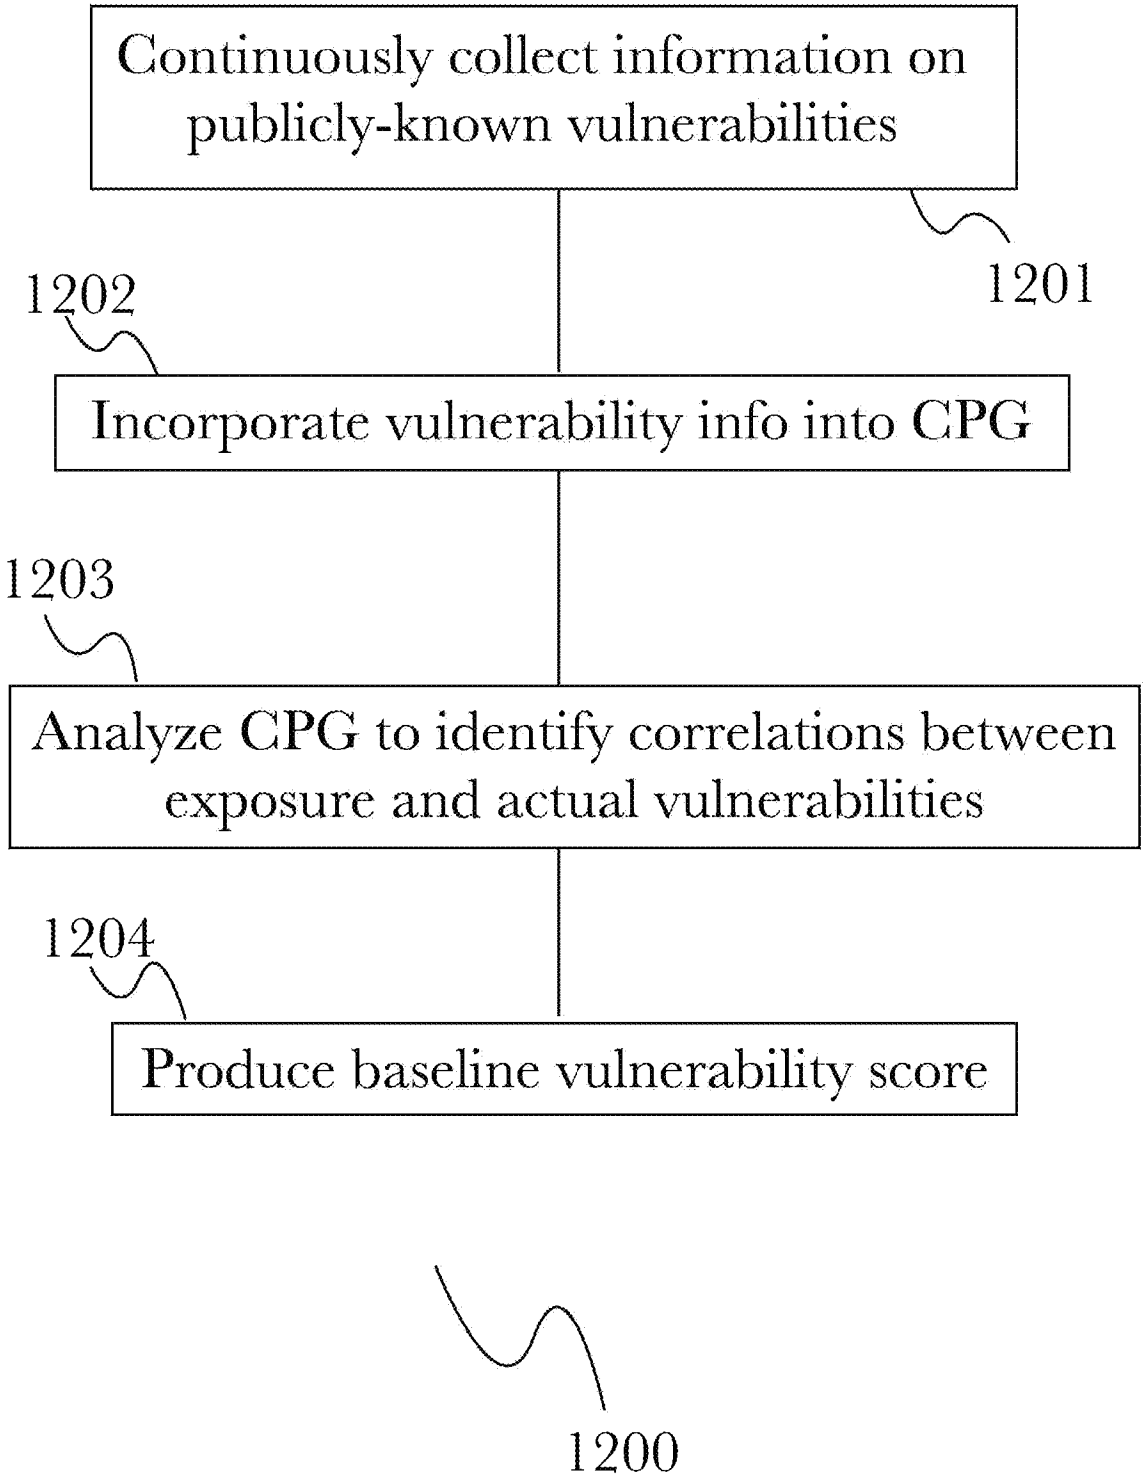
Continuously collect information on publicly-known vulnerabilities
1201
1202
Incorporate vulnerability info into CPG
1203
Analyze CPG to identify correlations between exposure and actual vulnerabilities
1204
Produce baseline vulnerability score
1200

Fig. 12

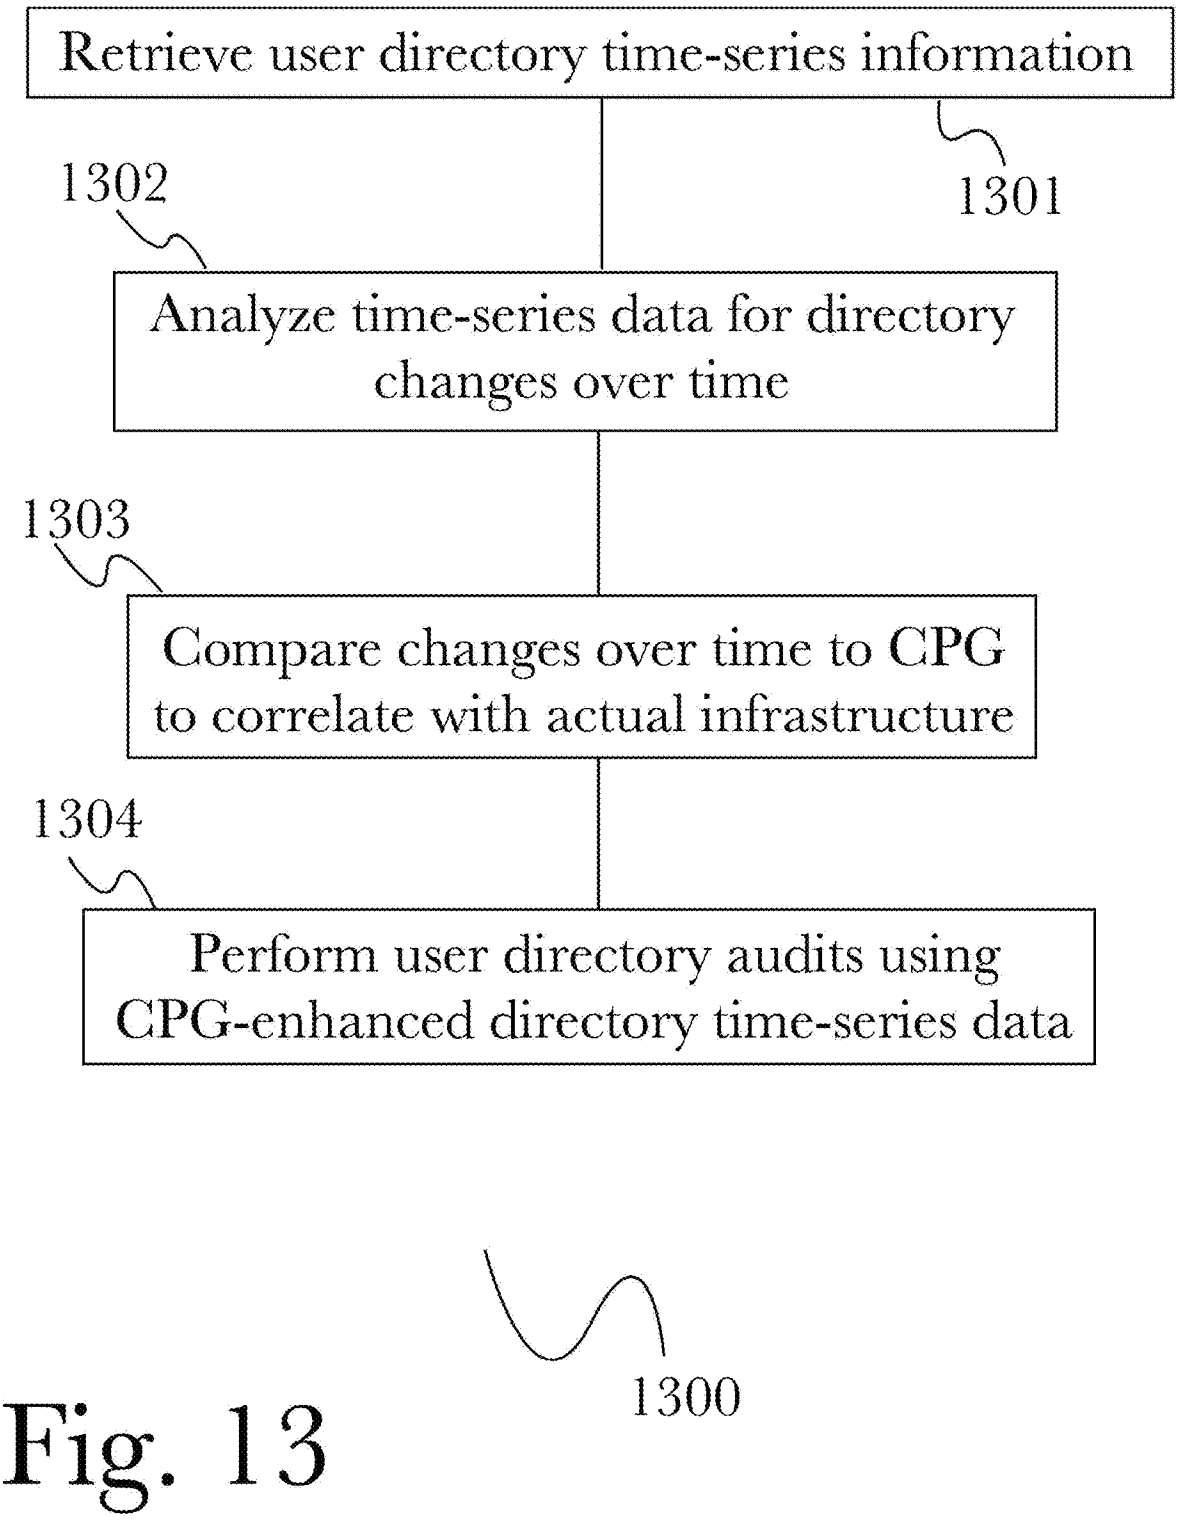
Retrieve user directory time-series information
1301
1302
Analyze time-series data for directory changes over time
1303
Compare changes over time to CPG to correlate with actual infrastructure
1304
Perform user directory audits using CPG-enhanced directory time-series data
1300

Fig. 13

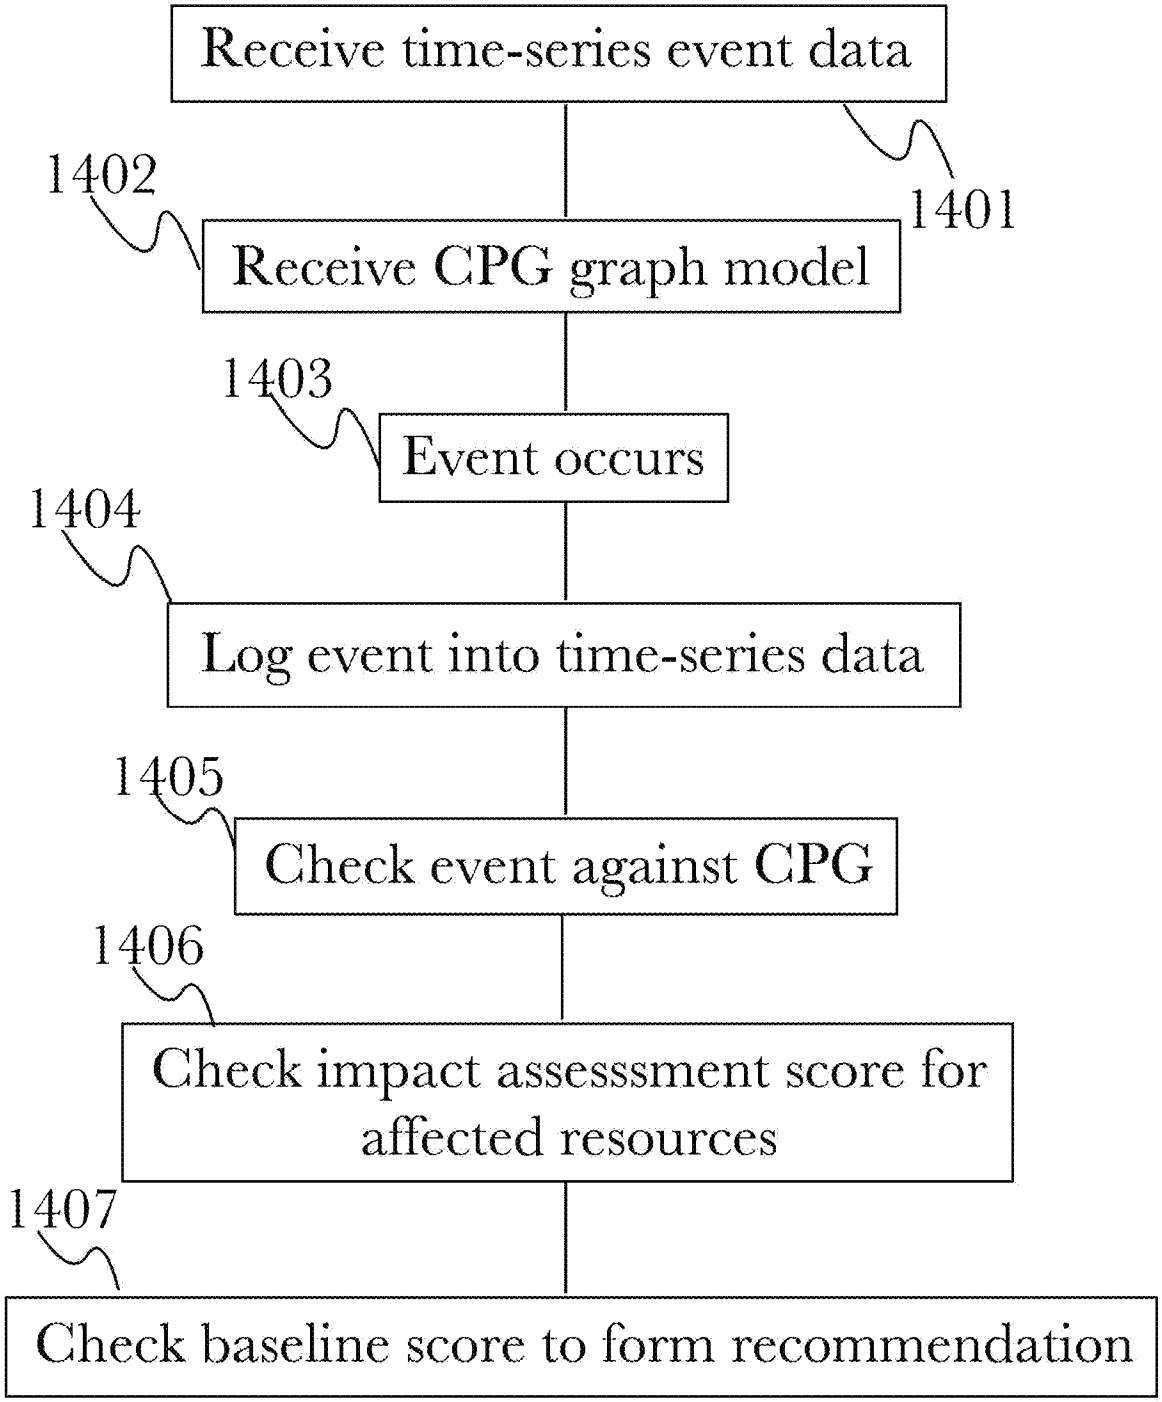
Receive time-series event data
1401
1402
Receive CPG graph model
1403
Event occurs
1404
Log event into time-series data
1405
Check event against CPG
1406
Check impact assesssment score for affected resources
1407
Check baseline score to form recommendation

Fig. 14

1400

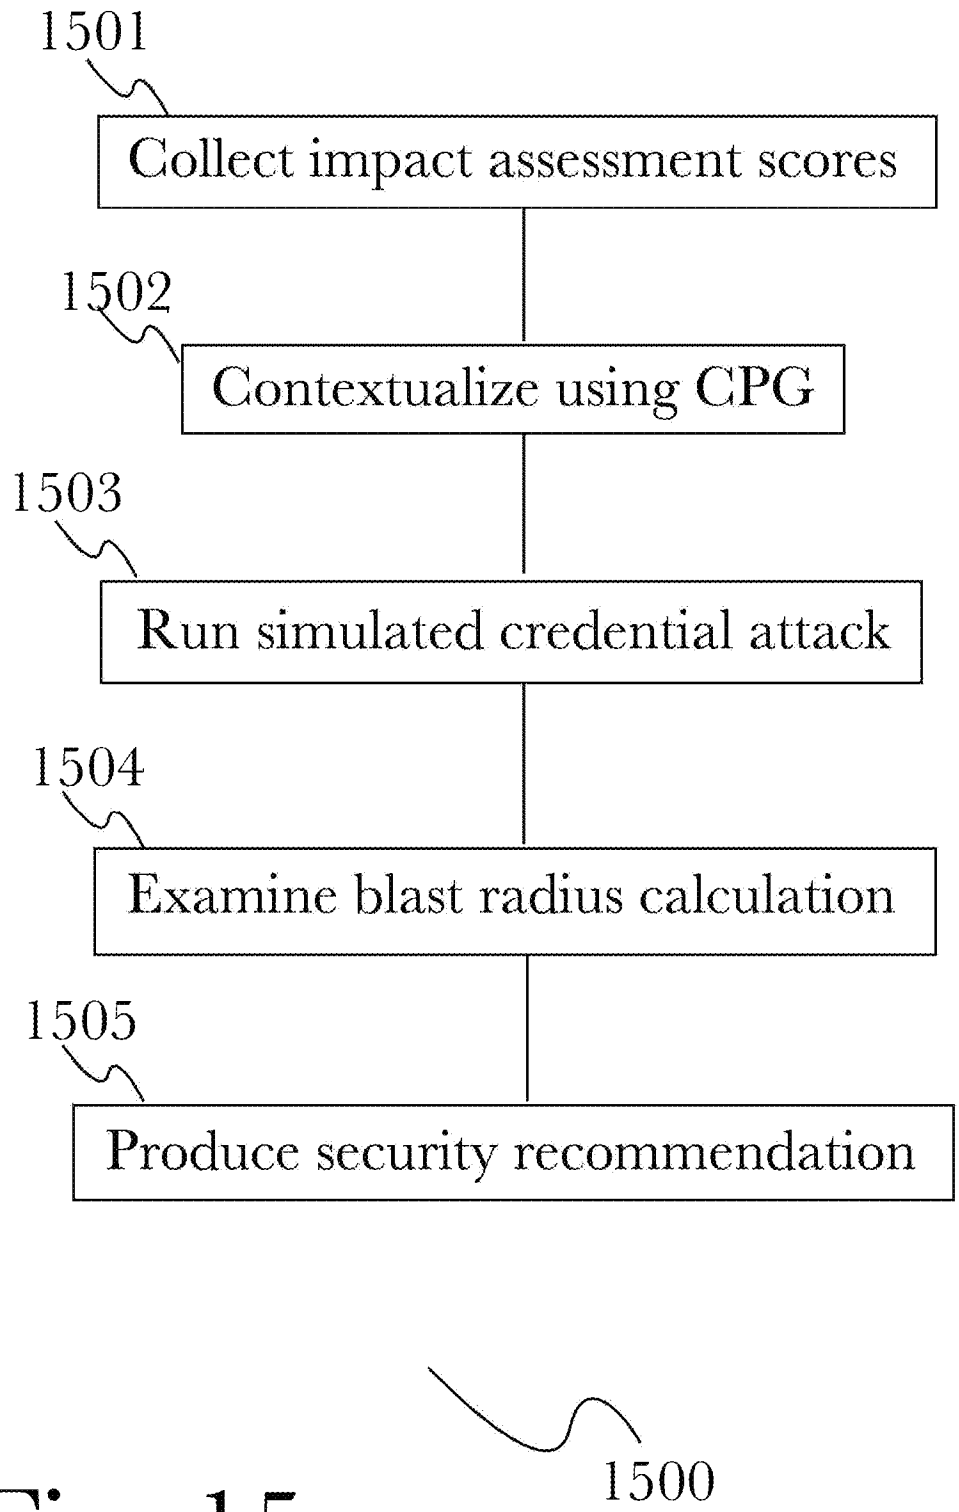
1501
Collect impact assessment scores
1502
Contextualize using CPG
1503
Run simulated credential attack
1504
Examine blast radius calculation
1505
Produce security recommendation
1500

Fig. 15

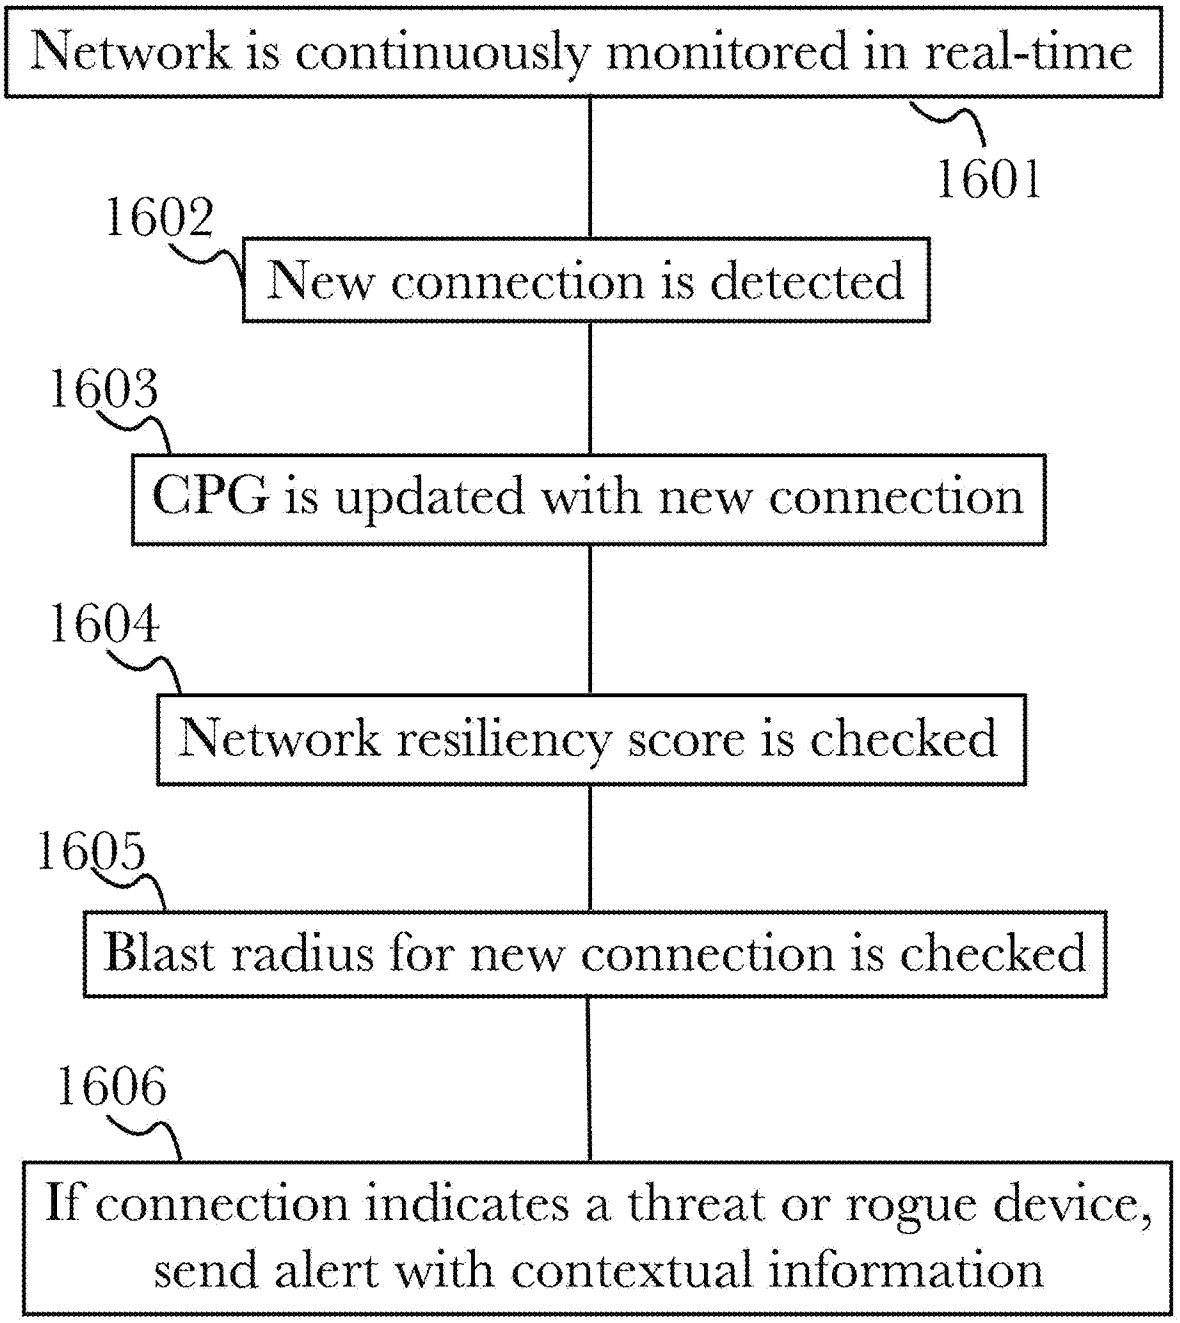
Network is continuously monitored in real-time
1601
1602
New connection is detected
1603
CPG is updated with new connection
1604
Network resiliency score is checked
1605
Blast radius for new connection is checked
1606
If connection indicates a threat or rogue device, send alert with contextual information

Fig. 16

1600

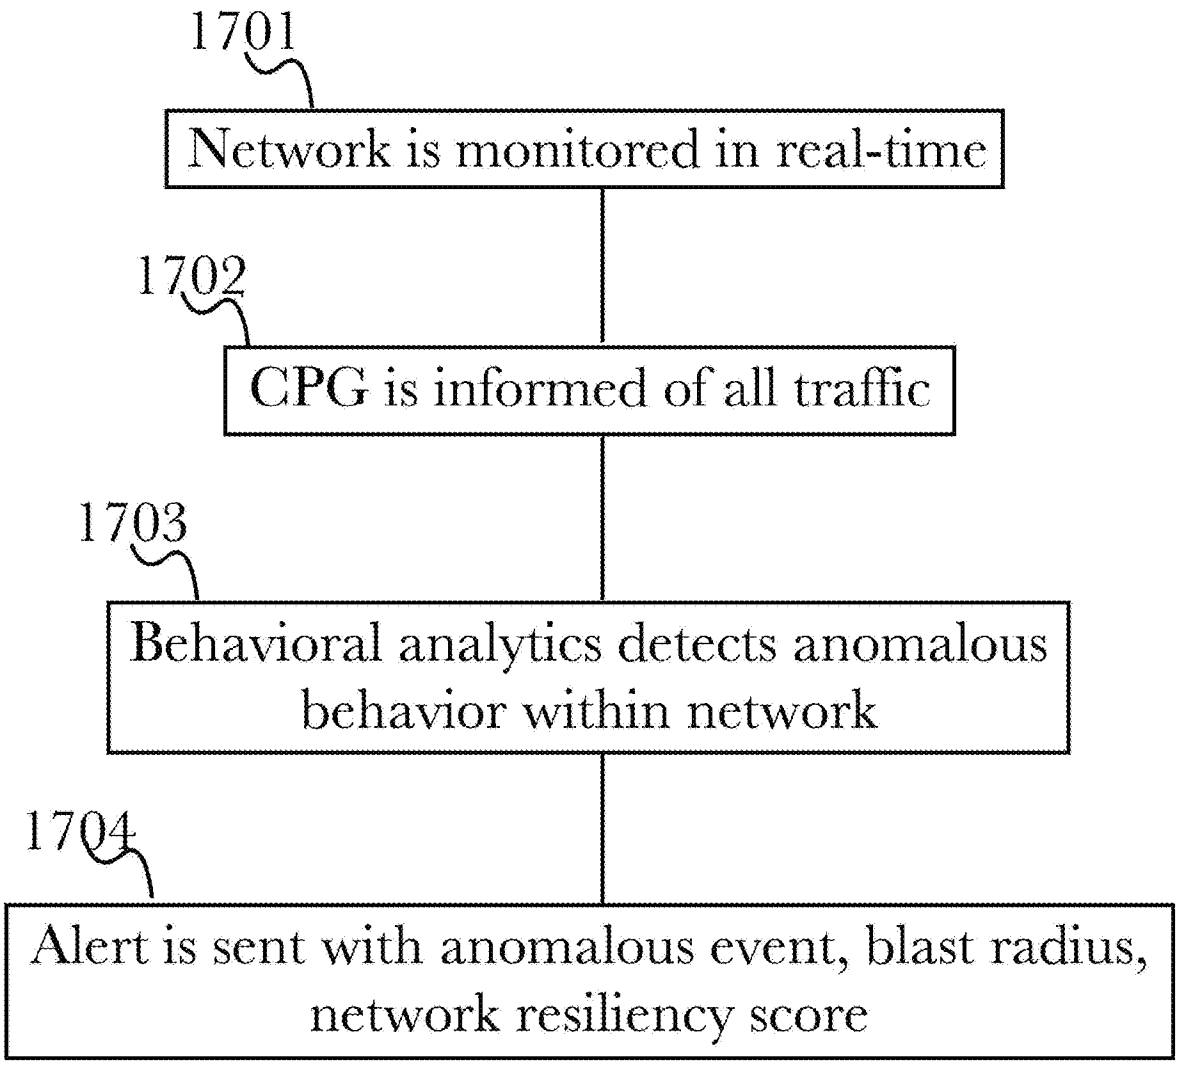
1701
Network is monitored in real-time
1702
CPG is informed of all traffic
1703
Behavioral analytics detects anomalous behavior within network
1704
Alert is sent with anomalous event, blast radius, network resiliency score

Fig. 17

1700

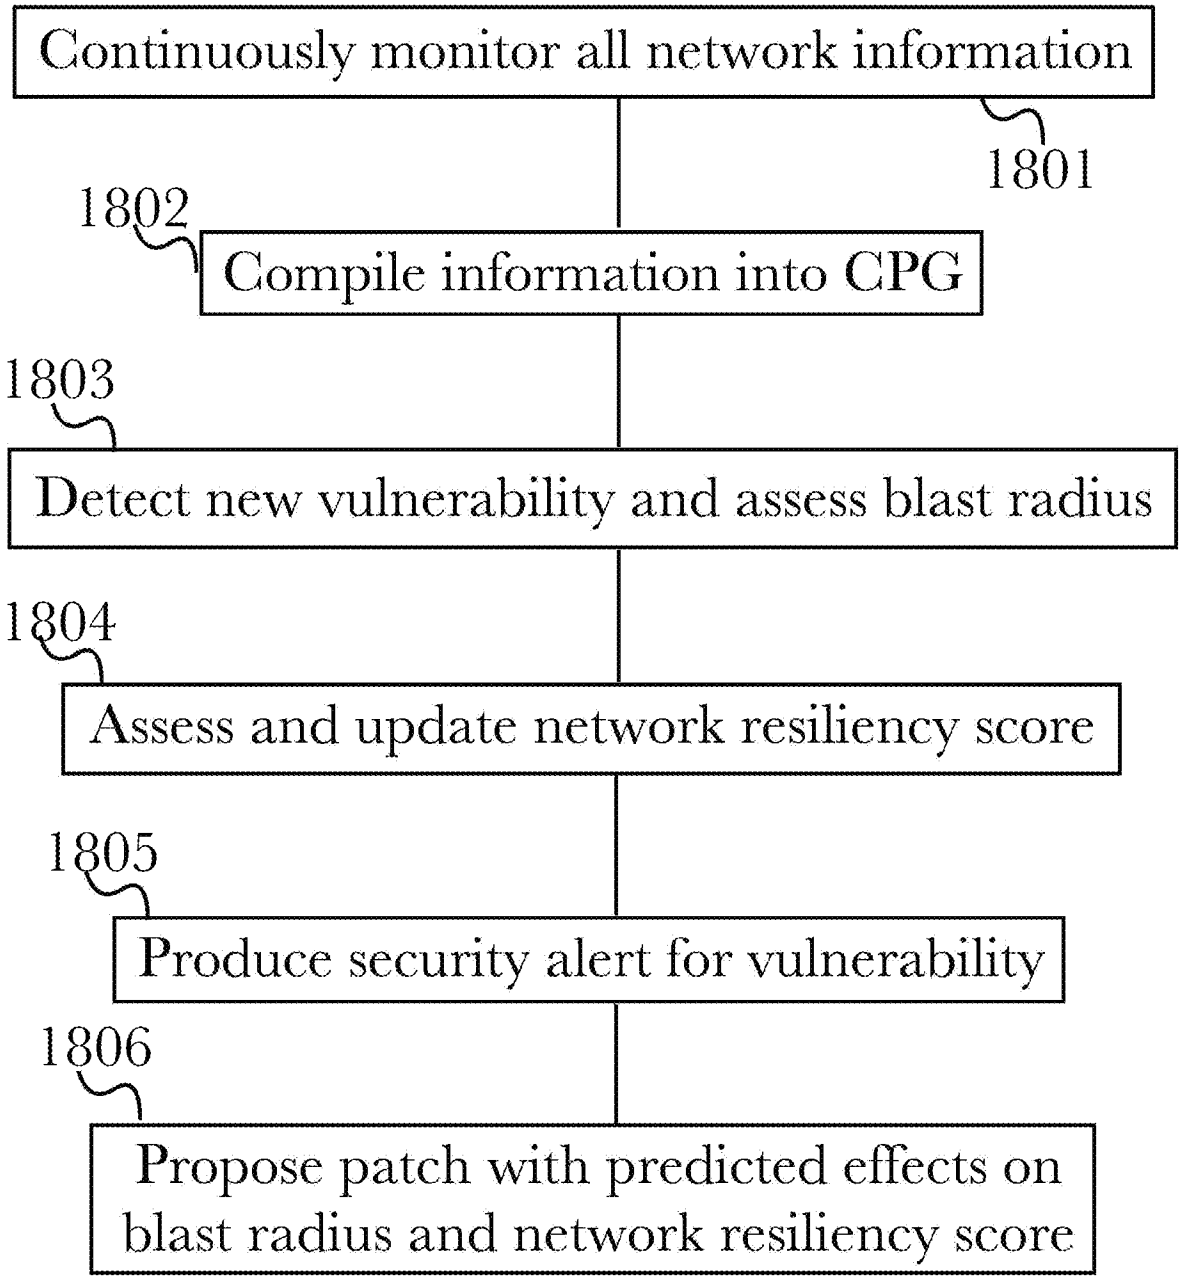
Continuously monitor all network information
1801
1802
Compile information into CPG
1803
Detect new vulnerability and assess blast radius
1804
Assess and update network resiliency score
1805
Produce security alert for vulnerability
1806
Propose patch with predicted effects on blast radius and network resiliency score

Fig. 18

1800

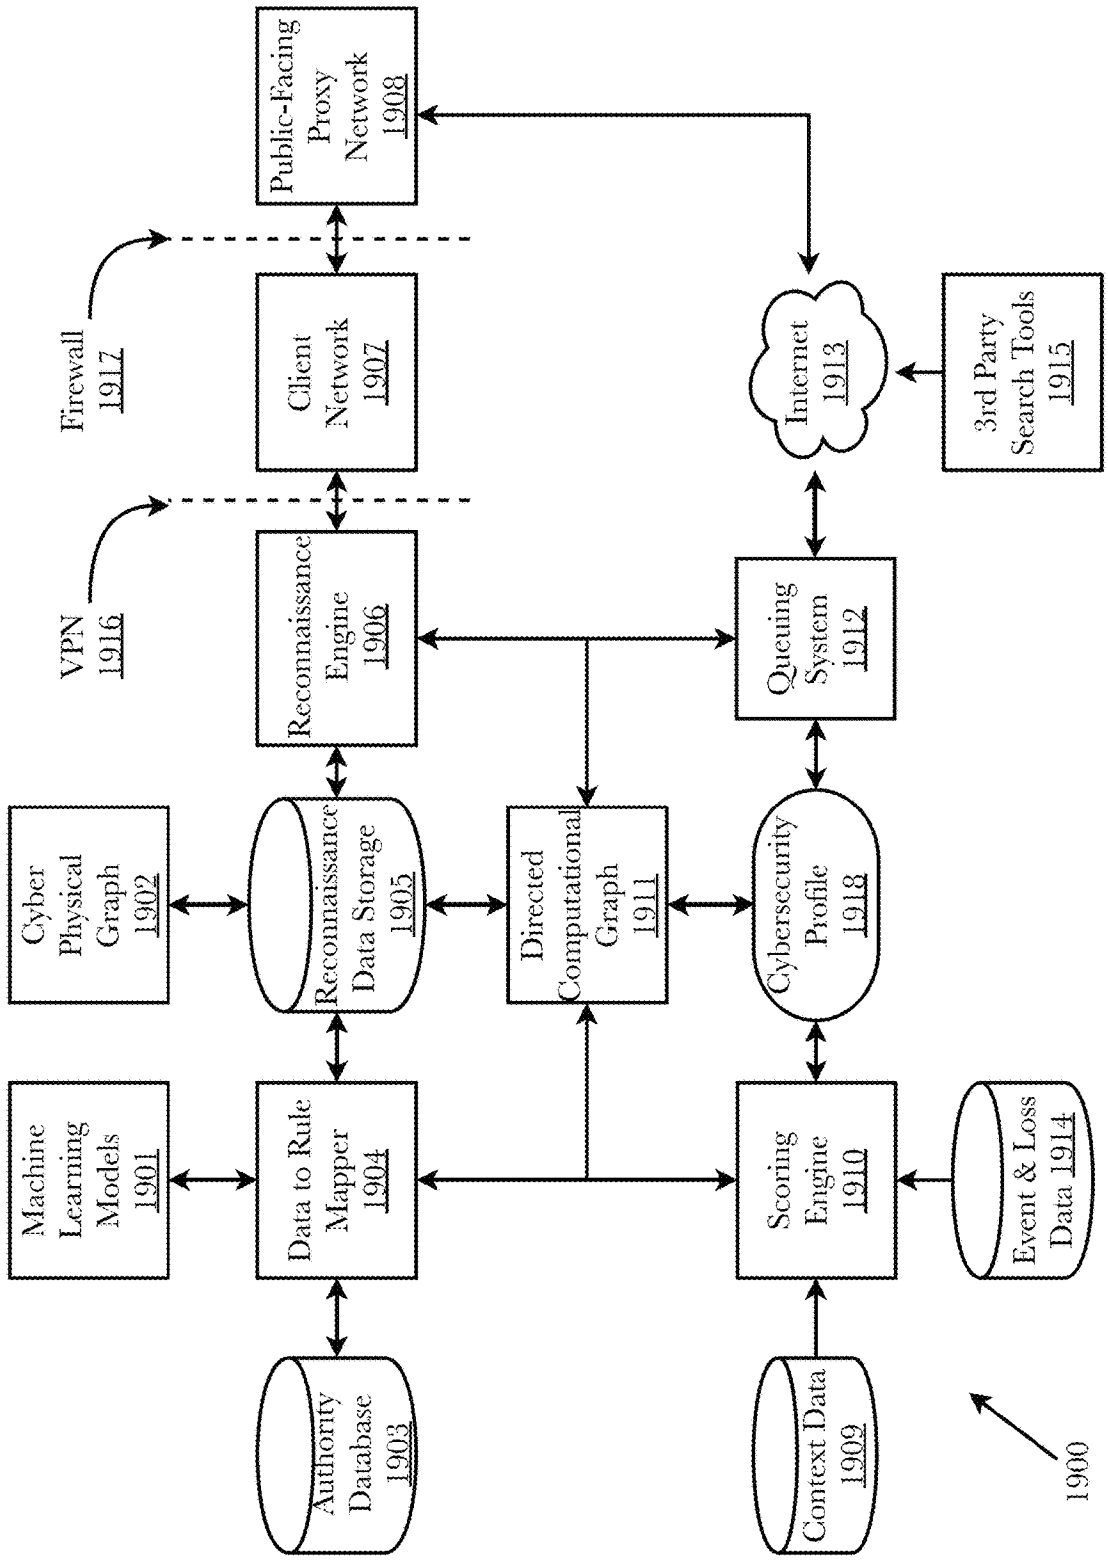
Machine Learning Models 1901
Cyber Physical Graph 1902
Authority Database 1903
Data to Rule Mapper 1904
Reconnaissance Data Storage 1905
Reconnaissance Engine 1906
Client Network 1907
Public-Facing Proxy Network 1908
Directed Computational Graph 1911
Context Data 1909
Scoring Engine 1910
Cybersecurity Profile 1918
Queuing System 1912
Event & Loss Data 1914
Internet 1913
3rd Party Search Tools 1915
VPN 1916
Firewall 1917
1900

Fig. 19

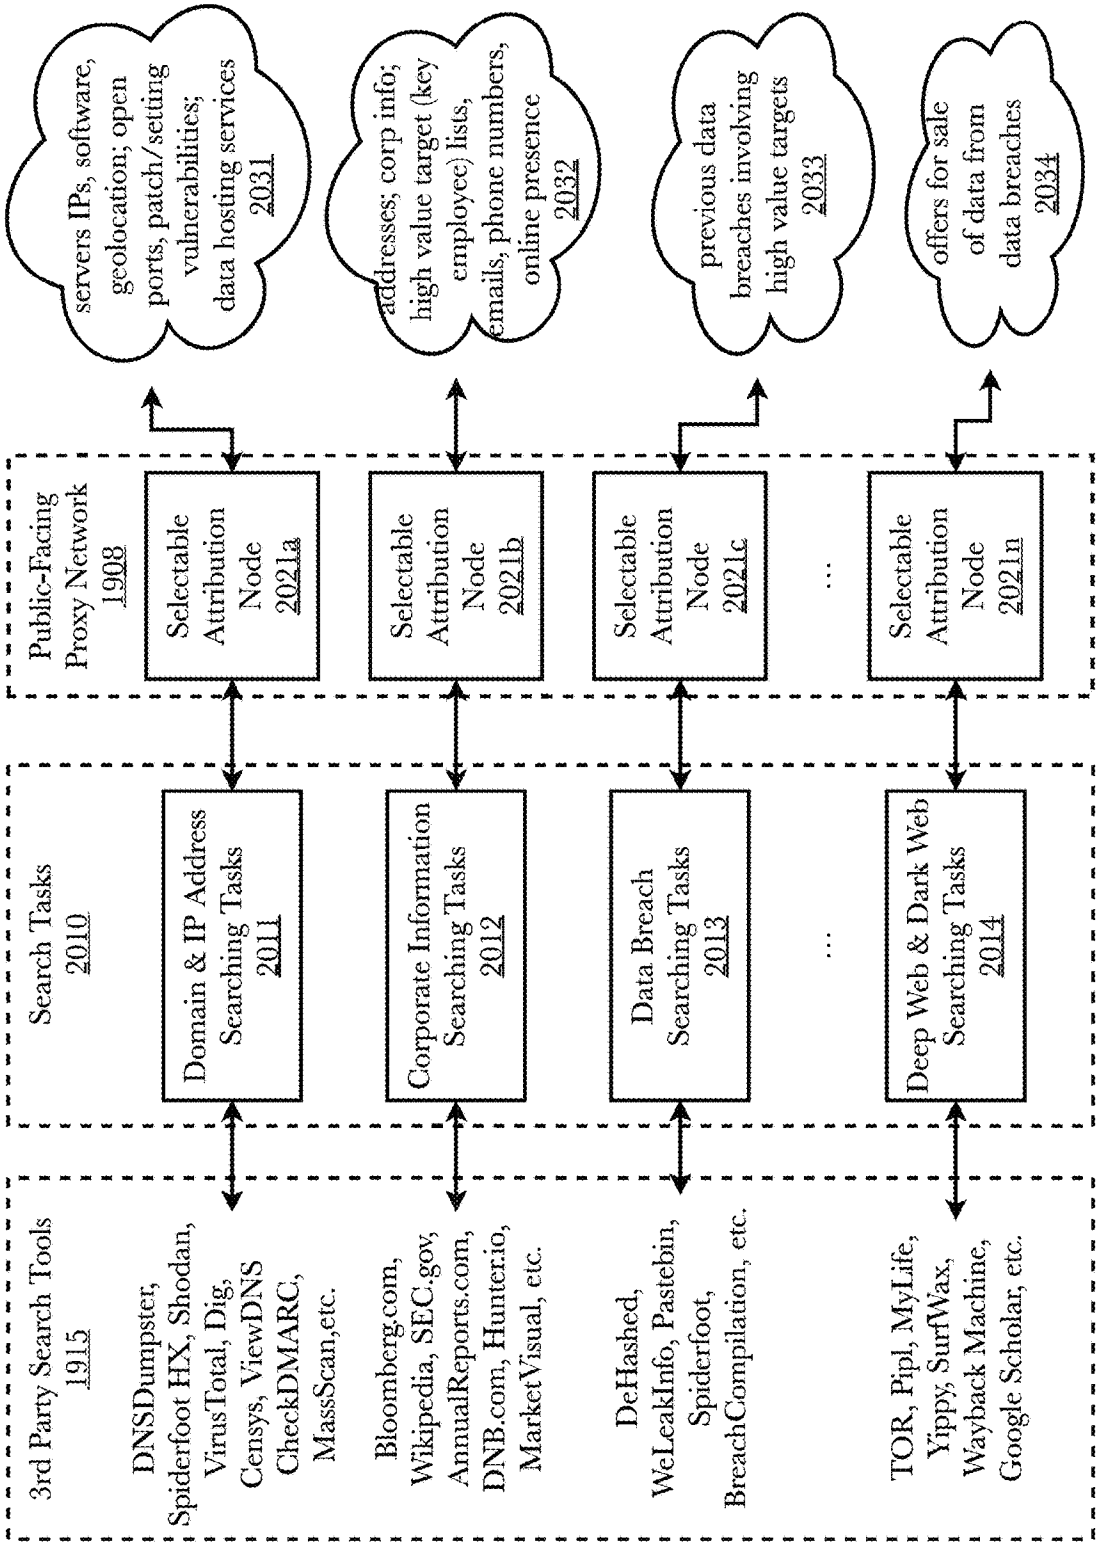
3rd Party Search Tools 1915
DNSDumpster, Spiderfoot HX, Shodan, VirusTotal, Dig, Censys, ViewDNS CheckDMARC, MassScan,etc.
Bloomberg.com, Wikipedia, SEC.gov, AnnualReports.com, DNB.com, Hunter.io, MarketVisual, etc.
DeHashed, WeLeakInfo, Pastebin, Spiderfoot, BreachCompilation, etc.
TOR, Pipl, MyLife, Yippy, SurfWax, Wayback Machine, Google Scholar, etc.
Search Tasks 2010
Domain & IP Address Searching Tasks 2011
Corporate Information Searching Tasks 2012
Data Breach Searching Tasks 2013
...
Deep Web & Dark Web Searching Tasks 2014
Public-Facing Proxy Network 1908
Selectable Attribution Node 2021a
Selectable Attribution Node 2021b
Selectable Attribution Node 2021c
...
Selectable Attribution Node 2021n
servers IPs, software, geolocation; open ports, patch/setting vulnerabilities; data hosting services 2031
addresses; corp info; high value target (key employee) lists, emails, phone numbers, online presence 2032
previous data breaches involving high value targets 2033
offers for sale of data from data breaches 2034

Fig. 20

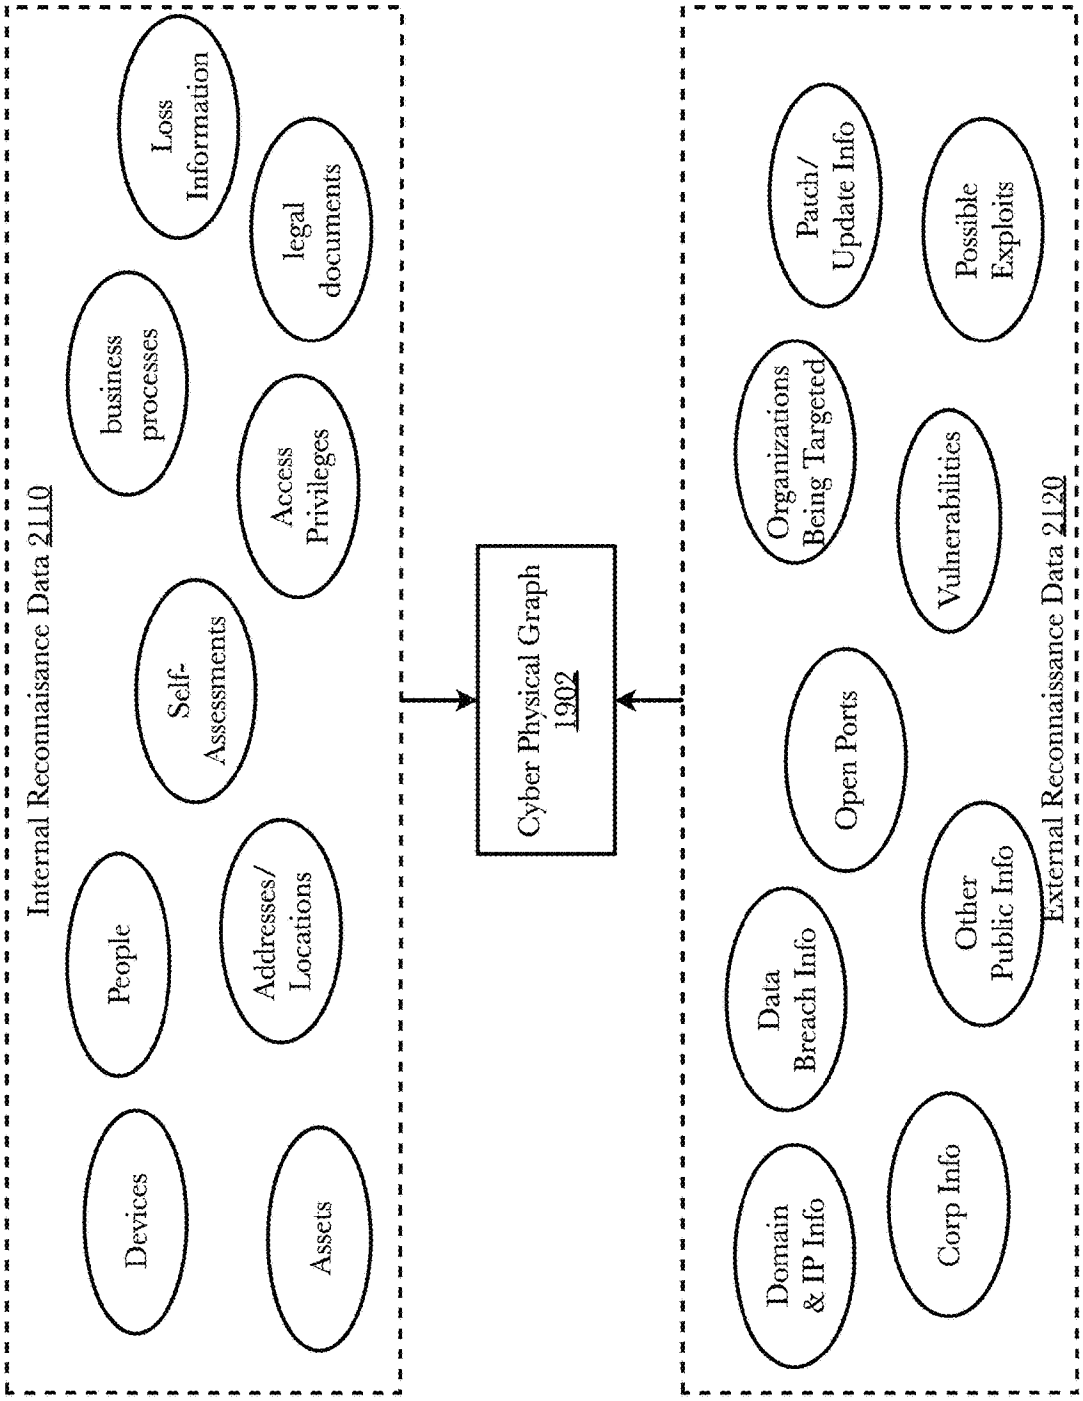
Internal Reconnaisance Data 2110
Devices
People
Assets
Addresses/ Locations
Self-Assessments
Access Privileges
business processes
legal documents
Loss Information
Cyber Physical Graph 1902
External Reconnaissance Data 2120
Domain & IP Info
Data Breach Info
Corp Info
Open Ports
Other Public Info
Organizations Being Targeted
Vulnerabilities
Patch/ Update Info
Possible Exploits

Fig. 21

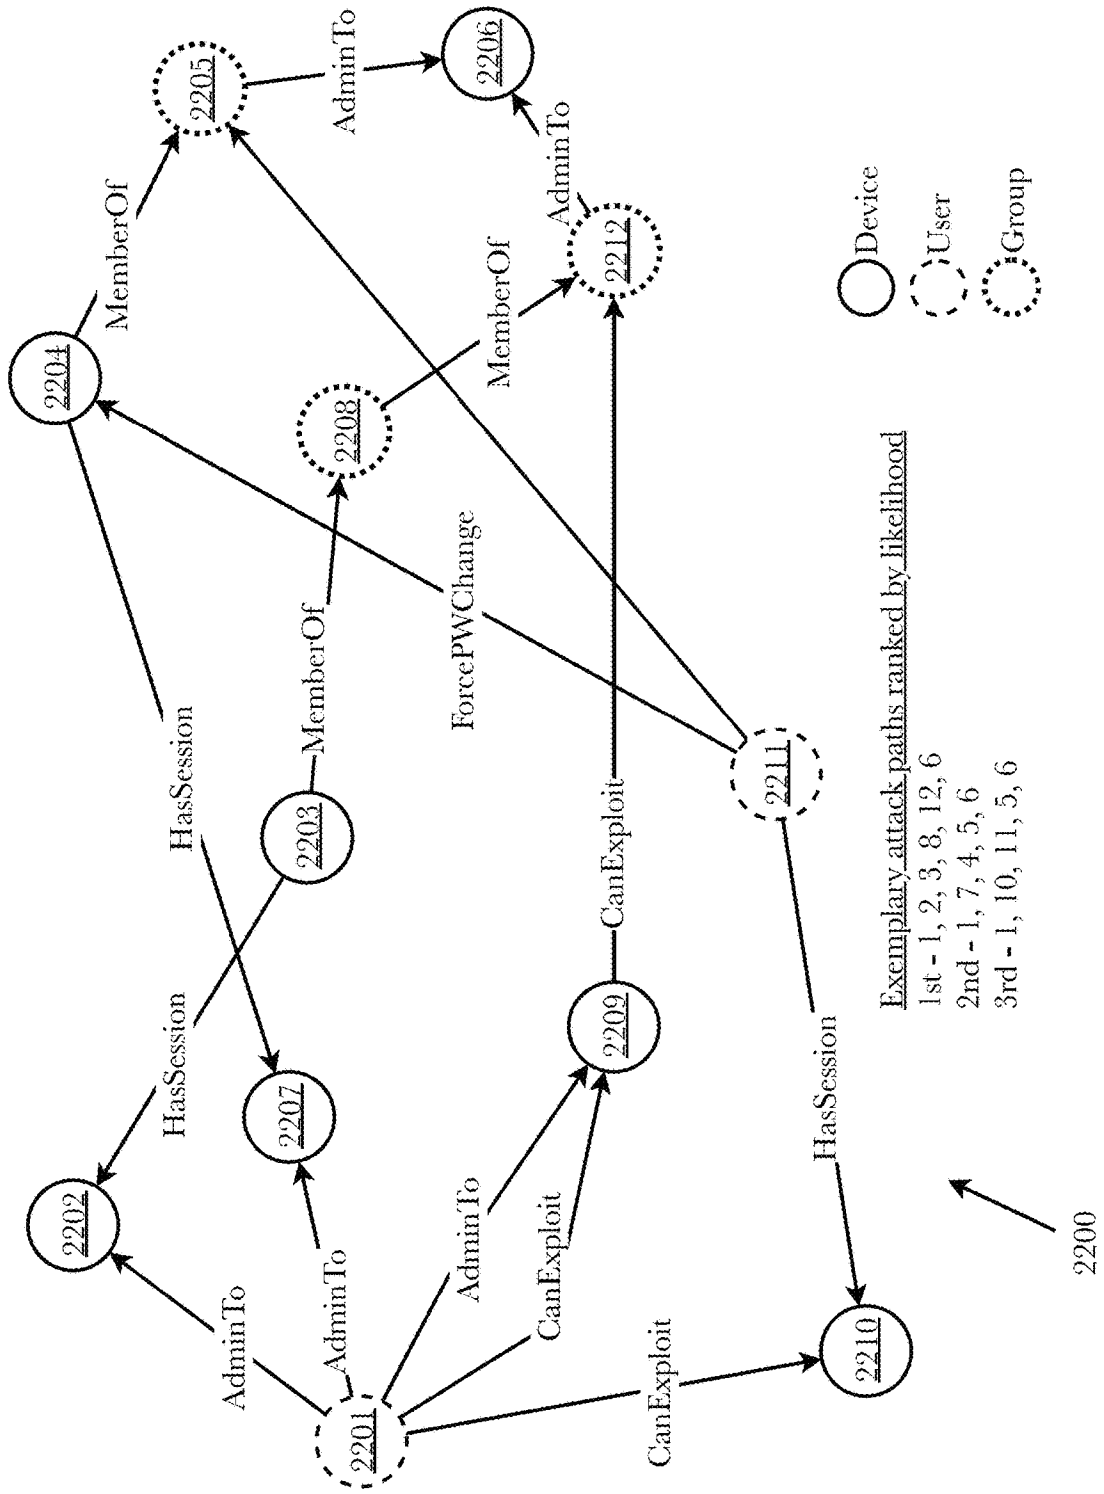
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
AdminTo
AdminTo
AdminTo
AdminTo
HasSession
HasSession
MemberOf
MemberOf
MemberOf
ForcePWChange
CanExploit
CanExploit
CanExploit
HasSession
Device
User
Group
Exemplary attack paths ranked by likelihood
1st - 1, 2, 3, 8, 12, 6
2nd - 1, 7, 4, 5, 6
3rd - 1, 10, 11, 5, 6
2200

Fig. 22

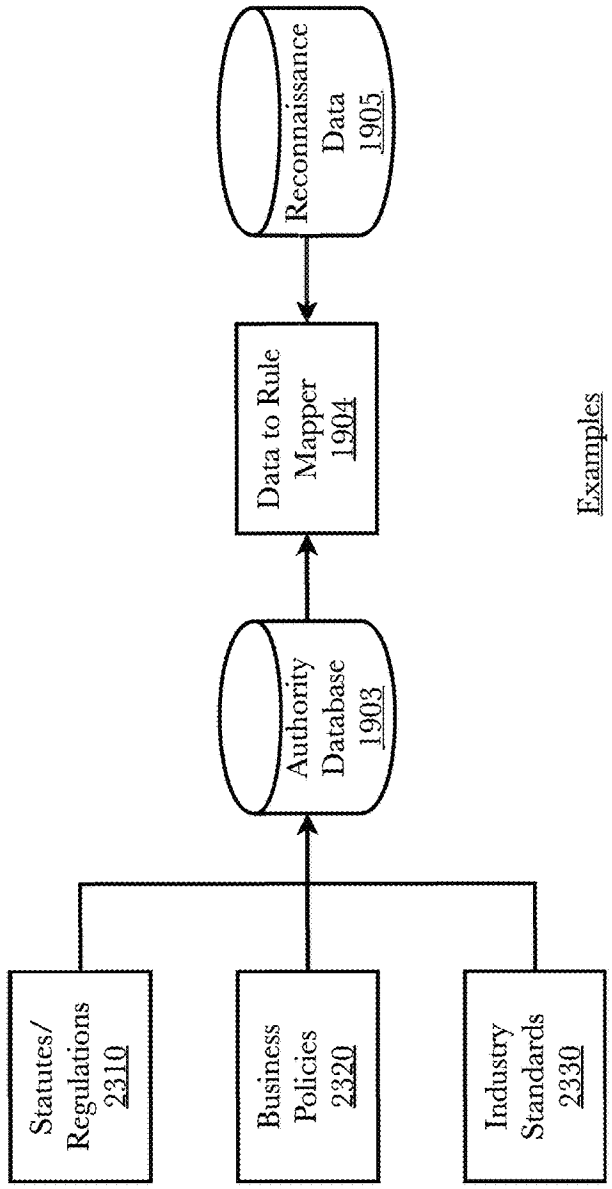

Examples

Authority	Rule	Data Point Received	Tagged with Impact
2311 GDPR Article 33, paragraph 1	Notify cognizant authority within 72 hours breach	Data breach reported online	Possible fines
2321 Corporate Policy	Do not access corp systems using personal computers	Employee account accessed using non-business computer	Possible data theft; security breach
2331 Industry Standards	No open ports accessible outside of network perimeter	Open ports detected	Possible data loss

Fig. 23

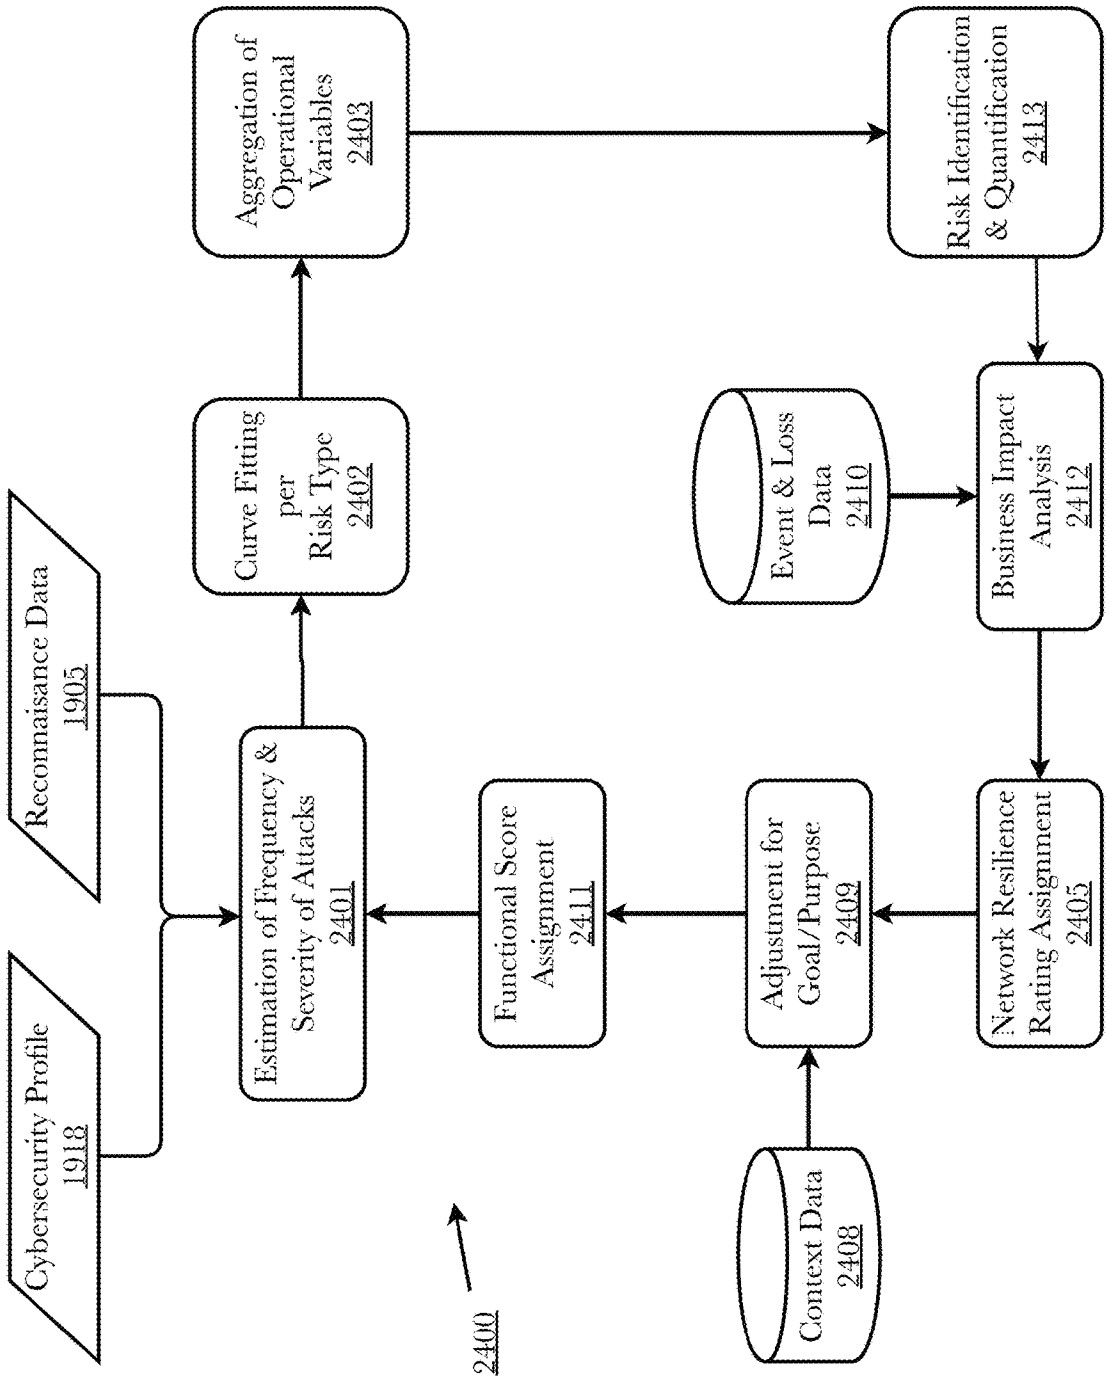
Cybersecurity Profile
1918
Reconnaisance Data
1905
Estimation of Frequency & Severity of Attacks
2401
Curve Fitting per Risk Type
2402
Aggregation of Operational Variables
2403
Risk Identification & Quantification
2413
Business Impact Analysis
2412
Event & Loss Data
2410
Network Resilience Rating Assignment
2405
Adjustment for Goal/Purpose
2409
Context Data
2408
Functional Score Assignment
2411
2400

Fig. 24

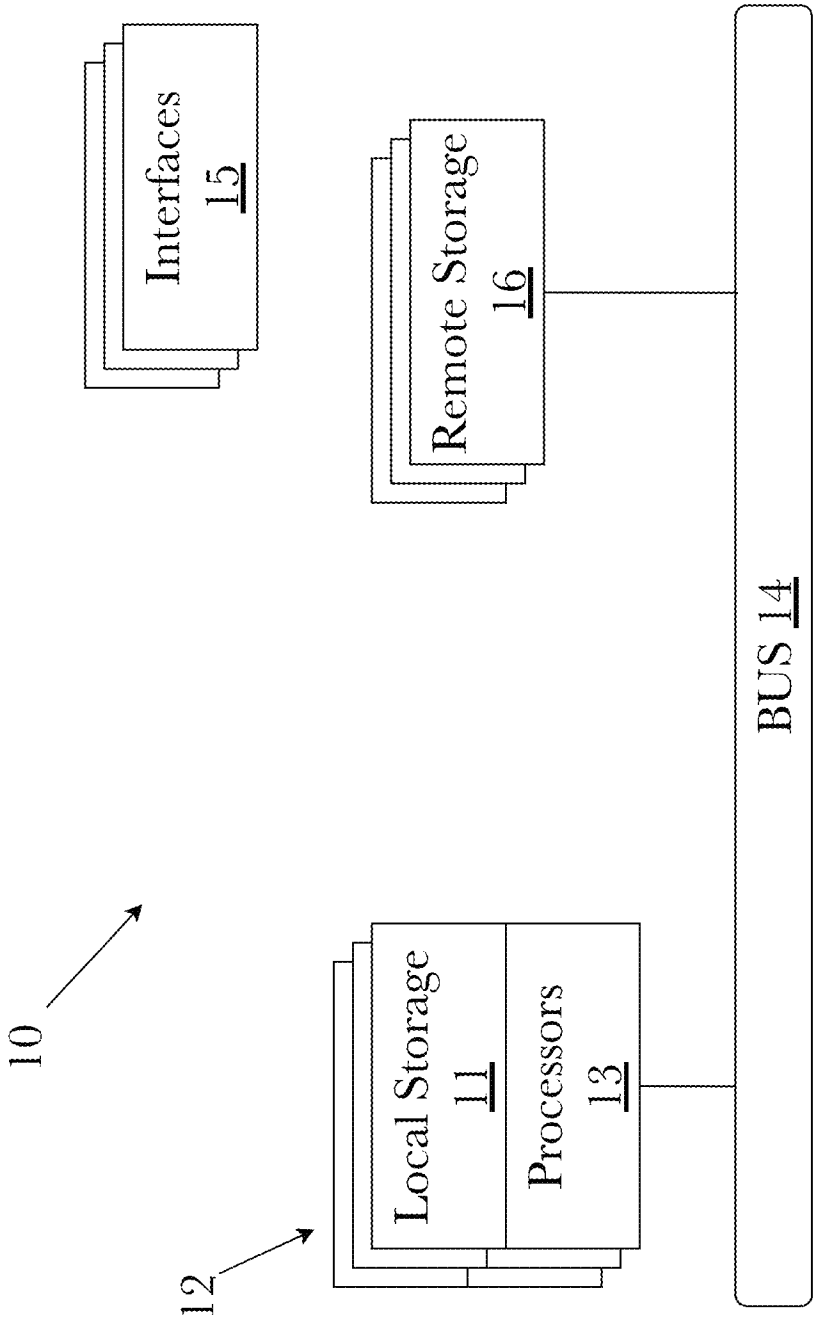
10
12
Local Storage
11
Processors
13
Interfaces
15
Remote Storage
16
BUS 14

Fig. 25

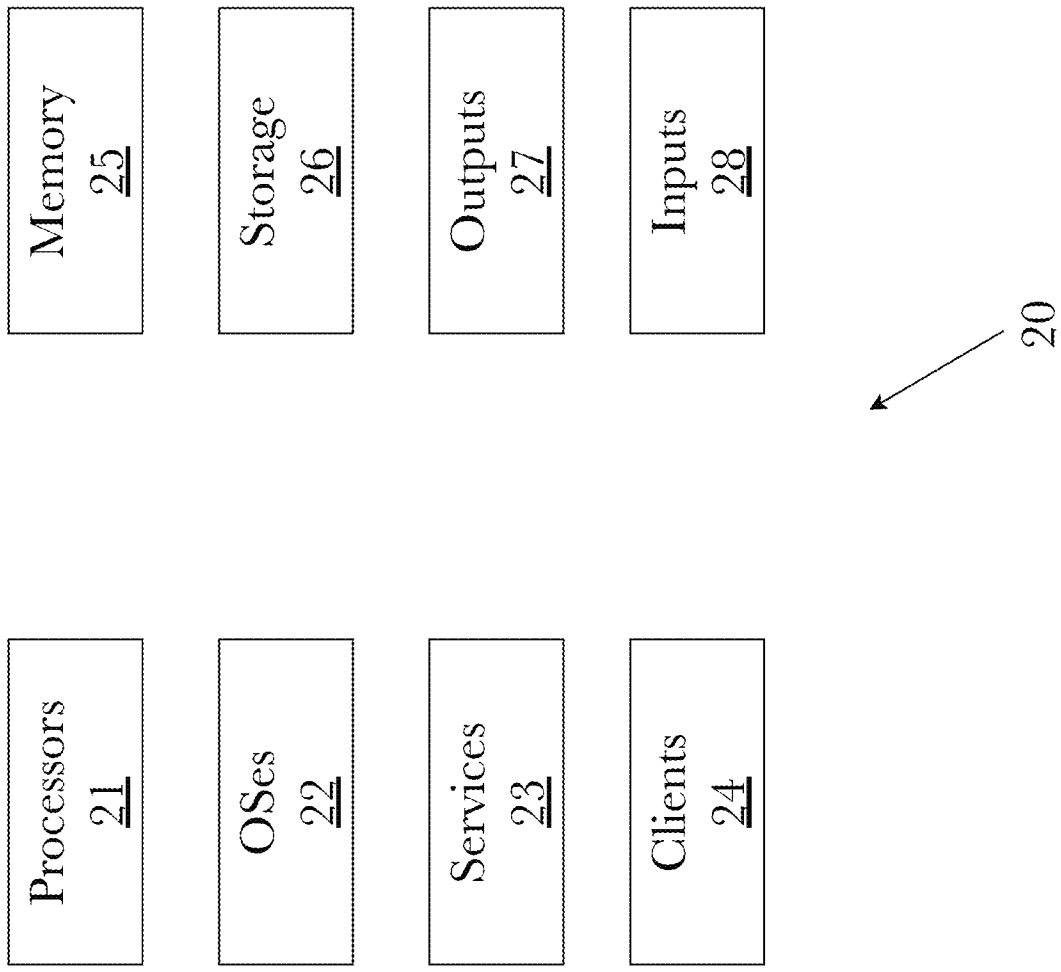
Memory
25
Storage
26
Outputs
27
Inputs
28
20
Processors
21
OSes
22
Services
23
Clients
24

Fig. 26

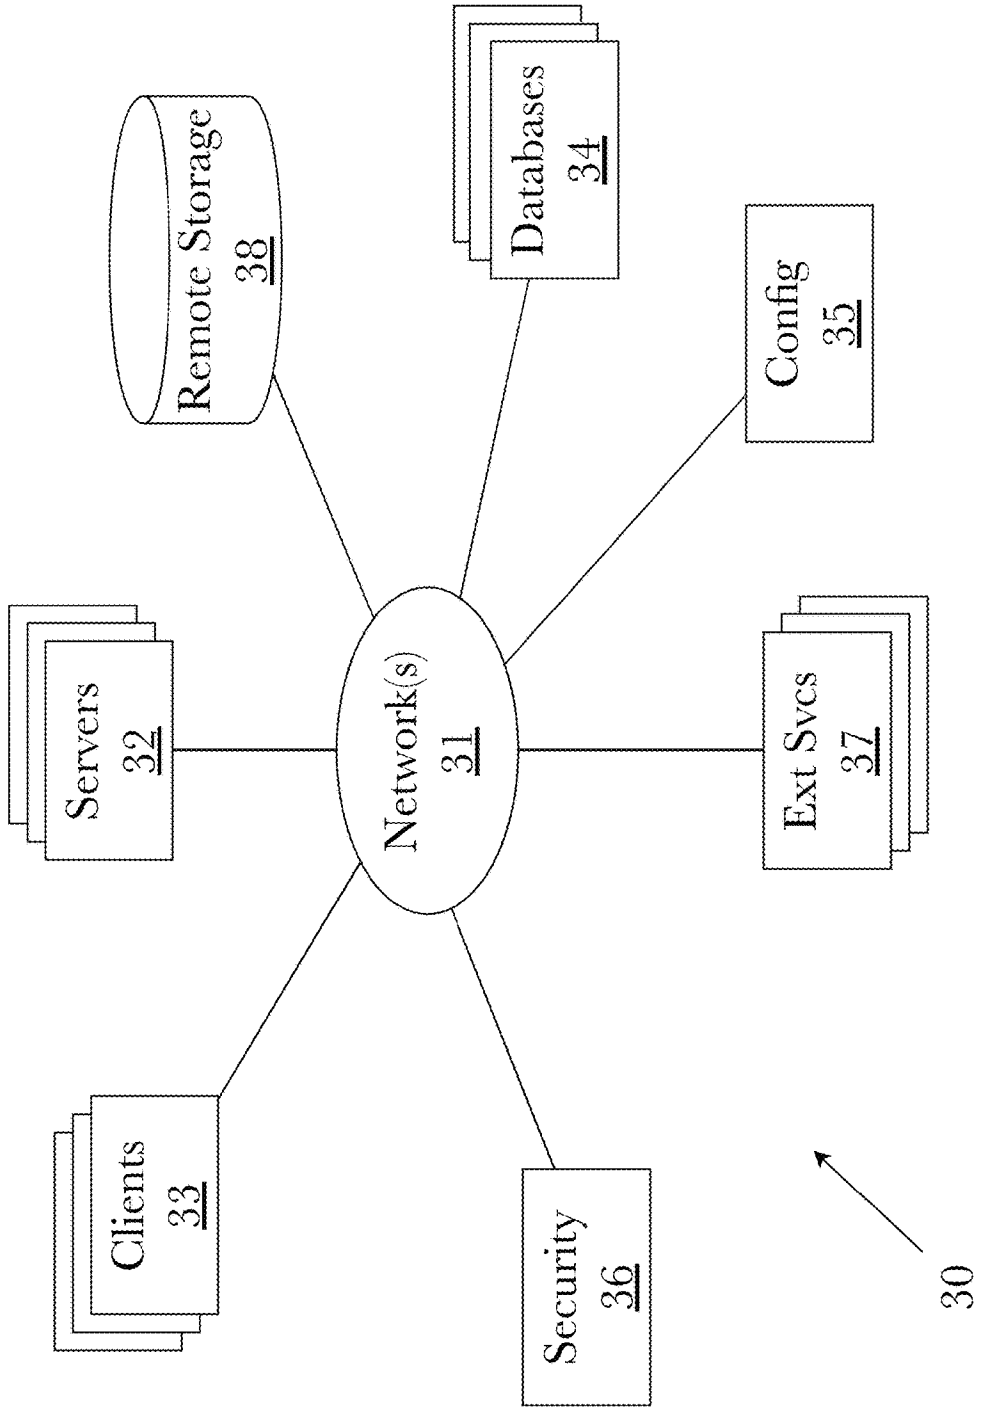
Remote Storage
38
Servers
32
Clients
33
Network(s)
31
Databases
34
Config
35
Ext Svcs
37
Security
36
30

Fig. 27

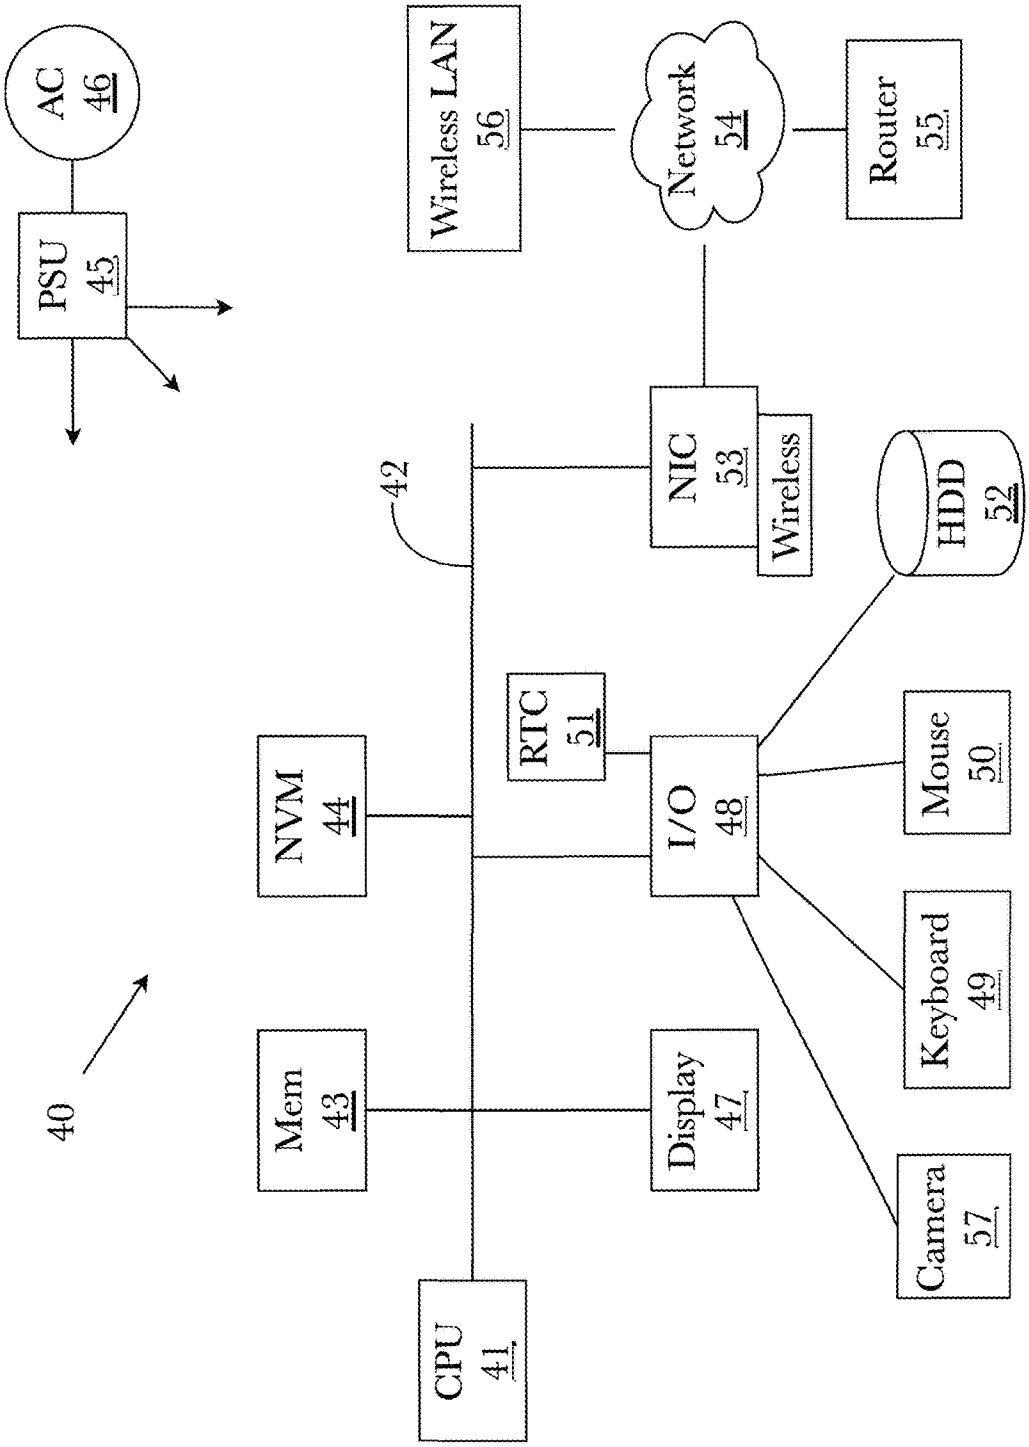
AC
46
PSU
45
Wireless LAN
56
Network
54
Router
55
40
42
NIC
53
Wireless
HDD
52
NVM
44
RTC
51
I/O
48
Mouse
50
Mem
43
Keyboard
49
Display
47
Camera
57
CPU
41

Fig. 28

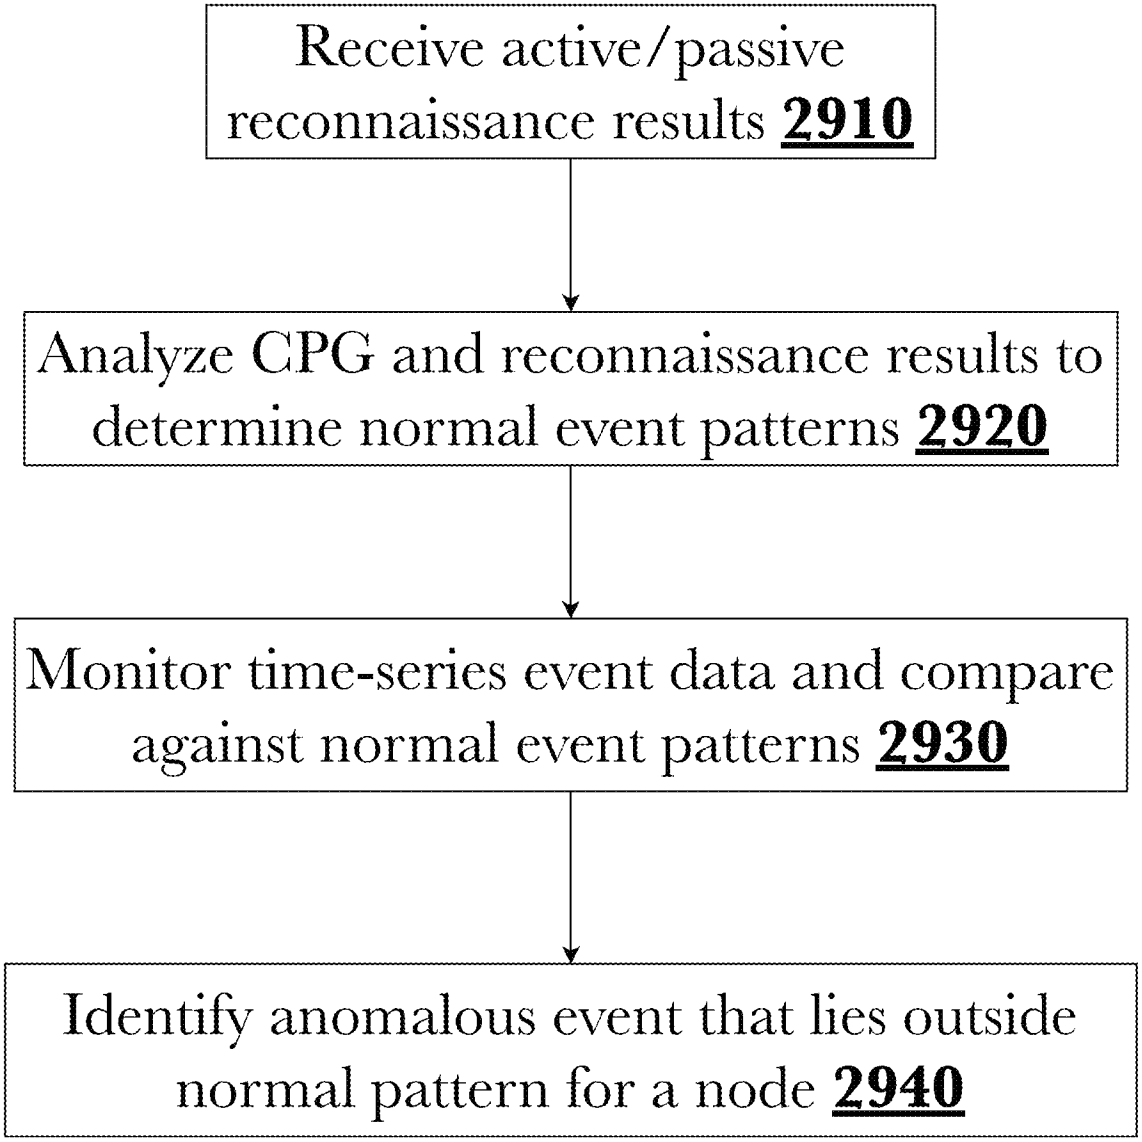
Receive active/passive reconnaissance results 2910
Analyze CPG and reconnaissance results to determine normal event patterns 2920
Monitor time-series event data and compare against normal event patterns 2930
Identify anomalous event that lies outside normal pattern for a node 2940

Fig. 29

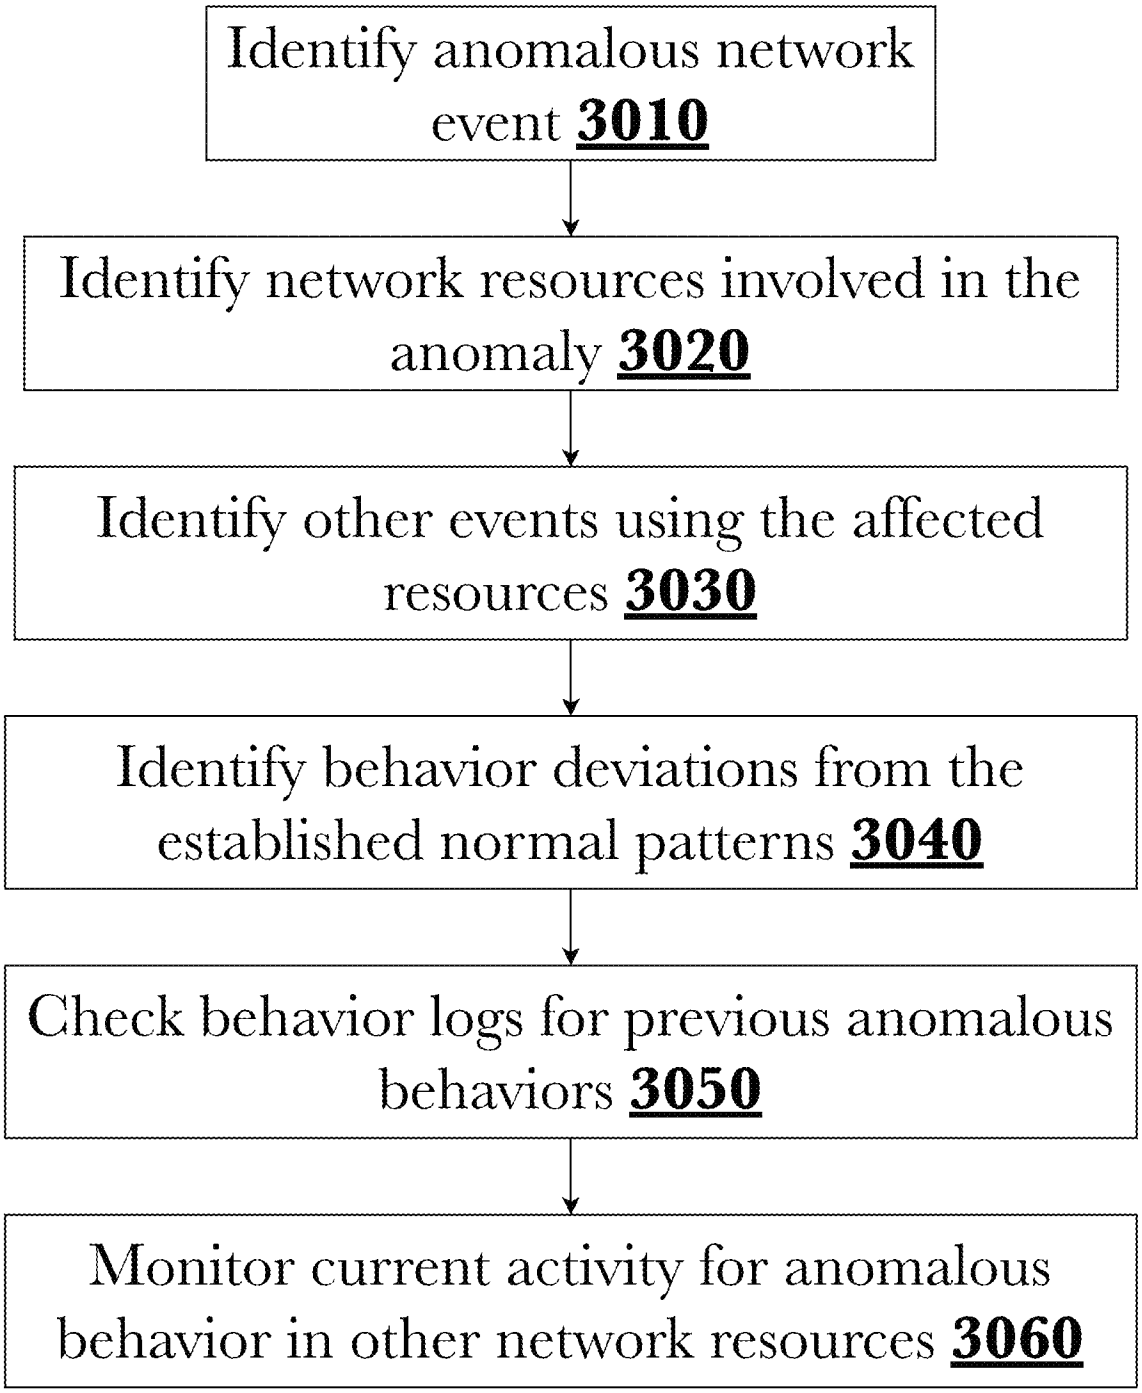
Identify anomalous network event 3010
Identify network resources involved in the anomaly 3020
Identify other events using the affected resources 3030
Identify behavior deviations from the established normal patterns 3040
Check behavior logs for previous anomalous behaviors 3050
Monitor current activity for anomalous behavior in other network resources 3060

Fig. 30

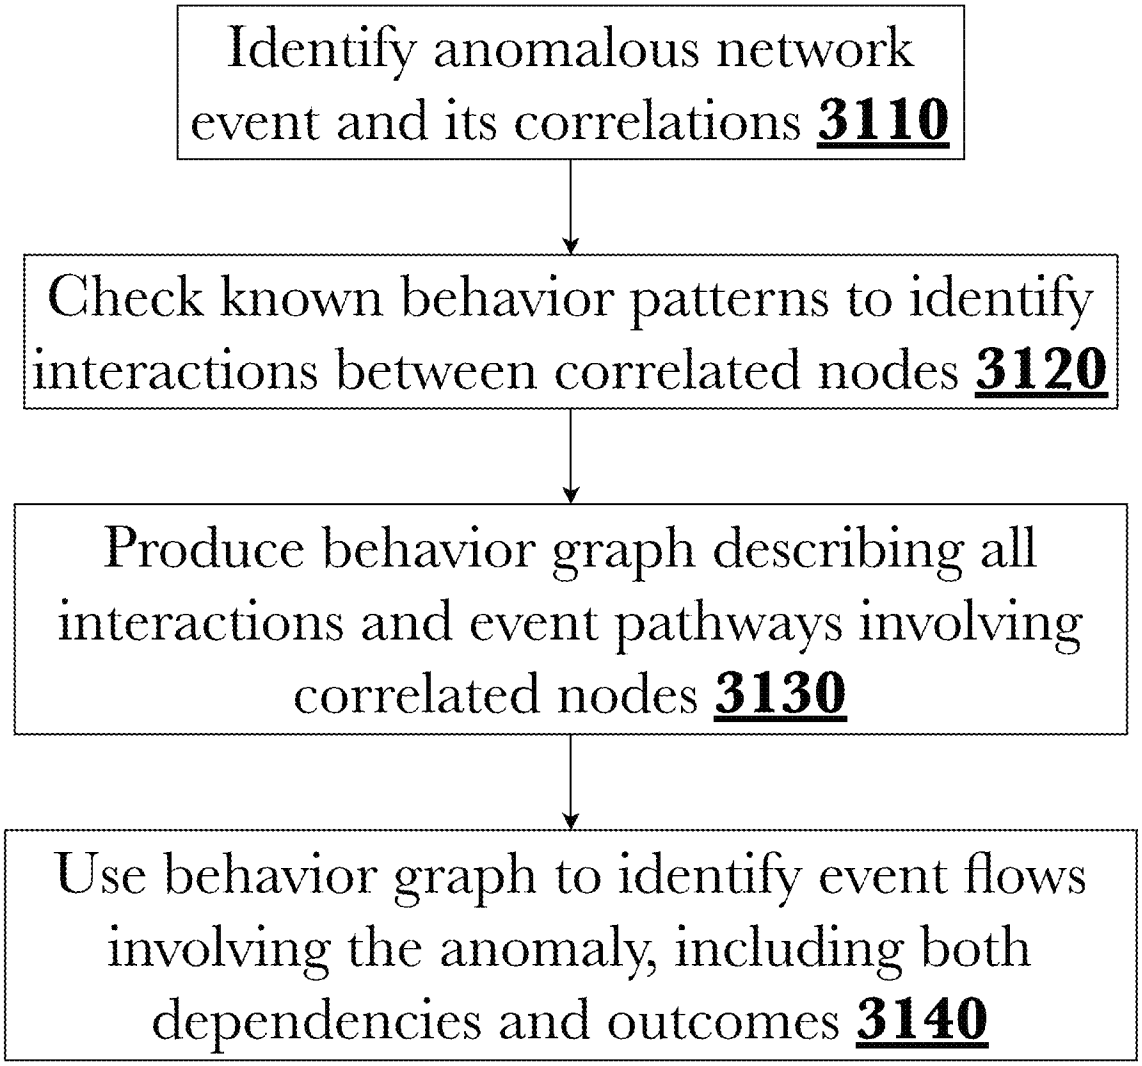
Identify anomalous network event and its correlations 3110
Check known behavior patterns to identify interactions between correlated nodes 3120
Produce behavior graph describing all interactions and event pathways involving correlated nodes 3130
Use behavior graph to identify event flows involving the anomaly, including both dependencies and outcomes 3140

Fig. 31

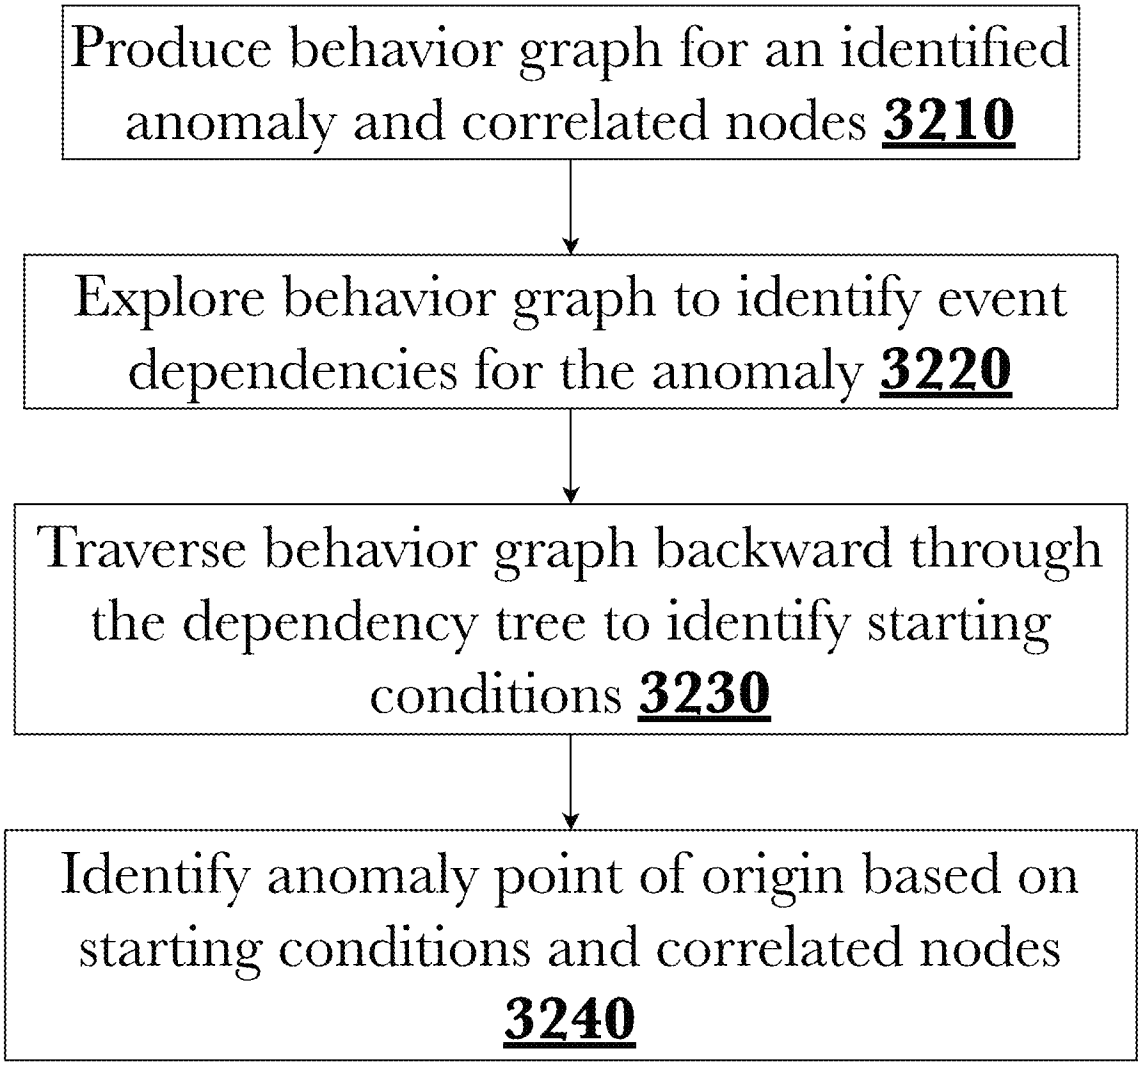
Produce behavior graph for an identified anomaly and correlated nodes 3210
Explore behavior graph to identify event dependencies for the anomaly 3220
Traverse behavior graph backward through the dependency tree to identify starting conditions 3230
Identify anomaly point of origin based on starting conditions and correlated nodes 3240

Fig. 32

CORRELATING NETWORK EVENT ANOMALIES USING ACTIVE AND PASSIVE EXTERNAL RECONNAISSANCE TO IDENTIFY ATTACK INFORMATION

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation of U.S. patent application Ser. No. 17/237,346, titled "CORRELATING NETWORK EVENT ANOMALIES USING ACTIVE AND PASSIVE EXTERNAL RECONNAISSANCE TO IDENTIFY ATTACK INFORMATION", filed on Apr. 22, 2021, which is a continuation-in-part of U.S. patent application Ser. No. 16/777,270, titled "CYBERSECURITY PROFILING AND RATING USING ACTIVE AND PASSIVE EXTERNAL RECONNAISSANCE", filed on Jan. 30, 2020, now issued as U.S. Pat. No. 11,025,674 on Jun. 1, 2021, which is a continuation-in-part of U.S. patent application Ser. No. 16/720,383, titled "RATING ORGANIZATION CYBERSECURITY USING ACTIVE AND PASSIVE EXTERNAL RECONNAISSANCE", filed on Dec. 19, 2019, now issued as U.S. Pat. No. 10,944,795 on Mar. 9, 2021, which is a continuation of U.S. patent application Ser. No. 15/823,363, filed on Nov. 27, 2017, now issued as U.S. Pat. No. 10,560,483 on Feb. 11, 2020, which is a continuation-in-part of U.S. patent application Ser. No. 15/725,274, titled "APPLICATION OF ADVANCED CYBERSECURITY THREAT MITIGATION TO ROGUE DEVICES, PRIVILEGE ESCALATION, AND RISK-BASED VULNERABILITY AND PATCH MANAGEMENT", filed on Oct. 4, 2017, now issued as U.S. Pat. No. 10,609,079 on Mar. 31, 2020, which is a continuation-in-part of U.S. patent application Ser. No. 15/655,113, titled "ADVANCED CYBERSECURITY THREAT MITIGATION USING BEHAVIORAL AND DEEP ANALYTICS", filed on Jul. 20, 2017, now issued as U.S. Pat. No. 10,735,456 on Aug. 4, 2020, which is a continuation-in-part of U.S. patent application Ser. No. 15/616,427, titled "RAPID PREDICTIVE ANALYSIS OF VERY LARGE DATA SETS USING AN ACTOR-DRIVEN DISTRIBUTED COMPUTATIONAL GRAPH", filed on Jun. 7, 2017, which is a continuation-in-part of U.S. patent application Ser. No. 14/925,974, titled "RAPID PREDICTIVE ANALYSIS OF VERY LARGE DATA SETS USING THE DISTRIBUTED COMPUTATIONAL GRAPH", filed on Oct. 28, 2015; said application Ser. No. 15/655,113 is also a continuation-in-part of U.S. patent application Ser. No. 15/237,625, titled "DETECTION MITIGATION AND REMEDIATION OF CYBERATTACKS EMPLOYING AN ADVANCED CYBER-DECISION PLATFORM", filed on Aug. 15, 2016, now issued as U.S. Pat. No. 10,248,910 on Apr. 2, 2019, which is a continuation-in-part of U.S. patent application Ser. No. 15/206,195, titled "ACCURATE AND DETAILED MODELING OF SYSTEMS WITH LARGE COMPLEX DATA SETS USING A DISTRIBUTED SIMULATION ENGINE", filed on Jul. 8, 2016, which is a continuation-in-part of U.S. patent application Ser. No. 15/186,453, titled "SYSTEM FOR AUTOMATED CAPTURE AND ANALYSIS OF BUSINESS INFORMATION FOR RELIABLE BUSINESS VENTURE OUTCOME PREDICTION", filed on Jun. 18, 2016, which is a continuation-in-part of U.S. patent application Ser. No. 15/166,158, titled "SYSTEM FOR AUTOMATED CAPTURE AND ANALYSIS OF BUSINESS INFORMATION FOR SECURITY AND CLIENT-FACING INFRASTRUCTURE RELIABILITY", filed on May 26, 2016, which is a continuation-in-part of U.S. patent application Ser. No. 15/141,752, titled "SYSTEM FOR FULLY INTEGRATED CAPTURE, AND ANALYSIS OF BUSINESS INFORMATION RESULTING IN PREDICTIVE DECISION MAKING AND SIMULATION", filed on Apr. 28, 2016, now issued as U.S. Pat. No. 10,860,962 on Dec. 8, 2020, which is a continuation-in-part of U.S. patent application Ser. No. 15/091,563, titled "SYSTEM FOR CAPTURE, ANALYSIS AND STORAGE OF TIME SERIES DATA FROM SENSORS WITH HETEROGENEOUS REPORT INTERVAL PROFILES", filed on Apr. 5, 2016, now issued as U.S. Pat. No. 10,204,147 on Feb. 12, 2019, and is also a continuation-in-part of U.S. patent application Ser. No. 14/986,536, titled "DISTRIBUTED SYSTEM FOR LARGE VOLUME DEEP WEB DATA EXTRACTION", filed on Dec. 31, 2015, now issued as U.S. Pat. No. 10,210,255 on Feb. 19, 2019, and is also a continuation-in-part of U.S. patent application Ser. No. 14/925,974, titled "RAPID PREDICTIVE ANALYSIS OF VERY LARGE DATA SETS USING THE DISTRIBUTED COMPUTATIONAL GRAPH", filed on Oct. 28, 2015, the entire contents of all of which are incorporated herein by reference in their entirety.

BACKGROUND OF THE INVENTION

Field of the Invention

The disclosure relates to the field of computer management, and more particularly to the field of cybersecurity and threat analytics.

Discussion of the State of the Art

Understanding the cybersecurity profile of an organization is a complex endeavor, and the larger the organization, the more difficult the task is. In fact, the complexity of the task increases exponentially with the size of the organization, not linearly, because each component of the organization's network connects to multiple other components. Existing cybersecurity rating methods (for example, the Common Vulnerability Scoring Systems, or CVSS) fail to adequately profile and rate the cybersecurity profiles of organizations because they do not incorporate sufficient information.

What is needed is a system and method for monitoring critical infrastructure entities within a network, identifying anomalous network events, and correlating those anomalous events to identify attack paths and points of origin.

SUMMARY OF THE INVENTION

Accordingly, the inventor has developed a system and method for correlating network event anomalies using active and passive reconnaissance to identify attack information, that can identify attack patterns and points of origin based on observed anomalies and their relationships to other observed network events. The system and method may further comprise gathering data about the totality of the organization's infrastructure and operations, particularly including the organization's network infrastructure, but also including such information as physical locations, data assets, corporate legal structure, etc., and gathering data about the organization's operations, including such information as business processes and policies, business process dependencies, prior data loss information and security breaches, and behavioral data for both devices and their users. The system and method may further comprise generating a cybersecurity profile

using active and passive internal and external reconnaissance of the organization to determine cybersecurity vulnerabilities and potential impacts to the business in light of the information gathered about the organization's infrastructure and operations. The system and method may further comprise generating cybersecurity scores and ratings that take into account all of the above information, plus the context and purpose of the score or rating to be generated based on the cybersecurity profile.

According to a preferred embodiment, a system for correlating network event anomalies to identify attack information, comprising: a cyber-physical graph module comprising a first plurality of programming instructions stored in a memory of, and operating on a processor of, a computing device, wherein the first plurality of programming instructions, when operating on the processor, cause the computing device to create a cyber-physical graph of the organization using the information, the cyber-physical graph comprising nodes representing entities associated with the organization and edges representing relationships between entities associated with the organization; a reconnaissance engine comprising a second plurality of programming instructions stored in the memory of, and operating on the processor of, the computing device, wherein the second plurality of programming instructions, when operating on the processor, cause the computing device to: perform a reconnaissance search using the cyber-physical graph; and apply some or all of the results of the reconnaissance search to the cyber-physical graph to create a normal behavior model for a plurality of nodes in the cyber-physical graph; and a directed computational graph engine comprising a third plurality of programming instructions stored in the memory of, and operating on the processor of, the computing device, wherein the third plurality of programming instructions, when operating on the processor, cause the computing device to: receive the cybersecurity profile and the normal behavior model; using the cyber-physical graph and the normal behavior model: identify an anomalous event based on analysis of cyber-physical graph and the normal behavior model; analyze the cyber-physical graph and the normal behavior model to identify correlations between affected nodes, the affected nodes being based on the anomalous event; generate a behavior graph based on the identified correlations; analyze the behavior graph to produce a dependency tree comprising causative relationships between events; and traverse the behavior tree backward in a temporal dimension to identify a plurality of starting conditions for the anomalous event, is disclosed.

According to another preferred embodiment, a method for correlating network event anomalies to identify attack information, comprising the steps of: creating a cyber-physical graph of an organization using entities associated with the organization and relationships between the entities associated with the organization, the cyber-physical graph comprising nodes representing the entities associated with the organization and edges representing the relationships between entities associated with the organization; performing a reconnaissance search using the cyber-physical graph; applying some or all of the results of the reconnaissance search to the cyber-physical graph to create a normal behavior model for a plurality of nodes in the cyber-physical graph; and using the cyber-physical graph and the normal behavior model: identifying an anomalous event based on analysis of cyber-physical graph and the normal behavior model; analyzing the cyber-physical graph and the normal behavior model to identify correlations between affected nodes, the affected nodes being based on the anomalous event; generating a behavior graph based on the identified correlations; analyzing the behavior graph to produce a dependency tree comprising causative relationships between events; and traversing the behavior tree backward in a temporal dimension to identify a plurality of starting conditions for the anomalous event, is disclosed.

According to an aspect of an embodiment, a directed computational graph is used, comprising nodes representing data transformations and edges representing communications between the nodes; wherein a plurality of nodes represents a data transformation pipeline; and wherein the directed computational graph is used to direct a workflow in the system.

BRIEF DESCRIPTION OF THE DRAWING FIGURES

The accompanying drawings illustrate several aspects and, together with the description, serve to explain the principles of the invention according to the aspects. It will be appreciated by one skilled in the art that the particular arrangements illustrated in the drawings are merely exemplary, and are not to be considered as limiting of the scope of the invention or the claims herein in any way.

FIG. **1** is a block diagram of an exemplary system architecture for an advanced cyber decision platform for external network reconnaissance and cybersecurity rating.

FIG. **2**A is a block diagram showing general steps for performing passive network reconnaissance.

FIG. **2**B is a process diagram showing a general flow of a process for performing active reconnaissance using DNS leak information collection.

FIG. **2**C is a process diagram showing a general flow of a process for performing active reconnaissance using web application and technology reconnaissance.

FIG. **2**D is a process diagram showing a general flow of a process for producing a cybersecurity rating using reconnaissance data.

FIGS. **3**A and **3**B are process diagrams showing business operating system functions in use to mitigate cyberattacks.

FIG. **4** is a process flow diagram of a method for segmenting cyberattack information to appropriate corporation parties.

FIG. **5** is a diagram of an exemplary architecture for a system for rapid predictive analysis of very large data sets using an actor-driven distributed computational graph, according to one aspect.

FIG. **6** is a diagram of an exemplary architecture for a system for rapid predictive analysis of very large data sets using an actor-driven distributed computational graph, according to one aspect.

FIG. **7** is a diagram of an exemplary architecture for a system for rapid predictive analysis of very large data sets using an actor-driven distributed computational graph, according to one aspect.

FIG. **8** is a flow diagram of an exemplary method for cybersecurity behavioral analytics, according to one aspect.

FIG. **9** is a flow diagram of an exemplary method for measuring the effects of cybersecurity attacks, according to one aspect.

FIG. **10** is a flow diagram of an exemplary method for continuous cybersecurity monitoring and exploration, according to one aspect.

FIG. **11** is a flow diagram of an exemplary method for mapping a cyber-physical system graph, according to one aspect.

FIG. **12** is a flow diagram of an exemplary method for continuous network resilience rating, according to one aspect.

FIG. **13** is a flow diagram of an exemplary method for cybersecurity privilege oversight, according to one aspect.

FIG. **14** is a flow diagram of an exemplary method for cybersecurity risk management, according to one aspect.

FIG. **15** is a flow diagram of an exemplary method for mitigating compromised credential threats, according to one aspect.

FIG. **16** is a flow diagram of an exemplary method for dynamic network and rogue device discovery, according to one aspect.

FIG. **17** is a flow diagram of an exemplary method for Kerberos "golden ticket" attack detection, according to one aspect.

FIG. **18** is a flow diagram of an exemplary method for risk-based vulnerability and patch management, according to one aspect.

FIG. **19** is a block diagram showing an exemplary system architecture for a system for cybersecurity profiling and rating.

FIG. **20** is a relational diagram showing the relationships between exemplary 3rd party search tools, search tasks that can be generated using such tools, and the types of information that may be gathered with those tasks.

FIG. **21** is a relational diagram showing the exemplary types and classifications of information that may be used in constructing a cyber-physical graph of an organization's infrastructure and operations.

FIG. **22** is a directed graph diagram showing an exemplary cyber-physical graph and its possible use in analyzing cybersecurity threats.

FIG. **23** is a block diagram showing exemplary operation of a data to rule mapper.

FIG. **24** is a block diagram showing an exemplary architecture diagram for a scoring engine.

FIG. **25** is a block diagram illustrating an exemplary hardware architecture of a computing device.

FIG. **26** is a block diagram illustrating an exemplary logical architecture for a client device.

FIG. **27** is a block diagram illustrating an exemplary architectural arrangement of clients, servers, and external services.

FIG. **28** is another block diagram illustrating an exemplary hardware architecture of a computing device.

FIG. **29** is a flow diagram illustrating an exemplary method for identifying an anomalous network event.

FIG. **30** is a flow diagram illustrating an exemplary method for identifying correlations between an anomaly and other network events and resources.

FIG. **31** is a flow diagram illustrating an exemplary method for generating a behavior graph based on an identified anomaly and its correlations.

FIG. **32** is a flow diagram illustrating an exemplary method for determining an attack's point of origin using a behavior graph.

DETAILED DESCRIPTION OF THE INVENTION

The inventor has conceived, and reduced to practice, a system and method for correlating network event anomalies using active and passive reconnaissance to identify attack information, that can identify attack patterns and points of origin based on observed anomalies and their relationships to other observed network events. The system and method may further comprise gathering data about the totality of the organization's infrastructure and operations, particularly including the organization's network infrastructure, but also including such information as physical locations, data assets, corporate legal structure, etc., and gathering data about the organization's operations, including such information as business processes and policies, business process dependencies, prior data loss information and security breaches, and behavioral data for both devices and their users. The system and method may further comprise generating a cybersecurity profile using active and passive internal and external reconnaissance of the organization to determine cybersecurity vulnerabilities and potential impacts to the business in light of the information gathered about the organization's infrastructure and operations. The system and method may further comprise generating cybersecurity scores and ratings that take into account all of the above information, plus the context and purpose of the score or rating to be generated based on the cybersecurity profile.

One or more different aspects may be described in the present application. Further, for one or more of the aspects described herein, numerous alternative arrangements may be described; it should be appreciated that these are presented for illustrative purposes only and are not limiting of the aspects contained herein or the claims presented herein in any way. One or more of the arrangements may be widely applicable to numerous aspects, as may be readily apparent from the disclosure. In general, arrangements are described in sufficient detail to enable those skilled in the art to practice one or more of the aspects, and it should be appreciated that other arrangements may be utilized and that structural, logical, software, electrical and other changes may be made without departing from the scope of the particular aspects. Particular features of one or more of the aspects described herein may be described with reference to one or more particular aspects or figures that form a part of the present disclosure, and in which are shown, by way of illustration, specific arrangements of one or more of the aspects. It should be appreciated, however, that such features are not limited to usage in the one or more particular aspects or figures with reference to which they are described. The present disclosure is neither a literal description of all arrangements of one or more of the aspects nor a listing of features of one or more of the aspects that must be present in all arrangements.

Headings of sections provided in this patent application and the title of this patent application are for convenience only, and are not to be taken as limiting the disclosure in any way.

Devices that are in communication with each other need not be in continuous communication with each other, unless expressly specified otherwise. In addition, devices that are in communication with each other may communicate directly or indirectly through one or more communication means or intermediaries, logical or physical.

A description of an aspect with several components in communication with each other does not imply that all such components are required. To the contrary, a variety of optional components may be described to illustrate a wide variety of possible aspects and in order to more fully illustrate one or more aspects. Similarly, although process steps, method steps, algorithms or the like may be described in a sequential order, such processes, methods and algorithms may generally be configured to work in alternate orders, unless specifically stated to the contrary. In other words, any sequence or order of steps that may be described in this patent application does not, in and of itself, indicate

a requirement that the steps be performed in that order. The steps of described processes may be performed in any order practical. Further, some steps may be performed simultaneously despite being described or implied as occurring non-simultaneously (e.g., because one step is described after the other step). Moreover, the illustration of a process by its depiction in a drawing does not imply that the illustrated process is exclusive of other variations and modifications thereto, does not imply that the illustrated process or any of its steps are necessary to one or more of the aspects, and does not imply that the illustrated process is preferred. Also, steps are generally described once per aspect, but this does not mean they must occur once, or that they may only occur once each time a process, method, or algorithm is carried out or executed. Some steps may be omitted in some aspects or some occurrences, or some steps may be executed more than once in a given aspect or occurrence.

When a single device or article is described herein, it will be readily apparent that more than one device or article may be used in place of a single device or article. Similarly, where more than one device or article is described herein, it will be readily apparent that a single device or article may be used in place of the more than one device or article.

The functionality or the features of a device may be alternatively embodied by one or more other devices that are not explicitly described as having such functionality or features. Thus, other aspects need not include the device itself.

Techniques and mechanisms described or referenced herein will sometimes be described in singular form for clarity. However, it should be appreciated that particular aspects may include multiple iterations of a technique or multiple instantiations of a mechanism unless noted otherwise. Process descriptions or blocks in figures should be understood as representing modules, segments, or portions of code which include one or more executable instructions for implementing specific logical functions or steps in the process. Alternate implementations are included within the scope of various aspects in which, for example, functions may be executed out of order from that shown or discussed, including substantially concurrently or in reverse order, depending on the functionality involved, as would be understood by those having ordinary skill in the art.

Definitions

As used herein, "graph" is a representation of information and relationships, where each primary unit of information makes up a "node" or "vertex" of the graph and the relationship between two nodes makes up an edge of the graph. Nodes can be further qualified by the connection of one or more descriptors or "properties" to that node. For example, given the node "James R," name information for a person, qualifying properties might be "183 cm tall", "DOB 8/13/1965" and "speaks English". Similar to the use of properties to further describe the information in a node, a relationship between two nodes that forms an edge can be qualified using a "label". Thus, given a second node "Thomas G," an edge between "James R" and "Thomas G" that indicates that the two people know each other might be labeled "knows." When graph theory notation (Graph=(Vertices, Edges)) is applied this situation, the set of nodes are used as one parameter of the ordered pair, V and the set of 2 element edge endpoints are used as the second parameter of the ordered pair, E. When the order of the edge endpoints within the pairs of E is not significant, for example, the edge James R, Thomas G is equivalent to Thomas G, James R, the graph is designated as "undirected." Under circumstances when a relationship flows from one node to another in one direction, for example James R is "taller" than Thomas G, the order of the endpoints is significant. Graphs with such edges are designated as "directed." In the distributed computational graph system, transformations within transformation pipeline are represented as directed graph with each transformation comprising a node and the output messages between transformations comprising edges. Distributed computational graph stipulates the potential use of non-linear transformation pipelines which are programmatically linearized. Such linearization can result in exponential growth of resource consumption. The most sensible approach to overcome possibility is to introduce new transformation pipelines just as they are needed, creating only those that are ready to compute. Such method results in transformation graphs which are highly variable in size and node, edge composition as the system processes data streams. Those familiar with the art will realize that transformation graph may assume many shapes and sizes with a vast topography of edge relationships. The examples given were chosen for illustrative purposes only and represent a small number of the simplest of possibilities. These examples should not be taken to define the possible graphs expected as part of operation of the invention

As used herein, "transformation" is a function performed on zero or more streams of input data which results in a single stream of output which may or may not then be used as input for another transformation. Transformations may comprise any combination of machine, human or machine-human interactions Transformations need not change data that enters them, one example of this type of transformation would be a storage transformation which would receive input and then act as a queue for that data for subsequent transformations. As implied above, a specific transformation may generate output data in the absence of input data. A time stamp serves as a example. In the invention, transformations are placed into pipelines such that the output of one transformation may serve as an input for another. These pipelines can consist of two or more transformations with the number of transformations limited only by the resources of the system. Historically, transformation pipelines have been linear with each transformation in the pipeline receiving input from one antecedent and providing output to one subsequent with no branching or iteration. Other pipeline configurations are possible. The invention is designed to permit several of these configurations including, but not limited to: linear, afferent branch, efferent branch and cyclical.

A "database" or "data storage subsystem" (these terms may be considered substantially synonymous), as used herein, is a system adapted for the long-term storage, indexing, and retrieval of data, the retrieval typically being via some sort of querying interface or language. "Database" may be used to refer to relational database management systems known in the art, but should not be considered to be limited to such systems. Many alternative database or data storage system technologies have been, and indeed are being, introduced in the art, including but not limited to distributed non-relational data storage systems such as Hadoop, column-oriented databases, in-memory databases, and the like. While various aspects may preferentially employ one or another of the various data storage subsystems available in the art (or available in the future), the invention should not be construed to be so limited, as any data storage architecture may be used according to the aspects. Similarly, while in some cases one or more particular data storage needs are described as being satisfied by separate components (for example, an expanded private

capital markets database and a configuration database), these descriptions refer to functional uses of data storage systems and do not refer to their physical architecture. For instance, any group of data storage systems of databases referred to herein may be included together in a single database management system operating on a single machine, or they may be included in a single database management system operating on a cluster of machines as is known in the art. Similarly, any single database (such as an expanded private capital markets database) may be implemented on a single machine, on a set of machines using clustering technology, on several machines connected by one or more messaging systems known in the art, or in a master/slave arrangement common in the art. These examples should make clear that no particular architectural approaches to database management is preferred according to the invention, and choice of data storage technology is at the discretion of each implementer, without departing from the scope of the invention as claimed.

A "data context", as used herein, refers to a set of arguments identifying the location of data. This could be a Rabbit queue, a .csv file in cloud-based storage, or any other such location reference except a single event or record. Activities may pass either events or data contexts to each other for processing. The nature of a pipeline allows for direct information passing between activities, and data locations or files do not need to be predetermined at pipeline start.

A "pipeline", as used herein and interchangeably referred to as a "data pipeline" or a "processing pipeline", refers to a set of data streaming activities and batch activities. Streaming and batch activities can be connected indiscriminately within a pipeline. Events will flow through the streaming activity actors in a reactive way. At the junction of a streaming activity to batch activity, there will exist a StreamBatchProtocol data object. This object is responsible for determining when and if the batch process is run. One or more of three possibilities can be used for processing triggers: regular timing interval, every N events, or optionally an external trigger. The events are held in a queue or similar until processing. Each batch activity may contain a "source" data context (this may be a streaming context if the upstream activities are streaming), and a "destination" data context (which is passed to the next activity). Streaming activities may have an optional "destination" streaming data context (optional meaning: caching/persistence of events vs. ephemeral), though this should not be part of the initial implementation.

Conceptual Architecture

FIG. **1** is a block diagram of an advanced cyber decision platform for external network reconnaissance and cybersecurity rating. Client access to the system **105** for specific data entry, system control and for interaction with system output such as automated predictive decision making and planning and alternate pathway simulations, occurs through the system's distributed, extensible high bandwidth cloud interface **110** which uses a versatile, robust web application driven interface for both input and display of client-facing information via network **107** and operates a data store **112** such as, but not limited to MONGODB™, COUCHDB™, CASSANDRA™ or REDIS™ according to various arrangements. Much of the business data analyzed by the system both from sources within the confines of the client business, and from cloud based sources, also enter the system through the cloud interface **110**, data being passed to the connector module **135** which may possess the API routines **135***a* needed to accept and convert the external data and then pass the normalized information to other analysis and transformation components of the system, the directed computational graph module **155**, high volume web crawler module **115**, multidimensional time series database (MDTSDB) **120** and the graph stack service **145**. The directed computational graph module **155** retrieves one or more streams of data from a plurality of sources, which includes, but is in no way not limited to, a plurality of physical sensors, network service providers, web based questionnaires and surveys, monitoring of electronic infrastructure, crowd sourcing campaigns, and human input device information. Within the directed computational graph module **155**, data may be split into two identical streams in a specialized pre-programmed data pipeline **155***a*, wherein one sub-stream may be sent for batch processing and storage while the other sub-stream may be reformatted for transformation pipeline analysis. The data is then transferred to the general transformer service module **160** for linear data transformation as part of analysis or the decomposable transformer service module **150** for branching or iterative transformations that are part of analysis. The directed computational graph module **155** represents all data as directed graphs where the transformations are nodes and the result messages between transformations edges of the graph. The high volume web crawling module **115** uses multiple server hosted preprogrammed web spiders, which while autonomously configured are deployed within a web scraping framework **115***a* of which SCRAPY™ is an example, to identify and retrieve data of interest from web based sources that are not well tagged by conventional web crawling technology. The multiple dimension time series data store module **120** may receive streaming data from a large plurality of sensors that may be of several different types. The multiple dimension time series data store module may also store any time series data encountered by the system such as but not limited to enterprise network usage data, component and system logs, performance data, network service information captures such as, but not limited to news and financial feeds, and sales and service related customer data. The module is designed to accommodate irregular and high volume surges by dynamically allotting network bandwidth and server processing channels to process the incoming data. Inclusion of programming wrappers **120***a* for languages examples of which are, but not limited to C++, PERL, PYTHON, and ERLANG™ allows sophisticated programming logic to be added to the default function of the multidimensional time series database **120** without intimate knowledge of the core programming, greatly extending breadth of function. Data retrieved by the multidimensional time series database (MDTSDB) **120** and the high volume web crawling module **115** may be further analyzed and transformed into task optimized results by the directed computational graph **155** and associated general transformer service **150** and decomposable transformer service **160** modules. Alternately, data from the multidimensional time series database and high volume web crawling modules may be sent, often with scripted cuing information determining important vertexes **145***a*, to the graph stack service module **145** which, employing standardized protocols for converting streams of information into graph representations of that data, for example, open graph internet technology although the invention is not reliant on any one standard. Through the steps, the graph stack service module **145** represents data in graphical form influenced by any pre-determined scripted modifications **145***a* and stores it in a graph-based data store **145***b* such as GIRAPH™ or a key

value pair type data store REDIS™, or RIAK™, among others, all of which are suitable for storing graph-based information.

Results of the transformative analysis process may then be combined with further client directives, additional business rules and practices relevant to the analysis and situational information external to the already available data in the automated planning service module **130** which also runs powerful information theory **130***a* based predictive statistics functions and machine learning algorithms to allow future trends and outcomes to be rapidly forecast based upon the current system derived results and choosing each a plurality of possible business decisions. The using all available data, the automated planning service module **130** may propose business decisions most likely to result is the most favorable business outcome with a usably high level of certainty. Closely related to the automated planning service module in the use of system derived results in conjunction with possible externally supplied additional information in the assistance of end user business decision making, the action outcome simulation module **125** with its discrete event simulator programming module **125***a* coupled with the end user facing observation and state estimation service **140** which is highly scriptable **140***b* as circumstances require and has a game engine **140***a* to more realistically stage possible outcomes of business decisions under consideration, allows business decision makers to investigate the probable outcomes of choosing one pending course of action over another based upon analysis of the current available data.

When performing external reconnaissance via a network **107**, web crawler **115** may be used to perform a variety of port and service scanning operations on a plurality of hosts. This may be used to target individual network hosts (for example, to examine a specific server or client device) or to broadly scan any number of hosts (such as all hosts within a particular domain, or any number of hosts up to the complete IPv4 address space). Port scanning is primarily used for gathering information about hosts and services connected to a network, using probe messages sent to hosts that prompt a response from that host. Port scanning is generally centered around the transmission control protocol (TCP), and using the information provided in a prompted response a port scan can provide information about network and application layers on the targeted host.

Port scan results can yield information on open, closed, or undetermined ports on a target host. An open port indicated that an application or service is accepting connections on this port (such as ports used for receiving customer web traffic on a web server), and these ports generally disclose the greatest quantity of useful information about the host. A closed port indicates that no application or service is listening for connections on that port, and still provides information about the host such as revealing the operating system of the host, which may discovered by fingerprinting the TCP/IP stack in a response. Different operating systems exhibit identifiable behaviors when populating TCP fields, and collecting multiple responses and matching the fields against a database of known fingerprints makes it possible to determine the OS of the host even when no ports are open. An undetermined port is one that does not produce a requested response, generally because the port is being filtered by a firewall on the host or between the host and the network (for example, a corporate firewall behind which all internal servers operate).

Scanning may be defined by scope to limit the scan according to two dimensions, hosts and ports. A horizontal scan checks the same port on multiple hosts, often used by attackers to check for an open port on any available hosts to select a target for an attack that exploits a vulnerability using that port. This type of scan is also useful for security audits, to ensure that vulnerabilities are not exposed on any of the target hosts. A vertical scan defines multiple ports to examine on a single host, for example a "vanilla scan" which targets every port of a single host, or a "strobe scan" that targets a small subset of ports on the host. This type of scan is usually performed for vulnerability detection on single systems, and due to the single-host nature is impractical for large network scans. A block scan combines elements of both horizontal and vertical scanning, to scan multiple ports on multiple hosts. This type of scan is useful for a variety of service discovery and data collection tasks, as it allows a broad scan of many hosts (up to the entire Internet, using the complete IPv4 address space) for a number of desired ports in a single sweep.

Large port scans involve quantitative research, and as such may be treated as experimental scientific measurement and are subject to measurement and quality standards to ensure the usefulness of results. To avoid observational errors during measurement, results must be precise (describing a degree of relative proximity between individual measured values), accurate (describing relative proximity of measured values to a reference value), preserve any metadata that accompanies the measured data, avoid misinterpretation of data due to faulty measurement execution, and must be well-calibrated to efficiently expose and address issues of inaccuracy or misinterpretation. In addition to these basic requirements, large volumes of data may lead to unexpected behavior of analysis tools, and extracting a subset to perform initial analysis may help to provide an initial overview before working with the complete data set. Analysis should also be reproducible, as with all experimental science, and should incorporate publicly-available data to add value to the comprehensibility of the research as well as contributing to a "common framework" that may be used to confirm results.

When performing a port scan, web crawler **115** may employ a variety of software suitable for the task, such as Nmap, ZMap, or masscan. Nmap is suitable for large scans as well as scanning individual hosts, and excels in offering a variety of diverse scanning techniques. ZMap is a newer application and unlike Nmap (which is more general-purpose), ZMap is designed specifically with Internet-wide scans as the intent. As a result, ZMap is far less customizable and relies on horizontal port scans for functionality, achieving fast scan times using techniques of probe randomization (randomizing the order in which probes are sent to hosts, minimizing network saturation) and asynchronous design (utilizing stateless operation to send and receive packets in separate processing threads). Masscan uses the same asynchronous operation model of ZMap, as well as probe randomization. In masscan however, a certain degree of statistical randomness is sacrificed to improve computation time for large scans (such as when scanning the entire IPv4 address space), using the BlackRock algorithm. This is a modified implementation of symmetric encryption algorithm DES, with fewer rounds and modulo operations in place of binary ones to allow for arbitrary ranges and achieve faster computation time for large data sets.

Received scan responses may be collected and processed through a plurality of data pipelines **155***a* to analyze the collected information. MDTSDB **120** and graph stack **145** may be used to produce a hybrid graph/time-series database using the analyzed data, forming a graph of Internet-accessible organization resources and their evolving state infor-

mation over time. Customer-specific profiling and scanning information may be linked to CPG graphs (as described below in detail, referring to FIG. **11**) for a particular customer, but this information may be further linked to the base-level graph of internet-accessible resources and information. Depending on customer authorizations and legal or regulatory restrictions and authorizations, techniques used may involve both passive, semi-passive and active scanning and reconnaissance.

FIG. **2**A is a block diagram showing general steps **200** for performing passive network reconnaissance. It should be appreciated that the steps illustrated and described may be performed in any order, and that steps may be added or omitted as needed for any particular reconnaissance operation. In a step **201**, network address ranges and domains or sub-domains associated with a plurality of targets may be identified, for example to collect information for defining the scope of further scanning operations. In another step **202**, external sites may be identified to understand relationships between targets and other third-party content providers, such as trust relationships or authoritative domain name service (DNS) resolution records. In another step **203**, individual people or groups may be identified using names, email addresses, phone numbers, or other identifying information that may be useful for a variety of social engineering activities. In another step **204**, technologies used may be identified, such as types or versions of hardware or software used by an organization, and this may include collecting and extracting information from job descriptions (for example) to identify technologies in use by an organization (for example, a job description for an administrator familiar with specific database software indicates that said software is in use within the organization). In another step **205**, content of interest may be identified, for example including web and email portals, log files, backup or archive files, and other forms of sensitive information that may be contained within HTML comments or client-side scripts, as may be useful for vulnerability discovery and penetration testing activities. In another step **206**, publicly-available information may be used to identify vulnerabilities that may be exploited with further active penetration testing.

FIG. **2**B is a process diagram showing a general flow of a process **210** for performing active reconnaissance using DNS leak information collection. In an initial step **211**, publicly-available DNS leak disclosure information may be collected to maintain current information regarding known leaks and vulnerabilities. In a next step **212**, third-level domain (TLDR) information may be collected and used to report domain risk factors, such as domains that do not resolve properly (due to malformed DNS records, for example). In a next step **213**, a DNS trust map may be created using a hybrid graph/time-series data structure, using a graph stack service **145** and MDTSDB **120**. This trust map may be produced as the output of an extraction process performed by a DCG **155** through a plurality of data pipelines **155***a*, analyzing collected data and mapping data points to produce hybrid structured output representing each data point over time. In a final step **214**, the trust map may then be analyzed to identify anomalies, for example using community detection algorithms that may discover when new references are being created, and this may be used to identify vulnerabilities that may arise as a byproduct of the referential nature of a DNS hierarchy. In this manner, DCG pipeline processing and time-series data graphing may be used to identify vulnerabilities that would otherwise be obscured within a large dataset.

FIG. **2**C is a process diagram showing a general flow of a process **220** for performing active reconnaissance using web application and technology reconnaissance. In an initial step **221**, a plurality of manual HTTP requests may be transmitted to a host, for example to determine if a web server is announcing itself, or to obtain an application version number from an HTTP response message. In a next step **222**, a robots.txt, used to identify and communicate with web crawlers and other automated "bots", may be searched for to identify portions of an application or site that robots are requested to ignore. In a next step **223**, the host application layer may be fingerprinted, for example using file extensions and response message fields to identify characteristic patterns or markers that may be used to identify host or application details. In a next step **224**, publicly-exposed/admin pages may be checked, to determine if any administrative portals are exposed and therefore potentially-vulnerable, as well as to potentially determine administration policies or capabilities based on exposed information. In a final step **225**, an application may be profiled according to a particular toolset in use, such as WORDPRESS™ (for example) or other specific tools or plugins.

FIG. **2**D is a process diagram showing a general flow of a process **230** for producing a cybersecurity rating using reconnaissance data. In an initial step **231**, external reconnaissance may be performed using DNS and IP information as described above (referring to FIG. **2**B), collecting information from DNS records, leak announcements, and publicly-available records to produce a DNS trust map from collected information and the DCG-driven analysis thereof. In a next step **232**, web and application recon may be performed (as described in FIG. **2**C), collecting information on applications, sites, and publicly-available records. In a next step **233**, collected information over time may be analyzed for software version numbers, revealing the patching frequency of target hosts and their respective applications and services. Using a hybrid time-series graph, timestamps may be associated with ongoing changes to reveal these updates over time. In a next step **234**, a plurality of additional endpoints may be scanned, such as (for example, including but not limited to) internet-of-things (IoT) devices that may be scanned and fingerprinted, end-user devices such as personal smartphones, tablets, or computers, or social network endpoints such as scraping content from user social media pages or feeds. User devices may be fingerprinted and analyzed similar to organization hosts, and social media content may be retrieved such as collecting sentiment from services like TWITTER™ or LINKEDIN™, or analyzing job description listings and other publicly-available information. In a next step **235**, open-source intelligence feeds may be checked, such as company IP address blacklists, search domains, or information leaks (for example, posted to public records such as PASTEBIN™). In a final step **236**, collected information from all sources may be scored according to a weighted system, producing an overall cybersecurity rating score based on the information collected and the analysis of that information to reveal additional insights, relationships, and vulnerabilities.

For example, in an exemplary scoring system similar to a credit rating, information from initial Internet recon operations may be assigned a score up to 400 points, along with up to 200 additional points for web/application recon results, 100 points for patch frequency, and 50 points each for additional endpoints and open-source intel results. This yields a weighted score incorporating all available information from all scanned sources, allowing a meaningful and readily-appreciable representation of an organization's over-

all cybersecurity strength. Additionally, as scanning may be performed repeatedly and results collected into a time-series hybrid data structure, this cybersecurity rating may evolve over time to continuously reflect the current state of the organization, reflecting any recent changes, newly-discovered or announced vulnerabilities, software or hardware updates, newly-added or removed devices or services, and any other changes that may occur.

FIGS. **3**A and **3**B are process diagrams showing business operating system functions in use to mitigate cyberattacks. Input network data which may include network flow patterns **321**, the origin and destination of each piece of measurable network traffic **322**, system logs from servers and workstations on the network **323**, endpoint data **329**, any security event log data from servers or available security information and event (SIEM) systems **324**, external threat intelligence feeds **324**, identity or assessment context **325**, external network health or cybersecurity feeds **326**, Kerberos domain controller or ACTIVE DIRECTORY™ server logs or instrumentation **327**, business unit performance related data **328**, endpoint data **329**, among many other possible data types for which the invention was designed to analyze and integrate, may pass into **315** the business operating system **310** for analysis as part of its cyber security function. These multiple types of data from a plurality of sources may be transformed for analysis **311**, **312** using at least one of the specialized cybersecurity, risk assessment or common functions of the business operating system in the role of cybersecurity system, such as, but not limited to network and system user privilege oversight **331**, network and system user behavior analytics **332**, attacker and defender action timeline **333**, SIEM integration and analysis **334**, dynamic benchmarking **335**, and incident identification and resolution performance analytics **336** among other possible cybersecurity functions; value at risk (VAR) modeling and simulation **341**, anticipatory vs. reactive cost estimations of different types of data breaches to establish priorities **342**, work factor analysis **343** and cyber event discovery rate **344** as part of the system's risk analytics capabilities; and the ability to format and deliver customized reports and dashboards **351**, perform generalized, ad hoc data analytics on demand **352**, continuously monitor, process and explore incoming data for subtle changes or diffuse informational threads **353** and generate cyber-physical systems graphing **354** as part of the business operating system's common capabilities. Output **317** can be used to configure network gateway security appliances **361**, to assist in preventing network intrusion through predictive change to infrastructure recommendations **362**, to alert an enterprise of ongoing cyberattack early in the attack cycle, possibly thwarting it but at least mitigating the damage **362**, to record compliance to standardized guidelines or SLA requirements **363**, to continuously probe existing network infrastructure and issue alerts to any changes which may make a breach more likely **364**, suggest solutions to any domain controller ticketing weaknesses detected **365**, detect presence of malware **366**, perform one time or continuous vulnerability scanning depending on client directives **367**, and thwart or mitigate damage from cyber attacks **368**. These examples are, of course, only a subset of the possible uses of the system, they are exemplary in nature and do not reflect any boundaries in the capabilities of the invention.

FIG. **4** is a process flow diagram of a method for segmenting cyberattack information to appropriate corporation parties **400**. As previously disclosed **200**, **351**, one of the strengths of the advanced cyber-decision platform is the ability to finely customize reports and dashboards to specific audiences, concurrently is appropriate. This customization is possible due to the devotion of a portion of the business operating system's programming specifically to outcome presentation by modules which include the observation and state estimation service **140** with its game engine **140***a* and script interpreter **140***b*. In the setting of cybersecurity, issuance of specialized alerts, updates and reports may significantly assist in getting the correct mitigating actions done in the most timely fashion while keeping all participants informed at predesignated, appropriate granularity. Upon the detection of a cyberattack by the system **401** all available information about the ongoing attack and existing cybersecurity knowledge are analyzed, including through predictive simulation in near real time **402** to develop both the most accurate appraisal of current events and actionable recommendations concerning where the attack may progress and how it may be mitigated. The information generated in totality is often more than any one group needs to perform their mitigation tasks. At this point, during a cyberattack, providing a single expansive and all inclusive alert, dashboard image, or report may make identification and action upon the crucial information by each participant more difficult, therefore the cybersecurity focused arrangement may create multiple targeted information streams each concurrently designed to produce most rapid and efficacious action throughout the enterprise during the attack and issue follow-up reports with and recommendations or information that may lead to long term changes afterward **403**. Examples of groups that may receive specialized information streams include but may not be limited to front line responders during the attack **404**, incident forensics support both during and after the attack **405**, chief information security officer **406** and chief risk officer **407** the information sent to the latter two focused to appraise overall damage and to implement both mitigating strategy and preventive changes after the attack. Front line responders may use the cyber-decision platform's analyzed, transformed and correlated information specifically sent to them **404** to probe the extent of the attack, isolate such things as: the predictive attacker's entry point onto the enterprise's network, the systems involved or the predictive ultimate targets of the attack and may use the simulation capabilities of the system to investigate alternate methods of successfully ending the attack and repelling the attackers in the most efficient manner, although many other queries known to those skilled in the art are also answerable by the invention. Simulations run may also include the predictive effects of any attack mitigating actions on normal and critical operation of the enterprise's IT systems and corporate users. Similarly, a chief information security officer may use the cyber-decision platform to predictively analyze **406** what corporate information has already been compromised, predictively simulate the ultimate information targets of the attack that may or may not have been compromised and the total impact of the attack what can be done now and in the near future to safeguard that information. Further, during retrospective forensic inspection of the attack, the forensic responder may use the cyber-decision platform **405***a* to clearly and completely map the extent of network infrastructure through predictive simulation and large volume data analysis. The forensic analyst may also use the platform's capabilities to perform a time series and infrastructural spatial analysis of the attack's progression with methods used to infiltrate the enterprise's subnets and servers. Again, the chief risk officer would perform analyses of what information **407***a* was stolen and predictive simulations on what the theft means to the enterprise as time progresses. Additionally, the system's predictive capabilities may be employed to assist in creation of a plan for changes

of the IT infrastructural that should be made that are optimal for remediation of cybersecurity risk under possibly limited enterprise budgetary constraints in place at the company so as to maximize financial outcome.

FIG. **5** is a diagram of an exemplary architecture for a system for rapid predictive analysis of very large data sets using an actor-driven distributed computational graph **500**, according to one aspect. According to the aspect, a DCG **500** may comprise a pipeline orchestrator **501** that may be used to perform a variety of data transformation functions on data within a processing pipeline, and may be used with a messaging system **510** that enables communication with any number of various services and protocols, relaying messages and translating them as needed into protocol-specific API system calls for interoperability with external systems (rather than requiring a particular protocol or service to be integrated into a DCG **500**).

Pipeline orchestrator **501** may spawn a plurality of child pipeline clusters **502***a-b*, which may be used as dedicated workers for streamlining parallel processing. In some arrangements, an entire data processing pipeline may be passed to a child cluster **502***a* for handling, rather than individual processing tasks, enabling each child cluster **502***a-b* to handle an entire data pipeline in a dedicated fashion to maintain isolated processing of different pipelines using different cluster nodes **502***a-b*. Pipeline orchestrator **501** may provide a software API for starting, stopping, submitting, or saving pipelines. When a pipeline is started, pipeline orchestrator **501** may send the pipeline information to an available worker node **502***a-b*, for example using AKKA™ clustering. For each pipeline initialized by pipeline orchestrator **501**, a reporting object with status information may be maintained. Streaming activities may report the last time an event was processed, and the number of events processed. Batch activities may report status messages as they occur. Pipeline orchestrator **501** may perform batch caching using, for example, an IGFS™ caching filesystem. This allows activities **512***a-d* within a pipeline **502***a-b* to pass data contexts to one another, with any necessary parameter configurations.

A pipeline manager **511***a-b* may be spawned for every new running pipeline, and may be used to send activity, status, lifecycle, and event count information to the pipeline orchestrator **501**. Within a particular pipeline, a plurality of activity actors **512***a-d* may be created by a pipeline manager **511***a-b* to handle individual tasks, and provide output to data services **522***a-d*. Data models used in a given pipeline may be determined by the specific pipeline and activities, as directed by a pipeline manager **511***a-b*. Each pipeline manager **511***a-b* controls and directs the operation of any activity actors **512***a-d* spawned by it. A pipeline process may need to coordinate streaming data between tasks. For this, a pipeline manager **511***a-b* may spawn service connectors to dynamically create TCP connections between activity instances **512***a-d*. Data contexts may be maintained for each individual activity **512***a-d*, and may be cached for provision to other activities **512***a-d* as needed. A data context defines how an activity accesses information, and an activity **512***a-d* may process data or simply forward it to a next step. Forwarding data between pipeline steps may route data through a streaming context or batch context.

A client service cluster **530** may operate a plurality of service actors **521***a-d* to serve the requests of activity actors **512***a-d*, ideally maintaining enough service actors **521***a-d* to support each activity per the service type. These may also be arranged within service clusters **520***a-d*, in a manner similar to the logical organization of activity actors **512***a-d* within clusters **502***a-b* in a data pipeline. A logging service **530** may be used to log and sample DCG requests and messages during operation while notification service **540** may be used to receive alerts and other notifications during operation (for example to alert on errors, which may then be diagnosed by reviewing records from logging service **530**), and by being connected externally to messaging system **510**, logging and notification services can be added, removed, or modified during operation without impacting DCG **500**. A plurality of DCG protocols **550***a-b* may be used to provide structured messaging between a DCG **500** and messaging system **510**, or to enable messaging system **510** to distribute DCG messages across service clusters **520***a-d* as shown. A service protocol **560** may be used to define service interactions so that a DCG **500** may be modified without impacting service implementations. In this manner it can be appreciated that the overall structure of a system using an actor-driven DCG **500** operates in a modular fashion, enabling modification and substitution of various components without impacting other operations or requiring additional reconfiguration.

FIG. **6** is a diagram of an exemplary architecture for a system for rapid predictive analysis of very large data sets using an actor-driven distributed computational graph **500**, according to one aspect. According to the aspect, a variant messaging arrangement may utilize messaging system **510** as a messaging broker using a streaming protocol **610**, transmitting and receiving messages immediately using messaging system **510** as a message broker to bridge communication between service actors **521***a-b* as needed. Alternately, individual services **522***a-b* may communicate directly in a batch context **620**, using a data context service **630** as a broker to batch-process and relay messages between services **522***a-b*.

FIG. **7** is a diagram of an exemplary architecture for a system for rapid predictive analysis of very large data sets using an actor-driven distributed computational graph **500**, according to one aspect. According to the aspect, a variant messaging arrangement may utilize a service connector **710** as a central message broker between a plurality of service actors **521***a-b*, bridging messages in a streaming context **610** while a data context service **630** continues to provide direct peer-to-peer messaging between individual services **522***a-b* in a batch context **620**.

It should be appreciated that various combinations and arrangements of the system variants described above (referring to FIGS. **1-7**) may be possible, for example using one particular messaging arrangement for one data pipeline directed by a pipeline manager **511***a-b*, while another pipeline may utilize a different messaging arrangement (or may not utilize messaging at all). In this manner, a single DCG **500** and pipeline orchestrator **501** may operate individual pipelines in the manner that is most suited to their particular needs, with dynamic arrangements being made possible through design modularity as described above in FIG. **5**.

FIG. **19** is block diagram showing an exemplary system architecture **1900** for a system for cybersecurity profiling and rating. The system in this example contains a cyber-physical graph **1902** which is used to represent a complete picture of an organization's infrastructure and operations including, importantly, the organization's computer network infrastructure particularly around system configurations that influence cybersecurity protections and resiliency. The system further contains a directed computational graph **1911**, which contains representations of complex processing pipelines and is used to control workflows through the system such as determining which 3rd party search tools **1915** to use, assigning search tasks, and analyzing the cyber-physical

graph **1902** and comparing results of the analysis against reconnaissance data received from the reconnaissance engine **1906** and stored in the reconnaissance data storage **1905**. In some embodiments, the determination of which 3rd party search tools **1915** to use and assignment of search tasks may be implemented by a reconnaissance engine **1906**. The cyber-physical graph **1902** plus the analyses of data directed by the directed computational graph on the reconnaissance data received from the reconnaissance engine **1906** are combined to represent the cyber-security profile of the client organization whose network **1907** is being evaluated. A queuing system **1912** is used to organize and schedule the search tasks requested by the reconnaissance engine **1906**. A data to rule mapper **1904** is used to retrieve laws, policies, and other rules from an authority database **1903** and compare reconnaissance data received from the reconnaissance engine **1906** and stored in the reconnaissance data storage **1905** against the rules in order to determine whether and to what extent the data received indicates a violation of the rules. Machine learning models **1901** may be used to identify patterns and trends in any aspect of the system, but in this case are being used to identify patterns and trends in the data which would help the data to rule mapper **1904** determine whether and to what extent certain data indicate a violation of certain rules. A scoring engine **1910** receives the data analyses performed by the directed computational graph **1911**, the output of the data to rule mapper **1904**, plus event and loss data **1914** and contextual data **1909** which defines a context in which the other data are to be scored and/or rated. A public-facing proxy network **1908** is established outside of a firewall **1917** around the client network **1907** both to control access to the client network from the Internet **1913**, and to provide the ability to change the outward presentation of the client network **1907** to the Internet **1913**, which may affect the data obtained by the reconnaissance engine **1906**. In some embodiments, certain components of the system may operate outside the client network **1907** and may access the client network through a secure, encrypted virtual private network (VPN) **1916**, as in a cloud-based or platform-as-a-service implementation, but in other embodiments some or all of these components may be installed and operated from within the client network **1907**.

As a brief overview of operation, information is obtained about the client network **1907** and the client organization's operations, which is used to construct a cyber-physical graph **1902** representing the relationships between devices, users, resources, and processes in the organization, and contextualizing cybersecurity information with physical and logical relationships that represent the flow of data and access to data within the organization including, in particular, network security protocols and procedures. The directed computational graph **1911** containing workflows and analysis processes, selects one or more analyses to be performed on the cyber-physical graph **1902**. Some analyses may be performed on the information contained in the cyber-physical graph, and some analyses may be performed on or against the cyber-physical graph using information obtained from the Internet **1913** from reconnaissance engine **1906**. The workflows contained in the directed computational graph **1911** select one or more search tools to obtain information about the organization from the Internet **1915**, and may comprise one or more third party search tools **1915** available on the Internet. As data are collected, they are fed into a reconnaissance data storage **1905**, from which they may be retrieved and further analyzed. Comparisons are made between the data obtained from the reconnaissance engine **1906**, the cyber-physical graph **1902**, the data to rule mapper, from which comparisons a cybersecurity profile of the organization is developed. The cybersecurity profile is sent to the scoring engine **1910** along with event and loss data **1914** and context data **1909** for the scoring engine **1910** to develop a score and/or rating for the organization that takes into consideration both the cybersecurity profile, context, and other information.

FIG. **24** is a block diagram showing an exemplary architecture **2400** for a scoring engine. Data fed into the scoring engine comprise the cybersecurity profile **1918** and reconnaissance data **1905** developed at earlier stages of system operation. Based on these data, a frequency and severity of attack is estimated **2408**. For each risk type, curve fitting **2402** may be performed on the data points to assign a "best fit" function along the range of data points, which captures trends in the data and allows for predictions of how similar data will behave in the future. Aggregations of operational variables **2403** may be applied to identify maxima, minima, counts, sums, and standard deviations of the data. Risk identification and quantification is then performed **2413**, and a business impact analysis is performed **2412** based on a totality of the predicted risks, their severity, business dependencies reflected in the cyber-physical graph, and prior event and loss data **2410**, among other variables. From this analysis of business impact **2412**, a network resilience rating is assigned **2405**, representing a weighted and adjusted total of relative exposure the organization has to various types of risks, each of which may be assigned a sub-rating. The network resilience rating **2405** may be a single score for all factors, a combination of scores, or a score for a particular risk or area of concern. The network resilience rating **2411** may then be adjusted or filtered depending on the context in which it is to be used **2409**. For example, context data received **2408** may indicate that the scores are to be used for compliance with internal theft policies, but the factors associated with the network resilience rating indicate that the highest risks are associated with cyber-attacks from external systems, which may cause the adjustment for goal/purpose **2409** to filter out the factors of the network resilience rating associated with risks from external cyber-attacks or reduce their contribution to a functional score. Finally, a functional cybersecurity score **2411** is assigned which takes into account the adjusted factors of the network resilience score and the context in which the functional score is to be applied. The process may be iterative, in that the network resilience rating **2405** from previous analyses may be fed back into the start of the process at estimation of frequency and severity of attacks **2401**.

Detailed Description of Exemplary Aspects

FIG. **29** is a flow diagram illustrating an exemplary method for identifying an anomalous event. In an initial step **2910**, active and passive reconnaissance results may be received at a scoring engine **2400**. This data may then be analyzed **2920** along with a cyber-physical graph to determine patterns of normal event behavior for various nodes within the CPG, based on observed events over time. Event logs, reconnaissance data, and other time-series data may then be monitored **2930** to track event patterns over time, comparing against normal patterns to identify anomalies as well as enabling refinement of the normal pattern records by incorporating additional data as it is produced. When an event occurs that is recognized to be outside the normal patterns, for example if the event exceeds a configured threshold for aberrance or involves a known-compromised node or external address or resource (such as a deliberate "honeypot" node for intrusion detection), it may then be

identified as an anomaly **2940** that may indicate a compromised node or an intrusion attempt. These anomalous events may be used to identify correlations in network behavior and identify attack patterns, as described below in FIGS. **30-32**.

FIG. **30** is a flow diagram illustrating an exemplary method for identifying correlations between an anomaly and other network events and resources. According to this method, when an anomaly is identified **3010** (as described above, with reference to FIG. **29**), it may be correlated with other events, network resources, behavior patterns, or any other available information that may be contextually linked to the anomaly. Using a CPG for the network, correlations may be identified by first comparing resources involved in the anomaly **3020**, for example any machines that participated in the anomalous event or any databases or other resources that were accessed. This information may then be used to further correlate with other events involving those resources **3030**, such as events taking place at the same (or nearly the same) time, or events with a recognizable causative relationship such as non-anomalous events that may have been triggered by the anomaly (for example, as would occur if an attacker accessed a resource and used it to perform routine functions that are not anomalous in themselves, but were clearly related to the anomalous event as they occurred immediately following it during a time when they would normally be unlikely to happen). An established normal behavior pattern (as described above, with reference to FIG. **29**) may also be used to identify additional anomalies **3040** that may not comprise anomalous events (that is, discrete network events that do not correspond to established normal event behavior) but rather deviations from normal behavior that occur around the same time or involving the same resources as an anomalous event. For example, if an anomalous event occurs involving a particular computer on the network, behavior logs may be reviewed **3050** to identify any recent anomalous behavior that may have gone unnoticed, and current behavior **3060** may be checked for any deviations from established normal patterns. These deviations may be determined using preconfigured thresholds, or ad-hoc thresholds that may be established on the fly based on the anomalous event and other available information pertaining to the incident as it unfolds. For example, an event involving a large number of resources or particularly sensitive data stores may cause a threshold for behavior deviation to be set lower than a comparatively minor anomaly involving only a small number of nonessential network resources.

In this manner, anomalous events may be rapidly correlated with any unusual behavior or activity on the network in order to gain a more-complete picture of the extent of any potential attack and reveal additional information that may be used in mitigating the attack, as described below in FIGS. **31-32**.

FIG. **31** is a flow diagram illustrating an exemplary method for generating a behavior graph based on an identified anomaly and its correlations. When an anomaly is detected and correlated **3110** as described above (with reference to FIGS. **29-30**), a DCG may be used to formulate a behavior graph representing the various behavioral connections and interactions between involved resources or processes pertaining to the anomaly, in order to identify possible event flows between affected nodes and gain further insight into the anomaly and any possible effects or outcomes it may produce. Initially **3120**, the DCG may be used to identify all behavior interactions between affected processes and resources, using established known behavior patterns and event correlations to iterate over the affected nodes in a CPG until all relationships have been identified. These behavioral relationships may then be used **3130** to produce an initial behavioral graph that describes the interconnectedness between nodes and the various event pathways that can occur, including potential (or actual) attach paths that an intruder may take advantage of when attempting to compromise the network or any particular nodes. These pathways indicate what dependencies must be satisfied for any given event or node behavior, what the effects of the event or behavior are, what nodes are in turn affected by those outcomes, and other such directed flow information. The graph may then be used to identify **3140** any possible event flows and interactions between the affected nodes, indicating causative effect pathways that are triggered by events and node behaviors, which in the case of an anomaly allows insight into what effects the anomaly may have or what may have led up to the event in question as is described below, in FIG. **32**.

FIG. **32** is a flow diagram illustrating an exemplary method for determining an attack's point of origin using a behavior graph. When a behavior graph is produced for an identified event anomaly **3210**, it may be used to determine event dependencies **3220** based on the identified behavioral patterns and causative relationships in the graph. Using these dependencies, a dependency tree may be produced that comprises the chain of causative events, behaviors, and conditions that led up to the anomaly. This dependency tree comprises the potential event flows and interaction pathways within and between the nodes, and causative relationships can be followed through time as branching paths through the tree. By traversing this dependency tree backward **3230** (that is, following the path of event flows back in time), it then becomes possible to determine a point of origin for the anomaly and the starting conditions that were in place **3240**. This enables more effective mitigation of attacks and resolution of flaws in a network, by identifying the precise causes of vulnerabilities and deficiencies so they may be addressed directly rather than traditional approaches that are inherently limited to mitigating the outcomes of an unknown attack; in other words, traditional approaches focus on current symptoms of an attack whereas this method is able to directly identify and address the root cause of the issue.

FIG. **8** is a flow diagram of an exemplary method **800** for cybersecurity behavioral analytics, according to one aspect. According to the aspect, behavior analytics may utilize passive information feeds from a plurality of existing endpoints (for example, including but not limited to user activity on a network, network performance, or device behavior) to generate security solutions. In an initial step **801**, a web crawler **115** may passively collect activity information, which may then be processed **802** using a DCG **155** to analyze behavior patterns. Based on this initial analysis, anomalous behavior may be recognized **803** (for example, based on a threshold of variance from an established pattern or trend) such as high-risk users or malicious software operators such as bots. These anomalous behaviors may then be used **804** to analyze potential angles of attack and then produce **805** security suggestions based on this second-level analysis and predictions generated by an action outcome simulation module **125** to determine the likely effects of the change. The suggested behaviors may then be automatically implemented **806** as needed. Passive monitoring **801** then continues, collecting information after new security solutions are implemented **806**, enabling machine learning to improve operation over time as the relationship between security changes and observed behaviors and threats are observed and analyzed.

This method **800** for behavioral analytics enables proactive and high-speed reactive defense capabilities against a variety of cyberattack threats, including anomalous human behaviors as well as nonhuman "bad actors" such as automated software bots that may probe for, and then exploit, existing vulnerabilities. Using automated behavioral learning in this manner provides a much more responsive solution than manual intervention, enabling rapid response to threats to mitigate any potential impact. Utilizing machine learning behavior further enhances this approach, providing additional proactive behavior that is not possible in simple automated approaches that merely react to threats as they occur.

FIG. **9** is a flow diagram of an exemplary method **900** for measuring the effects of cybersecurity attacks, according to one aspect. According to the aspect, impact assessment of an attack may be measured using a DCG **155** to analyze a user account and identify its access capabilities **901** (for example, what files, directories, devices or domains an account may have access to). This may then be used to generate **902** an impact assessment score for the account, representing the potential risk should that account be compromised. In the event of an incident, the impact assessment score for any compromised accounts may be used to produce a "blast radius" calculation **903**, identifying exactly what resources are at risk as a result of the intrusion and where security personnel should focus their attention. To provide proactive security recommendations through a simulation module **125**, simulated intrusions may be run **904** to identify potential blast radius calculations for a variety of attacks and to determine **905** high risk accounts or resources so that security may be improved in those key areas rather than focusing on reactive solutions.

FIG. **10** is a flow diagram of an exemplary method **1000** for continuous cybersecurity monitoring and exploration, according to one aspect. According to the aspect, a state observation service **140** may receive data from a variety of connected systems **1001** such as (for example, including but not limited to) servers, domains, databases, or user directories. This information may be received continuously, passively collecting events and monitoring activity over time while feeding **1002** collected information into a graphing service **145** for use in producing time-series graphs **1003** of states and changes over time. This collated time-series data may then be used to produce a visualization **1004** of changes over time, quantifying collected data into a meaningful and understandable format. As new events are recorded, such as changing user roles or permissions, modifying servers or data structures, or other changes within a security infrastructure, these events are automatically incorporated into the time-series data and visualizations are updated accordingly, providing live monitoring of a wealth of information in a way that highlights meaningful data without losing detail due to the quantity of data points under examination.

FIG. **11** is a flow diagram of an exemplary method **1100** for mapping a cyber-physical system graph (CPG), according to one aspect. According to the aspect, a cyber-physical system graph may comprise a visualization of hierarchies and relationships between devices and resources in a security infrastructure, contextualizing security information with physical device relationships that are easily understandable for security personnel and users. In an initial step **1101**, behavior analytics information (as described previously, referring to FIG. **8**) may be received at a graphing service **145** for inclusion in a CPG. In a next step **1102**, impact assessment scores (as described previously, referring to FIG. **9**) may be received and incorporated in the CPG information, adding risk assessment context to the behavior information. In a next step **1103**, time-series information (as described previously, referring to FIG. **10**) may be received and incorporated, updating CPG information as changes occur and events are logged. This information may then be used to produce **1104** a graph visualization of users, servers, devices, and other resources correlating physical relationships (such as a user's personal computer or smartphone, or physical connections between servers) with logical relationships (such as access privileges or database connections), to produce a meaningful and contextualized visualization of a security infrastructure that reflects the current state of the internal relationships present in the infrastructure.

FIG. **12** is a flow diagram of an exemplary method **1200** for continuous network resilience rating, according to one aspect. According to the aspect, a baseline score can be used to measure an overall level of risk for a network infrastructure, and may be compiled by first collecting **1201** information on publicly-disclosed vulnerabilities, such as (for example) using the Internet or common vulnerabilities and exploits (CVE) process. This information may then **1202** be incorporated into a CPG as described previously in FIG. **11**, and the combined data of the CPG and the known vulnerabilities may then be analyzed **1203** to identify the relationships between known vulnerabilities and risks exposed by components of the infrastructure. This produces a combined CPG **1204** that incorporates both the internal risk level of network resources, user accounts, and devices as well as the actual risk level based on the analysis of known vulnerabilities and security risks.

FIG. **13** is a flow diagram of an exemplary method **1300** for cybersecurity privilege oversight, according to one aspect. According to the aspect, time-series data (as described above, referring to FIG. **10**) may be collected **1301** for user accounts, credentials, directories, and other user-based privilege and access information. This data may then **1302** be analyzed to identify changes over time that may affect security, such as modifying user access privileges or adding new users. The results of analysis may be checked **1303** against a CPG (as described previously in FIG. **11**), to compare and correlate user directory changes with the actual infrastructure state. This comparison may be used to perform accurate and context-enhanced user directory audits **1304** that identify not only current user credentials and other user-specific information, but changes to this information over time and how the user information relates to the actual infrastructure (for example, credentials that grant access to devices and may therefore implicitly grant additional access due to device relationships that were not immediately apparent from the user directory alone).

FIG. **14** is a flow diagram of an exemplary method **1400** for cybersecurity risk management, according to one aspect. According to the aspect, multiple methods described previously may be combined to provide live assessment of attacks as they occur, by first receiving **1401** time-series data for an infrastructure (as described previously, in FIG. **10**) to provide live monitoring of network events. This data is then enhanced **1402** with a CPG (as described above in FIG. **11**) to correlate events with actual infrastructure elements, such as servers or accounts. When an event (for example, an attempted attack against a vulnerable system or resource) occurs **1403**, the event is logged in the time-series data **1404**, and compared against the CPG **1405** to determine the impact. This is enhanced with the inclusion of impact assessment information **1406** for any affected resources, and the attack is then checked against a baseline score **1407** to

determine the full extent of the impact of the attack and any necessary modifications to the infrastructure or policies.

FIG. **15** is a flow diagram of an exemplary method **1500** for mitigating compromised credential threats, according to one aspect. According to the aspect, impact assessment scores (as described previously, referring to FIG. **9**) may be collected **1501** for user accounts in a directory, so that the potential impact of any given credential attack is known in advance of an actual attack event. This information may be combined with a CPG **1502** as described previously in FIG. **11**, to contextualize impact assessment scores within the infrastructure (for example, so that it may be predicted what systems or resources might be at risk for any given credential attack). A simulated attack may then be performed **1503** to use machine learning to improve security without waiting for actual attacks to trigger a reactive response. A blast radius assessment (as described above in FIG. **9**) may be used in response **1504** to determine the effects of the simulated attack and identify points of weakness, and produce a recommendation report **1505** for improving and hardening the infrastructure against future attacks.

FIG. **16** is a flow diagram of an exemplary method **1600** for dynamic network and rogue device discovery, according to one aspect. According to the aspect, an advanced cyber decision platform may continuously monitor a network in real-time **1601**, detecting any changes as they occur. When a new connection is detected **1602**, a CPG may be updated **1603** with the new connection information, which may then be compared against the network's resiliency score **1604** to examine for potential risk. The blast radius metric for any other devices involved in the connection may also be checked **1605**, to examine the context of the connection for risk potential (for example, an unknown connection to an internal data server with sensitive information may be considered a much higher risk than an unknown connection to an externally-facing web server). If the connection is a risk, an alert may be sent to an administrator **1606** with the contextual information for the connection to provide a concise notification of relevant details for quick handling.

FIG. **17** is a flow diagram of an exemplary method **1700** for Kerberos "golden ticket" attack detection, according to one aspect. Kerberos is a network authentication protocol employed across many enterprise networks to enable single sign-on and authentication for enterprise services. This makes it an attractive target for attacks, which can result in persistent, undetected access to services within a network in what is known as a "golden ticket" attack. To detect this form of attack, behavioral analytics may be employed to detect forged authentication tickets resulting from an attack. According to the aspect, an advanced cyber decision platform may continuously monitor a network **1701**, informing a CPG in real-time of all traffic associated with entities in an organization, for example, people, places, devices, or services **1702**. Machine learning algorithms detect behavioral anomalies as they occur in real-time **1703**, notifying administrators with an assessment of the anomalous event **1704** as well as a blast radius score for the particular event and a network resiliency score to advise of the overall health of the network. By automatically detecting unusual behavior and informing an administrator of the anomaly along with contextual information for the event and network, a compromised ticket is immediately detected when a new authentication connection is made.

FIG. **18** is a flow diagram of an exemplary method **1800** for risk-based vulnerability and patch management, according to one aspect. According to the aspect, an advanced cyber decision platform may monitor all information about a network **1801**, including (but not limited to) device telemetry data, log files, connections and network events, deployed software versions, or contextual user activity information. This information is incorporated into a CPG **1802** to maintain an up-to-date model of the network in real-time. When a new vulnerability is discovered, a blast radius score may be assessed **1803** and the network's resiliency score may be updated **1804** as needed. A security alert may then be produced **1805** to notify an administrator of the vulnerability and its impact, and a proposed patch may be presented **1806** along with the predicted effects of the patch on the vulnerability's blast radius and the overall network resiliency score. This determines both the total impact risk of any particular vulnerability, as well as the overall effect of each vulnerability on the network as a whole. This continuous network assessment may be used to collect information about new vulnerabilities and exploits to provide proactive solutions with clear result predictions, before attacks occur.

FIG. **20** is a relational diagram showing the relationships between exemplary 3rd party search tools **1915**, search tasks **2010** that can be generated using such tools, and the types of information that may be gathered with those tasks **2011**-**2014**, and how a public-facing proxy network **1908** may be used to influence the search task results. While the use of 3rd party search tools **1915** is in no way required, and proprietary or other self-developed search tools may be used, there are numerous 3rd party search tools **1915** available on the Internet, many of them available for use free of charge, that are convenient for purposes of performing external and internal reconnaissance of an organization's infrastructure. Because they are well-known, they are included here as examples of the types of search tools that may be used and the reconnaissance data that may be gathered using such tools. The search tasks **2010** that may be generated may be classified into several categories. While this category list is by no means exhaustive, several important categories of reconnaissance data are domain and internet protocol (IP) address searching tasks **2011**, corporate information searching tasks **2012**, data breach searching tasks **2013**, and dark web searching tasks **2014**. Third party search tools **1915** for domain and IP address searching tasks **2011** include, for example, DNSDumpster, Spiderfoot HX, Shodan, VirusTotal, Dig, Censys, ViewDNS, and CheckDMARC, among others. These tools may be used to obtain reconnaissance data about an organization's server IPs, software, geolocation; open ports, patch/setting vulnerabilities; data hosting services, among other data **2031**. Third party search tools **1915** for corporate information searching tasks **2012** include, for example, Bloomberg.com, Wikipedia, SEC.gov, AnnualReports.com, DNB.com, Hunter.io, and MarketVisual, among others. These tools may be used to obtain reconnaissance data about an organization's addresses; corp info; high value target (key employee or key data assets) lists, emails, phone numbers, online presence **2032**. Third party search tools **1915** for data breach searching tasks **2013** include, for example, DeHashed, WeLeakInfo, Pastebin, Spiderfoot, and BreachCompilation, among others. These tools may be used to obtain reconnaissance data about an organization's previous data breaches, especially those involving high value targets, and similar data loss information **2033**. Third party search tools **1915** for deep web (reports, records, and other documents linked to in web pages, but not indexed in search results . . . estimated to be 90% of available web content) and dark web (websites accessible only through anonymizers such as TOR . . . estimated to be about 6% of available web content) search-

ing tasks **2014** include, for example, Pipl, MyLife, Yippy, SurfWax, Wayback machine, Google Scholar, DuckDuckGo, Fazzle, Not Evil, and Start Page, among others. These tools may be used to obtain reconnaissance data about an organization's lost and stolen data such as customer credit card numbers, stolen subscription credentials, hacked accounts, software tools designed for certain exploits, which organizations are being targeted for certain attacks, and similar information **2034**. A public-facing proxy network **1908** may be used to change the outward presentation of the organization's network by conducting the searches through selectable attribution nodes **2021***a-n*, which are configurable to present the network to the Internet in different ways such as, but not limited to, presenting the organization network as a commercial IP address, a residential IP address, or as an IP address from a particular country, all of which may influence the reconnaissance data received using certain search tools.

FIG. **21** is a relational diagram showing the exemplary types and classifications of information that may be used in constructing a cyber-physical graph **1902** of an organization's infrastructure and operations. The cyber-physical graph **1902** is a directed graph that represents a comprehensive picture of an organization's infrastructure and operations. A cyber-physical graph **1902** represents the relationships between entities associated with an organization, for example, devices, users, resources, groups, and computing services, the relationships between the entities defining relationships and processes in an organization's infrastructure, thereby contextualizing security information with physical and logical relationships that represent the flow of data and access to data within the organization including, in particular, network security protocols and procedures. Data that may be incorporated into a cyber-physical graph may be any data relating to an organization's infrastructure and operations, and two primary categories of data that may be incorporated are internal reconnaissance data **2110** and external reconnaissance data **2120**. Non-limiting examples of internal reconnaissance data **2110** include computers and devices, physical and intangible (data) assets, people (employees, contractors, etc.), addresses and locations of buildings, servers, etc., business processes, access privileges, loss information, legal documents, and self-assessments of cybersecurity. Non-limiting examples of external reconnaissance data **2120** include domains and IP information, data breach information, organization information such as corporate structures, key employees, etc., open port information, information regarding which organizations are current targets of cyber-attacks, network vulnerability information, system version and patch/update information, known and possible exploits, and publicly available information.

In an initial step **1101**, behavior analytics information (as described previously, referring to FIG. **8**) may be received at a graphing service **145** for inclusion in a CPG. In a next step **1102**, impact assessment scores (as described previously, referring to FIG. **9**) may be received and incorporated in the CPG information, adding risk assessment context to the behavior information. In a next step **1103**, time-series information (as described previously, referring to FIG. **10**) may be received and incorporated, updating CPG information as changes occur and events are logged. This information may then be used to produce **1104** a graph visualization of users, servers, devices, and other resources correlating physical relationships (such as a user's personal computer or smartphone, or physical connections between servers) with logical relationships (such as access privileges or database connections), to produce a meaningful and contextualized visualization of a security infrastructure that reflects the current state of the internal relationships present in the infrastructure.

FIG. **22** is a directed graph diagram showing an exemplary cyber-physical graph **2200** and its possible use in creating cybersecurity profiles and ratings. A cyber-physical graph **1902** represents the relationships between entities associated with an organization, for example, devices, users, resources, groups, and computing services, the relationships between the entities defining relationships and processes in an organization's infrastructure, thereby contextualizing security information with physical and logical relationships that represent the flow of data and access to data within the organization including, in particular, network security protocols and procedures. A cyber-physical graph, in its most basic form, represents the network devices comprising an organization's network infrastructure as nodes (also called vertices) in the graph and the physical or logical connections between them as edges between the nodes. The cyber-physical graph may be expanded to include network information and processes such as data flow, security protocols and procedures, and software versions and patch information. Further, human users and their access privileges to devices and assets may be included. A cyber-security graph may be further expanded to include internal process information such as business processes, loss information, and legal requirements and documents; external information such as domain and IP information, data breach information; and generated information such as open port information from external network scans, and vulnerabilities and avenues of attack. Thus, a cyber-physical graph may be used to represent a complete picture of an organization's infrastructure and operations.

In this example, which is necessarily simplified for clarity, the cyber-physical graph **2200** contains 12 nodes (vertices) comprising: seven computers and devices designated by solid circles **2202**, **2203**, **2204**, **2206**, **2207**, **2209**, **2210**, two users designated by dashed-line circles **2201**, **2211**, and three functional groups designated by dotted-line circles **2205**, **2208**, and **2212**. The edges (lines) between the nodes indicate relationships between the nodes, and have a direction and relationship indicator such as "AdminTo," "MemberOf," etc. While not shown here, the edges may also be assigned numerical weights or probabilities, indicating, for example, the likelihood of a successful attack gaining access from one node to another. Possible attack paths may be analyzed using the cyber-physical graph by running graph analysis algorithms such as shortest path algorithms, minimum cost/maximum flow algorithms, strongly connected node algorithms, etc. In this example, several exemplary attack paths are ranked by likelihood. In the most likely attack path, user **2201** is an administrator to device **2202** to which device **2203** has connected. Device **2203** is a member of functional group **2208**, which has a member of group **2212**. Functional group **2212** is an administrator to the target **2206**. In a second most likely attack path, user **2201** is an administrator to device **2207** to which device **2204** has connected. Device **2204** is a member of functional group **2205**, which is an administrator to the target device **2206**. In a third most likely attack path, a flaw in the security protocols allow the credentials of user **2201** to be used to gain access to device **2210**. User **2211** who is working on device **2210** may be tricked into providing access to functional group **2205**, which is an administrator to the target device **2206**.

FIG. **23** is a block diagram showing exemplary operation of a data to rule mapper. Laws, policies, standards, and other

rules are gathered and stored in an authority database **1903**. Non-limiting examples of such rules include federal, state, and local statutes, regulations, case law interpretations, and other laws **2310**, business policies and procedures **2320**, and industry standards (as one example, cybersecurity industry standards for network security) **2330**. Reconnaissance data are stored in a database **1905**. A data to rule mapper **1904** retrieves the reconnaissance data **1905** and matches it to rules from the authority database **1903**. An example of this operation for statues/regulations **2310** is shown in **2311**, where Article 33, paragraph 1 of the European Union's General Data Protection Regulation (GDPR) requires that an organization notify a cognizant authority of a data breach within 72 hours of knowledge of the breach. If a data point indicates that a data breach has been discovered because data of the organization is found online, the data point is associated with that rule, and tagged with the possible impact of fines if the rule is not followed. An example of this operation for business policies **2320** is shown in **2321**, where a corporate policy prohibits access of the organization's systems using personal computers. If a data point indicates that an employee account is accessed using a non-business-owned computer, the data point is associated with the rule, and tagged with possible data theft and/or security breach. An example of this operation for industry standards **2330** is shown in **2331**, where an industry standard prohibits open ports accessible from outside the network perimeter. If a data point indicates an open port, the data point is associated with the rule, and tagged with possible data loss and/or security breach.

Hardware Architecture

Generally, the techniques disclosed herein may be implemented on hardware or a combination of software and hardware. For example, they may be implemented in an operating system kernel, in a separate user process, in a library package bound into network applications, on a specially constructed machine, on an application-specific integrated circuit (ASIC), or on a network interface card.

Software/hardware hybrid implementations of at least some of the aspects disclosed herein may be implemented on a programmable network-resident machine (which should be understood to include intermittently connected network-aware machines) selectively activated or reconfigured by a computer program stored in memory. Such network devices may have multiple network interfaces that may be configured or designed to utilize different types of network communication protocols. A general architecture for some of these machines may be described herein in order to illustrate one or more exemplary means by which a given unit of functionality may be implemented. According to specific aspects, at least some of the features or functionalities of the various aspects disclosed herein may be implemented on one or more general-purpose computers associated with one or more networks, such as for example an end-user computer system, a client computer, a network server or other server system, a mobile computing device (e.g., tablet computing device, mobile phone, smartphone, laptop, or other appropriate computing device), a consumer electronic device, a music player, or any other suitable electronic device, router, switch, or other suitable device, or any combination thereof. In at least some aspects, at least some of the features or functionalities of the various aspects disclosed herein may be implemented in one or more virtualized computing environments (e.g., network computing clouds, virtual machines hosted on one or more physical computing machines, or other appropriate virtual environments).

Referring now to FIG. **25**, there is shown a block diagram depicting an exemplary computing device **10** suitable for implementing at least a portion of the features or functionalities disclosed herein. Computing device **10** may be, for example, any one of the computing machines listed in the previous paragraph, or indeed any other electronic device capable of executing software- or hardware-based instructions according to one or more programs stored in memory. Computing device **10** may be configured to communicate with a plurality of other computing devices, such as clients or servers, over communications networks such as a wide area network a metropolitan area network, a local area network, a wireless network, the Internet, or any other network, using known protocols for such communication, whether wireless or wired.

In one aspect, computing device **10** includes one or more central processing units (CPU) **12**, one or more interfaces **15**, and one or more busses **14** (such as a peripheral component interconnect (PCI) bus). When acting under the control of appropriate software or firmware, CPU **12** may be responsible for implementing specific functions associated with the functions of a specifically configured computing device or machine. For example, in at least one aspect, a computing device **10** may be configured or designed to function as a server system utilizing CPU **12**, local memory **11** and/or remote memory **16**, and interface(s) **15**. In at least one aspect, CPU **12** may be caused to perform one or more of the different types of functions and/or operations under the control of software modules or components, which for example, may include an operating system and any appropriate applications software, drivers, and the like.

CPU **12** may include one or more processors **13** such as, for example, a processor from one of the Intel, ARM, Qualcomm, and AMD families of microprocessors. In some aspects, processors **13** may include specially designed hardware such as application-specific integrated circuits (ASICs), electrically erasable programmable read-only memories (EEPROMs), field-programmable gate arrays (FPGAs), and so forth, for controlling operations of computing device **10**. In a particular aspect, a local memory **11** (such as non-volatile random access memory (RAM) and/or read-only memory (ROM), including for example one or more levels of cached memory) may also form part of CPU **12**. However, there are many different ways in which memory may be coupled to system **10**. Memory **11** may be used for a variety of purposes such as, for example, caching and/or storing data, programming instructions, and the like. It should be further appreciated that CPU **12** may be one of a variety of system-on-a-chip (SOC) type hardware that may include additional hardware such as memory or graphics processing chips, such as a QUALCOMM SNAPDRAGON™ or SAMSUNG EXYNOS™ CPU as are becoming increasingly common in the art, such as for use in mobile devices or integrated devices.

As used herein, the term "processor" is not limited merely to those integrated circuits referred to in the art as a processor, a mobile processor, or a microprocessor, but broadly refers to a microcontroller, a microcomputer, a programmable logic controller, an application-specific integrated circuit, and any other programmable circuit.

In one aspect, interfaces **15** are provided as network interface cards (NICs). Generally, NICs control the sending and receiving of data packets over a computer network; other types of interfaces **15** may for example support other peripherals used with computing device **10**. Among the interfaces that may be provided are Ethernet interfaces, frame relay interfaces, cable interfaces, DSL interfaces,

token ring interfaces, graphics interfaces, and the like. In addition, various types of interfaces may be provided such as, for example, universal serial bus (USB), Serial, Ethernet, FIREWIRE™, THUNDERBOLT™, PCI, parallel, radio frequency (RF), BLUETOOTH™, near-field communications (e.g., using near-field magnetics), 802.11 (WiFi), frame relay, TCP/IP, ISDN, fast Ethernet interfaces, Gigabit Ethernet interfaces, Serial ATA (SATA) or external SATA (ESATA) interfaces, high-definition multimedia interface (HDMI), digital visual interface (DVI), analog or digital audio interfaces, asynchronous transfer mode (ATM) interfaces, high-speed serial interface (HSSI) interfaces, Point of Sale (POS) interfaces, fiber data distributed interfaces (FDDIs), and the like. Generally, such interfaces **15** may include physical ports appropriate for communication with appropriate media. In some cases, they may also include an independent processor (such as a dedicated audio or video processor, as is common in the art for high-fidelity A/V hardware interfaces) and, in some instances, volatile and/or non-volatile memory (e.g., RAM).

Although the system shown in FIG. **25** illustrates one specific architecture for a computing device **10** for implementing one or more of the aspects described herein, it is by no means the only device architecture on which at least a portion of the features and techniques described herein may be implemented. For example, architectures having one or any number of processors **13** may be used, and such processors **13** may be present in a single device or distributed among any number of devices. In one aspect, a single processor **13** handles communications as well as routing computations, while in other aspects a separate dedicated communications processor may be provided. In various aspects, different types of features or functionalities may be implemented in a system according to the aspect that includes a client device (such as a tablet device or smartphone running client software) and server systems (such as a server system described in more detail below).

Regardless of network device configuration, the system of an aspect may employ one or more memories or memory modules (such as, for example, remote memory block **16** and local memory **11**) configured to store data, program instructions for the general-purpose network operations, or other information relating to the functionality of the aspects described herein (or any combinations of the above). Program instructions may control execution of or comprise an operating system and/or one or more applications, for example. Memory **16** or memories **11**, **16** may also be configured to store data structures, configuration data, encryption data, historical system operations information, or any other specific or generic non-program information described herein.

Because such information and program instructions may be employed to implement one or more systems or methods described herein, at least some network device aspects may include nontransitory machine-readable storage media, which, for example, may be configured or designed to store program instructions, state information, and the like for performing various operations described herein. Examples of such nontransitory machine-readable storage media include, but are not limited to, magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROM disks; magneto-optical media such as optical disks, and hardware devices that are specially configured to store and perform program instructions, such as read-only memory devices (ROM), flash memory (as is common in mobile devices and integrated systems), solid state drives (SSD) and "hybrid SSD" storage drives that may combine physical components of solid state and hard disk drives in a single hardware device (as are becoming increasingly common in the art with regard to personal computers), memristor memory, random access memory (RAM), and the like. It should be appreciated that such storage means may be integral and non-removable (such as RAM hardware modules that may be soldered onto a motherboard or otherwise integrated into an electronic device), or they may be removable such as swappable flash memory modules (such as "thumb drives" or other removable media designed for rapidly exchanging physical storage devices), "hot-swappable" hard disk drives or solid state drives, removable optical storage discs, or other such removable media, and that such integral and removable storage media may be utilized interchangeably. Examples of program instructions include both object code, such as may be produced by a compiler, machine code, such as may be produced by an assembler or a linker, byte code, such as may be generated by for example a JAVA™ compiler and may be executed using a Java virtual machine or equivalent, or files containing higher level code that may be executed by the computer using an interpreter (for example, scripts written in Python, Perl, Ruby, Groovy, or any other scripting language).

In some aspects, systems may be implemented on a standalone computing system. Referring now to FIG. **26**, there is shown a block diagram depicting a typical exemplary architecture of one or more aspects or components thereof on a standalone computing system. Computing device **20** includes processors **21** that may run software that carry out one or more functions or applications of aspects, such as for example a client application **24**. Processors **21** may carry out computing instructions under control of an operating system **22** such as, for example, a version of MICROSOFT WINDOWS™ operating system, APPLE macOS™ or iOS™ operating systems, some variety of the Linux operating system, ANDROID™ operating system, or the like. In many cases, one or more shared services **23** may be operable in system **20**, and may be useful for providing common services to client applications **24**. Services **23** may for example be WINDOWS™ services, user-space common services in a Linux environment, or any other type of common service architecture used with operating system **21**. Input devices **28** may be of any type suitable for receiving user input, including for example a keyboard, touchscreen, microphone (for example, for voice input), mouse, touchpad, trackball, or any combination thereof. Output devices **27** may be of any type suitable for providing output to one or more users, whether remote or local to system **20**, and may include for example one or more screens for visual output, speakers, printers, or any combination thereof. Memory **25** may be random-access memory having any structure and architecture known in the art, for use by processors **21**, for example to run software. Storage devices **26** may be any magnetic, optical, mechanical, memristor, or electrical storage device for storage of data in digital form (such as those described above, referring to FIG. **25**). Examples of storage devices **26** include flash memory, magnetic hard drive, CD-ROM, and/or the like.

In some aspects, systems may be implemented on a distributed computing network, such as one having any number of clients and/or servers. Referring now to FIG. **27**, there is shown a block diagram depicting an exemplary architecture **30** for implementing at least a portion of a system according to one aspect on a distributed computing network. According to the aspect, any number of clients **33** may be provided. Each client **33** may run software for implementing client-side portions of a system; clients may

comprise a system **20** such as that illustrated in FIG. **26**. In addition, any number of servers **32** may be provided for handling requests received from one or more clients **33**. Clients **33** and servers **32** may communicate with one another via one or more electronic networks **31**, which may be in various aspects any of the Internet, a wide area network, a mobile telephony network (such as CDMA or GSM cellular networks), a wireless network (such as WiFi, WiMAX, LTE, and so forth), or a local area network (or indeed any network topology known in the art; the aspect does not prefer any one network topology over any other). Networks **31** may be implemented using any known network protocols, including for example wired and/or wireless protocols.

In addition, in some aspects, servers **32** may call external services **37** when needed to obtain additional information, or to refer to additional data concerning a particular call. Communications with external services **37** may take place, for example, via one or more networks **31**. In various aspects, external services **37** may comprise web-enabled services or functionality related to or installed on the hardware device itself. For example, in one aspect where client applications **24** are implemented on a smartphone or other electronic device, client applications **24** may obtain information stored in a server system **32** in the cloud or on an external service **37** deployed on one or more of a particular enterprise's or user's premises. In addition to local storage on servers **32**, remote storage **38** may be accessible through the network(s) **31**.

In some aspects, clients **33** or servers **32** (or both) may make use of one or more specialized services or appliances that may be deployed locally or remotely across one or more networks **31**. For example, one or more databases **34** in either local or remote storage **38** may be used or referred to by one or more aspects. It should be understood by one having ordinary skill in the art that databases in storage **34** may be arranged in a wide variety of architectures and using a wide variety of data access and manipulation means. For example, in various aspects one or more databases in storage **34** may comprise a relational database system using a structured query language (SQL), while others may comprise an alternative data storage technology such as those referred to in the art as "NoSQL" (for example, HADOOP CASSANDRA™, GOOGLE BIGTABLE™, and so forth). In some aspects, variant database architectures such as column-oriented databases, in-memory databases, clustered databases, distributed databases, or even flat file data repositories may be used according to the aspect. It will be appreciated by one having ordinary skill in the art that any combination of known or future database technologies may be used as appropriate, unless a specific database technology or a specific arrangement of components is specified for a particular aspect described herein. Moreover, it should be appreciated that the term "database" as used herein may refer to a physical database machine, a cluster of machines acting as a single database system, or a logical database within an overall database management system. Unless a specific meaning is specified for a given use of the term "database", it should be construed to mean any of these senses of the word, all of which are understood as a plain meaning of the term "database" by those having ordinary skill in the art.

Similarly, some aspects may make use of one or more security systems **36** and configuration systems **35**. Security and configuration management are common information technology (IT) and web functions, and some amount of each are generally associated with any IT or web systems. It should be understood by one having ordinary skill in the art that any configuration or security subsystems known in the art now or in the future may be used in conjunction with aspects without limitation, unless a specific security **36** or configuration system **35** or approach is specifically required by the description of any specific aspect.

FIG. **28** shows an exemplary overview of a computer system **40** as may be used in any of the various locations throughout the system. It is exemplary of any computer that may execute code to process data. Various modifications and changes may be made to computer system **40** without departing from the broader scope of the system and method disclosed herein. Central processor unit (CPU) **41** is connected to bus **42**, to which bus is also connected memory **43**, nonvolatile memory **44**, display **47**, input/output (I/O) unit **48**, and network interface card (NIC) **53**. I/O unit **48** may, typically, be connected to peripherals such as a keyboard **49**, pointing device **50**, hard disk **52**, real-time clock **51**, a camera **57**, and other peripheral devices. NIC **53** connects to network **54**, which may be the Internet or a local network, which local network may or may not have connections to the Internet. The system may be connected to other computing devices through the network via a router **55**, wireless local area network **56**, or any other network connection. Also shown as part of system **40** is power supply unit **45** connected, in this example, to a main alternating current (AC) supply **46**. Not shown are batteries that could be present, and many other devices and modifications that are well known but are not applicable to the specific novel functions of the current system and method disclosed herein. It should be appreciated that some or all components illustrated may be combined, such as in various integrated applications, for example Qualcomm or Samsung system-on-a-chip (SOC) devices, or whenever it may be appropriate to combine multiple capabilities or functions into a single hardware device (for instance, in mobile devices such as smartphones, video game consoles, in-vehicle computer systems such as navigation or multimedia systems in automobiles, or other integrated hardware devices).

In various aspects, functionality for implementing systems or methods of various aspects may be distributed among any number of client and/or server components. For example, various software modules may be implemented for performing various functions in connection with the system of any particular aspect, and such modules may be variously implemented to run on server and/or client components.

The skilled person will be aware of a range of possible modifications of the various aspects described above. Accordingly, the present invention is defined by the claims and their equivalents.

Where a method is claimed below, the steps of the method are intended to be construed without any order unless it is impossible to perform a given step before another.

What is claimed is:

1. A computer system comprising:

a hardware memory, wherein the computer system is configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that:

store in the hardware memory a representation of a first graph, wherein the representation of the first graph comprises representations of a first plurality of nodes corresponding to a first plurality of entities and further comprises representations of a first plurality of edges, wherein the first graph is a directed graph,

wherein the first plurality of entities comprises a plurality of accounts and a plurality of resources, and

wherein each edge of the first plurality of edges corresponds to a respective relationship between a respective pair of entities;

receive streaming data comprising time-stamped data about events relating to one or more entities of the first plurality of entities,

based on a first portion of the streaming data, identify a first entity that does not correspond to any of the first plurality of nodes, wherein the first entity is not of the first plurality of entities,

based on a second portion of the streaming data, wherein the second portion is not identical to the first portion, identify a first relationship between a pair of entities of the first plurality of entities that does not correspond to any of the first plurality of edges,

modify, in the hardware memory, the representation of the first graph to generate a modified representation of the first graph, wherein the modified representation of the first graph comprises a representation of a first node corresponding to the first entity and a representation of a first edge corresponding to the first relationship, wherein the first node is not of the first plurality of nodes and the first edge is not of the first plurality of edges,

identify, based on the modified representation of the first graph, an attack path that could be involved in an attack involving the first entity, wherein identifying the attack path comprises:

identifying a second entity that can be reached using the first entity, wherein the second entity corresponds to a second node, and the second node is related by one or more edges to the first node corresponding to the first entity in the modified representation of the first graph; and,

identifying a third entity that can be reached using the second entity, wherein the third entity corresponds to a third node, and the third node is related by one or more edges to the second node in the modified representation of the first graph; and

generate a report comprising an identification of the first entity and at least one of the second entity and the third entity.

2. The computer system of claim **1**, wherein identifying the attack path further comprises:

identifying a first plurality of event flows that include a first anomalous event associated with the first entity.

3. The system of claim **2**, wherein identifying the attack path further comprises:

identifying a point of origin for the first anomalous event based on the first plurality of event flows.

4. The computer system of claim **3**,

wherein identifying the point of origin for the first anomalous event comprises:

traversing a second graph, wherein a representation of the second graph is stored in the hardware memory, wherein the representation of the second graph comprises a representation of a second plurality of nodes and a representation of a second plurality of edges, wherein each node of the second plurality of nodes corresponds to a respective node of the first plurality of nodes;

processing a representation of a first node of the second plurality of nodes,

wherein the first node corresponds to a second node of the first plurality of nodes;

identifying an edge of the second plurality of edges, wherein the edge connects a third node of the first plurality of nodes with the first node, and

processing a representation of the third node.

5. The computer system of claim **1**, wherein the modified representation of the first graph comprises a representation of a node corresponding to the first entity, wherein the first entity is identified based on active reconnaissance results.

6. The computer system of claim **1**, wherein the modified representation of the first graph comprises a representation of a node corresponding to the first entity, wherein the first entity is identified based on passive reconnaissance results.

7. The computer system of claim **1**, wherein the attack path is identified at least in part by processing at least a portion of the modified representation of the first graph using at least one transformation pipeline.

8. The computer system of claim **1**, wherein the computer system is further configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that:

based on a third portion of the streaming data, detect a change in a relationship corresponding to an existing edge of the first plurality of edges,

wherein the modified representation of the first graph further comprises a modified representation of the first plurality of edges that accounts for the detected change in the relationship, and

wherein the report further comprises a visual representation of the detected change in the relationship.

9. The computer system of claim **1**, wherein the first graph is a cyber-physical graph.

10. The computer system of claim **9**, wherein the report is an alert.

11. The computer system of claim **1**, wherein the computer system is further configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that determine whether an event, of the events relating to one or more entities of the first plurality of entities, is anomalous, wherein determining whether the event is anomalous comprises:

determining that the event relates to the first entity,

determining at least one behavior pattern associated with the first entity, and

comparing the event to the at least one behavior pattern.

12. The computer system of claim **11**, wherein comparing the event to the at least one behavior pattern comprises using a threshold.

13. The system of claim **11**, wherein comparing the event to the at least one behavior pattern further comprises:

determining that the first entity is associated with a resource,

wherein the resource is a sensitive resource.

14. The computer system of claim **1**, wherein the computer system comprises a plurality of physical computing machines.

15. The computer system of claim **1**, wherein the computer system comprises a plurality of virtual computing machines.

16. The computer system of claim **1**, wherein at least one entity of the first plurality of entities is at least one of a user, a place, a device, a resource, a group, or a service.

17. The computer system of claim **1**, wherein the representation of the first graph comprises a representation of at least one node that does not correspond to an entity.

18. A computer system comprising:

a hardware memory, wherein the computer system is configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that:

store in the hardware memory a representation of a first graph, wherein the representation of the first graph comprises representations of a first plurality of nodes corresponding to a first plurality of entities and further comprises representations of a first plurality of edges, wherein the first graph is a directed graph, wherein the first plurality of entities comprises a plurality of accounts and a plurality of resources, and

wherein each edge of the first plurality of edges corresponds to a respective relationship between a respective pair of entities;

identify, based on the representation of the first graph, a first plurality of attack paths comprising a first entity of the first plurality of entities, wherein each attack path of the first plurality of attack paths targets a second entity of the first plurality of entities,

receive streaming data comprising time-stamped data about events relating to one or more entities of the first plurality of entities,

based on a first portion of the streaming data, identify a third entity that does not correspond to any of the first plurality of nodes, wherein the third entity is not of the first plurality of entities,

based on a second portion of the streaming data, wherein the second portion is not identical to the first portion, identify a first relationship between a pair of entities of the first plurality of entities that does not correspond to any of the first plurality of edges,

modify, in the hardware memory, the representation of the first graph to generate a modified representation of the first graph, wherein the modified representation of the first graph comprises a representation of a node corresponding to the third entity and an edge corresponding to the first relationship, wherein the node is not of the first plurality of nodes and the edge is not of the first plurality of edges, and

identify, based on the modified representation of the first graph, a second plurality of attack paths comprising the first entity, wherein each attack path of the second plurality of attack paths targets the second entity, and wherein an attack path of the second plurality of attack paths comprises the third entity.

19. The computer system of claim 18, wherein the first graph is a cyber-physical graph.

20. The computer system of claim 18, wherein identifying the second plurality of attack paths comprises:

identifying a first plurality of event flows that could involve the first entity, and

identifying a first plurality of dependencies associated with the first plurality of event flows.

21. The computer system of claim 20, further configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that:

based on a third portion of the streaming data, identify an anomalous event,

store in the hardware memory a dependency tree for the anomalous event, and

based on the dependency tree, determine a point of origin for the anomalous event.

22. The computer system of claim 18, wherein a portion of the representation of the first graph is derived from reconnaissance data received by the computer system from a plurality of computer systems,

wherein a first computer system of the plurality of computer systems performs passive reconnaissance, and

wherein a second computer system of the plurality of computer systems performs active reconnaissance.

23. The computer system of claim 18, further configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that:

identify a third plurality of attack paths, comprising attack paths of the second plurality of attack paths that are not in the first plurality of attack paths, and

generate a report comprising information about the third plurality of attack paths.

24. The computer system of claim 18, further configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that analyze an attack path of the plurality of attack paths using a graph analysis algorithm.

25. The computer system of claim 24, wherein the graph analysis algorithm comprises a shortest-path algorithm.

26. The computer system of claim 25, further configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that generate a report comprising information obtained from the graph analysis algorithm.

27. The computer system of claim 18, wherein the representation of the first graph comprises an identification of the third entity as a sensitive resource, and wherein the computer system is further configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that generate a report that identifies the third entity as a sensitive resource.

28. The computer system of claim 18, wherein the representation of the first graph comprises an identification of a security vulnerability associated with the third entity, and wherein an attack path of the second plurality of attack paths is based on the security vulnerability.

* * * * *