

CONTINUOUS VULNERABILITY MANAGEMENT SYSTEM FOR BLOCKCHAIN SMART CONTRACT BASED DIGITAL ASSET USING SANDBOX AND ARTIFICIAL INTELLIGENCE

BACKGROUND

[0001] The field of smart contract auditing, in its relatively brief history, already presents a number of challenges to ensure protection and/or rapid reaction to vulnerabilities. The motivation for such protection and reaction is manifest: smart contract vulnerabilities can cost financial institutions millions of dollars. Two cases in the last few years -- the BeautyChainTM token (2018) containing an overflow bug leading to more than two billion USD of market value and a DAO hack (2016), due to reentrancy vulnerability, causing a multimillion-dollar loss -- provide evidence of smart contract vulnerability challenges. A smart contract of the crypto asset has risk exposure to vulnerabilities, and therefore needs continuous protection.

[0002] Existing solutions to smart contract vulnerabilities have attempted to detect, isolate, control and report on vulnerabilities, but they generally suffer from the same defects allowing the vulnerabilities to be introduced in the first place. Auditing too often takes the form of manual expert auditing, e.g., ad hoc checks and "eyeballing." Formal verification processes, i.e., mathematically proving correct code, may be in place, but are highly prone to deterministic complexity, or an inability to check if varying inputs and initial conditions lead to differing results. Moreover, the algorithms often used for formal verifications require manual annotations in the code to guide the mathematical prover algorithm, thereby greatly increasing the likelihood of manual error(s). Some open source tools may be able to analyze code text but they lack systematic design and maintainability, and therefore don't meet an enterprise standard.

[0003] There is, therefore, a need for a smart contract monitoring solution that meets the exacting financial enterprise grade of a vulnerability management standard, including automation of monitoring and running algorithms, along with systematic design and maintainability.

BRIEF SUMMARY

[0004] A continuous vulnerability management system for blockchain smart contract based digital assets using a sandbox and artificial intelligence is disclosed. In one embodiment, a

vulnerability management system for blockchain digital assets includes a sandbox engine configured to scan digital assets for vulnerabilities. The system also includes a knowledge base engine including artificial intelligence, where the knowledge base engine is configured to acquire information related to new vulnerabilities and threat intelligence. A risk scoring engine including machine learning is also part of the system, where the risk scoring engine is configured to analyze sandbox engine outcomes and assign a risk score to each vulnerable asset. A security control system is configured to take an action on an identified vulnerable asset based on the risk score assigned to the identified vulnerable asset.

[0005] In another embodiment, a method of protecting digital assets is disclosed, the method comprising scanning digital assets for vulnerabilities using a sandbox engine and acquiring information related to new vulnerabilities and threat intelligence using a knowledge base engine including artificial intelligence. The method includes analyzing sandbox engine outcomes and assigning a risk score to each vulnerable asset using a risk scoring engine including machine learning. Additionally, the method includes taking an action on an identified vulnerable asset based on the risk score assigned to the identified vulnerable asset using a security control system.

[0006] In still another embodiment, a non-transitory computer-readable storage medium is disclosed, the computer-readable storage medium including instructions that when executed by a computer, cause the computer to scan digital assets for vulnerabilities using a sandbox engine, and acquire information related to new vulnerabilities and threat intelligence using a knowledge base engine including artificial intelligence. The executed instructions also cause the computer to analyze sandbox engine outcomes and assign a risk score to each vulnerable asset using a risk scoring engine including machine learning. Additionally, the executed instructions cause the computer to take an action on an identified vulnerable asset based on the risk score assigned to the identified vulnerable asset using a security control system.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

[0007] To easily identify the discussion of any particular element or act, the most significant digit or digits in a reference number refer to the figure number in which that element is first introduced.

[0008] FIG. 1 illustrates a continuous vulnerability management system 100 in accordance with one embodiment.

- [0009] FIG. 2 illustrates sandbox engine units 200 in accordance with one embodiment.
- [0010] FIG. 3 illustrates a dynamic execution in virtual machine 300 in accordance with one embodiment.
- [0011] FIG. 4 illustrates a machine learning 400 in accordance with one embodiment.
- [0012] FIG. 5 illustrates a risk scoring engine 500 in accordance with one embodiment.
- [0013] FIG. 6 illustrates a digital asset trading exchange 600 in accordance with one embodiment.
- [0014] FIG. 7 illustrates a proactive scanning 700 in accordance with one embodiment.
- [0015] FIG. 8 illustrates a smart contract 800 in accordance with one embodiment.
- [0016] FIG. 9 illustrates a directed acyclic graph 900 in accordance with one embodiment.
- [0017] FIG. 10 illustrates an audit report 1000 in accordance with one embodiment.
- [0018] FIG. 11 illustrates a method 1100 of protecting digital assets in accordance with one embodiment.
- [0019] FIG. 12 illustrates a blockchain transaction process 1200 in accordance with one embodiment.
- [0020] FIG. 13 illustrates a blockchain formation 1300 in accordance with one embodiment.
- [0021] FIG. 14 illustrates a blockchain 1400 in accordance with one embodiment.
- [0022] FIG. 15 illustrates a basic deep neural network 1500 in accordance with one embodiment.
- [0023] FIG. 16 illustrates an artificial neuron 1600 in accordance with one embodiment.
- [0024] FIG. 17 depicts an illustrative computer system architecture that may be used in accordance with one or more illustrative aspects described herein.

DETAILED DESCRIPTION

[0025] "Accidental fork" refers to a branching in the chain that happens when two or more miners find a block at nearly the same time. The fork is resolved when subsequent block(s) are added and one of the chains becomes longer than the alternative(s). The network abandons the blocks that are not in the longest chain (they are called orphaned blocks).

[0026] "Action" refers to a control signal response to discovery of a vulnerable asset by the security control system, e.g., quarantining, adding to known vulnerable digital assets.

[0027] "Artificial intelligence" refers to the theory and development of computer systems able to perform tasks that normally require human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.

[0028] "Artificial Intelligence system" refers to a system including artificial intelligence methods and/or algorithms.

[0029] "At least one database of digital assets" refers to two or more of a collection of digital assets retained in a database, such as the database of global statistics.

[0030] "Audit report generation unit" refers to a means of creating an audit report. Embodiments of the audit report may include the location of vulnerabilities within a digital asset using a heatmap based on the statistical vulnerability data of the statistical analysis unit.

[0031] "Block time" refers to the average time it takes for the network to generate one extra block in the blockchain.

[0032] "Blockchain" refers to a distributed, decentralized, immutable database used to store encrypted data.

[0033] "Code" refers to a set of instructions forming a computer program which is executed by a computer. Code may take the form of source code, compiled code, machine code, etc.

[0034] "Compile" refers to a program or application that translates human-readable source code into computer-executable machine code.

[0035] "Compiled code" refers to computer-readable instructions translated from human-readable source code into computer-executable machine code.

[0036] "Crypto asset" refers to a digital asset designed to work as a medium of exchange that uses strong cryptography to secure financial transactions, control the creation of additional units, and verify the transfer of assets.

[0037] "Cryptocurrency" refers to a digital currency in which encryption techniques are used to regulate the generation of units of currency and verify the transfer of funds, operating independently of a central bank.

[0038] "Database of global statistics" refers to collection of statistical analysis on known digital assets and digital asset vulnerabilities.

[0039] "Digital asset" refers to any text or media that is formatted into a binary source and includes the right to use. An exemplary digital asset is a smart contract.

[0040] "Digital asset trading exchange" refers to a platform that allows customers to trade cryptocurrencies or digital currencies for other assets, such as conventional fiat money or other digital currencies.

[0041] "Dynamic analysis" refers to a testing procedure comprising part of the software debugging process and used to evaluate a program during real-time execution.

[0042] "Dynamic analysis capability" refers to the ability to perform dynamic analysis.

[0043] "Dynamic analysis unit" refers to the portion of the sandbox engine dedicated to dynamic analysis.

[0044] "Fork" refers to what happens when a blockchain diverges into two potential paths forward.

[0045] "Fuzzing" refers to an automated software testing technique that involves providing invalid, unexpected, or random data as inputs to a computer program. May also be referred to as fuzz testing.

[0046] "Hard fork" refers to a rule change such that the software validating blocks according to the old rules will detect the blocks produced according to the new rules as invalid.

[0047] "Heatmap" refers to a representation of data in the form of a map or diagram in which data values are represented as shades or colors.

[0048] "Heatmap data" refers to numeric or other data used to generate a heatmap.

[0049] "Identified vulnerable asset" refers to a digital asset found to include vulnerabilities and included in known vulnerable digital assets.

[0050] "Key-value pairs" refers to an abstract data type that includes a group of key identifiers and a set of associated values. Key-value pairs are frequently used in lookup tables, hash tables and configuration files.

[0051] "Knowledge base" refers to a technology used to store complex structured and unstructured information used by a computer system.

[0052] "Knowledge base engine" refers to a form of artificial intelligence (AI) that aims to capture the knowledge of human experts to support decision-making.

[0053] "Known legitimate digital assets" refers to a collection of digital assets, typically retained in a database, proven to be without vulnerabilities.

[0054] "Known vulnerable digital assets" refers to a collection of digital assets, typically retained in a database, proven to include vulnerabilities.

[0055] "Legitimate digital assets" refers to Digital assets proven to be without vulnerabilities.

[0056] "Machine learning" refers to an application of artificial intelligence (AI) that provides systems the ability to automatically learn and improve from experience without being explicitly programmed.

[0057] "New vulnerabilities" refers to Vulnerabilities previously undiscovered and not registered in a database of known vulnerable digital assets.

[0058] "Offline training engine" refers to a non-executable version of a software process in which a model, typically based on artificial intelligence, learns patterns or sequences based on specific inputs.

[0059] "Online prediction engine" refers to a dynamic version of a software process in which a model, typically based on artificial intelligence, makes predictions based on learning from at least an offline training engine.

[0060] "Orphaned blocks" refers to blocks in an abandoned fork.

[0061] "Private blockchain" refers to (i.e., permissioned blockchains) blockchains that use an access control layer to govern who has access to the network. In contrast to public blockchain networks, validators on private blockchain networks are vetted by the network owner. They do not rely on anonymous nodes to validate transactions nor do they benefit from the network effect. Private blockchains can also go by the name of 'consortium' or 'hybrid' blockchains.

[0062] "Provided training data" refers to inputs to an offline training engine, typically comprising feature vectors from units within the sandbox engine.

[0063] "Reentrancy vulnerability pattern" refers to a pattern in the software of digital assets in which vulnerability consists of multiple invocations being able to run concurrently.

[0064] "Regression model" refers to a statistical model used to investigate the relationship between two or more variables and estimate one variable based on the others.

[0065] "Risk score" refers to the numeric output of a risk scoring engine.

[0066] "Risk scoring engine" refers to a software application or process assigning a numeric score within a standardized scale that assesses the potential vulnerability of a digital asset or set of digital assets.

[0067] "Sandbox" refers to a testing environment that isolates untested code changes and experimentation from the production environment or repository.

[0068] "Sandbox engine" refers to a software application for executing a sandbox.

[0069] "Sandbox engine outcomes" refers to data, evaluations, and/or analyses as generated by a sandbox engine. Specifically, the results of the static analysis, dynamic analysis, and statistical analysis of the sandbox engine.

[0070] "Security control system" refers to an access control system able to characterize activity and incoming data to identify patterns and recognize certain conditions that increase the threat to an asset.

[0071] "Similar digital assets" refers to digital assets exhibiting similar qualities based on static analysis, dynamic analysis, and statistical analysis.

[0072] "Smart contract" refers to contracts that can be partially or fully executed or enforced without human interaction. One of the main objectives of a smart contract is automated escrow. The IMF believes smart contracts based on blockchain technology could reduce moral hazards and optimize the use of contracts in general. Due to the lack of widespread use their legal status is unclear. a blockchain implementation that enables the coding of contracts that execute when specified conditions are met. A blockchain smart contract is enabled by logic that defines and executes an agreement. For example, ethereum solidity is an open-source blockchain project that was built specifically to realize this possibility by implementing a Turing-complete programming language capability to implement smart contracts.

[0073] "Soft fork" refers to a change of rules that creates blocks recognized as valid by the old software, i.e. it is backwards-compatible.

[0074] "Static analysis" refers to a method of analyzing and evaluating search code without executing a program.

[0075] "Static analysis unit" refers to an area of the sandbox engine dedicated to static analysis.

[0076] "Static code scanning" refers to the act of searching static code for known or unknown vulnerabilities.

[0077] "Statistical analysis unit" refers to an area of the sandbox engine dedicated to analyzing risk scores against a database of global statistics.

[0078] "Statistical vulnerability data" refers to a collection of scored vulnerabilities for known vulnerable digital assets from the database of global statistics.

[0079] "Threat intelligence" refers to information an organization uses to understand the threats that have, will, or are currently targeting the organization. This information is used to prepare, prevent, and identify cyber threats looking to take advantage of valuable resources.

[0080] "Trade" refers to an exchange of digital assets, typically within an exchange.

[0081] "Trading" refers to the act of exchanging trades.

[0082] "Virtual machine" refers to a software program or operating system exhibiting the behavior of a separate computer, but capable of performing tasks such as running applications and programs like a separate computer.

[0083] "Vulnerabilities" refers to a cyber-security term that refers to flaws in a system that can leave it open to attack. A vulnerability may also refer to any type of weakness in a computer system itself, in a set of procedures, or in anything that leaves information security exposed to a threat.

[0084] "Vulnerable asset" refers to a digital asset found to include vulnerabilities.

[0085] Given the challenges described above, an innovative system design and method for continuous protection of smart contracts against known and detectable vulnerabilities comprises a wide-ranging set of characteristics. First, the design and method may include behavioral accuracy. Because smart contracts often perform automated functions (e.g., on a blockchain), how these contracts act may present at least the same threat of vulnerabilities as that of source code text. Additionally, analysis may require automation to reduce the likelihood of manual errors inherent in code review by human experts or adding auxiliary code comments to help analysis such as formal verification.

[0086] Next, since new vulnerabilities appear almost constantly, the design and method may require adaptability: a vulnerabilities knowledge base and a blockchain smart database may expand as new vulnerabilities are detected allowing customers to stay ahead of risk factors on their digital assets. Along with adaptability, the engines and analysis may also require learning, e.g., machine learning / artificial intelligence, through model training and similar methods to intelligently learn from an abundance of data. Finally, the design and method may require extensibility: the various engines and analysis processes deployed to detect digital asset vulnerabilities may be extended, replaced, upgraded or otherwise replaced as new threats arise or new vulnerability detection technologies become available.

[0087] Throughout this disclosure, specific examples using Solidity Smart Contract for Ethereum blockchain are used, but the disclosure is not limited thereto. The systems and methods in this disclosure may be applied to any compatible blockchains such as HYPERLEDGER (IBM), QUORUM (JP Morgan Chase), CORDA (R3 LLC), EOSIO (C++ and Web Assembly smart contract), TRON (Raybo Technology Co.), and emerging Ethereum smart contract languages such as Ethereum WebAssembly (eWASM).

[0088] Referring to FIG. 1, a continuous vulnerability management system 100 is shown according to an embodiment. The continuous vulnerability management system 104 comprises a number of components that pass a crypto asset smart contract 106 through various algorithms and processes to gauge the vulnerabilities of the crypto asset smart contract 106, assign a risk score, report on the score and potentially feed back a control signal 132 advising actions such as suspension (of the contract), quarantining/resuming, and so on.

[0089] The continuous vulnerability management system 104 first receives a crypto asset smart contract 106 from a financial institution 102 that has arranged for checking vulnerabilities of its smart contracts. Smart contracts, generally, are self-executing contracts, e.g., can be sent automatically without third party intervention, with the terms of the agreement between buyer and seller directly written into lines of code. A crypto asset is a digital asset 130 generally designed to work as a medium of exchange that uses strong cryptography to secure financial transactions, control the creation of additional units, and verify the transfer of assets. The code and the agreements contained herein exist across a distributed, decentralized blockchain network, the latter facilitating the self-executing nature of the smart contracts.

[0090] A sandbox engine 110, continually scans, e.g., without manual intervention, for vulnerabilities in digital assets such as the crypto asset smart contract 106 since a portfolio of active trades from a financial institution 102 may include between, for example, 50 and 100 digital assets. The sandbox engine 110 algorithms make minute comparisons of code -- both static and compiled -- between the received crypto asset smart contract 106 and similar digital assets in a knowledge base engine 122. The sandbox engine 110 performs static code scanning, e.g., scanning for anomalies and errors in static text code via searching and pattern matching techniques well known to those skilled in the art. The sandbox engine 110 may additionally perform a dynamic analysis on the results of compiling the smart contract source code into compiled code and executing the compiled code in a virtual machine environment, e.g., a software program or operating system capable of performing tasks such as running applications

and programs like a separate computer. The knowledge base engine 122 may include outputs from threat intelligence 126, a continuously updated and published database of source and compiled code threats, including new vulnerabilities 136. The knowledge base engine 122 may receive such outputs from the threat intelligence 126 by way of a standard "heartbeat" program, i.e., a program configured to transfer file records from a remote database on a flexibly periodic basis. The threat intelligence 126 may be filtered or otherwise correlated by use of artificial intelligence 134. A discussion and illustration in FIG. 2 describe these various processes of the sandbox engine 110 in greater detail below.

[0091] Once the crypto asset smart contract 106 has been scanned using static analysis, dynamic analysis or both, results pass to a risk scoring engine 112 to assess the vulnerabilities of the digital asset. The risk scoring engine 112 determines a risk score for the asset at least partly based on comparisons to records in a crypto asset smart contract database 124, which comprises a database of global statistics received from a smart contract data source 128. By checking a crypto asset smart contract 106, i.e., digital asset, against a population of similar digital assets from the crypto asset smart contract database 124, the risk scoring engine 112 generates statistical vulnerability data compared to the digital asset's vulnerabilities. Beyond these comparisons, the risk scoring engine 112 generates a risk score for the digital asset by additionally making use of a regression model, an offline training engine, an online prediction engine as described below and illustrated in FIG. 2.

[0092] A scheduler 114 determines the periodicity of ingesting digital assets according to at least date and time. The time resolution (e.g., to the minute, second, and so on) and periodicity of the scanning schedule may be determined by the financial institution 102 or the continuous vulnerability management system 104 or the two in agreement. Additionally, a policy manager 116 specifies rules and standards determining the actions to be taken with vulnerabilities and directs such decisions to the security control system 120 to be passed as a control signal 132 back to the financial institution 102. Once vulnerabilities have been identified via the risk score, reporting 118 of the vulnerabilities passes to a security control system 120, which then alerts the security controller client 108 at the financial institution 102 as to whether the assessed digital asset or digital assets may be subject to suspension, quarantining/resuming, or other security measures.

[0093] Referring to FIG. 2, sandbox engine units 200 according to various embodiments are described and illustrated. As noted above, the sandbox engine continually scans for

vulnerabilities in digital assets 202 such as in a smart contract crypto asset 212. As also noted, the sandbox engine may perform two types of scanning on differing stages of code: a static analysis as shown by the static analysis unit 204 path and a dynamic analysis as shown by the dynamic analysis unit 206 path. The digital asset source code may be transferred to the sandbox engine using a file transfer protocol (e.g., FTP or similar) or other electronic file copying means known to those skilled in the art for moving files or data between computer devices. Static code scanning 216 takes the digital assets and searches them for anomalies and errors, e.g., deprecated syntax that might cause vulnerability, in their static text code via searching and pattern matching techniques. The anomalies and errors may represent software routines intended to cause harm when the source code is compiled and executed. While they may be otherwise difficult to identify, particularly in source code form, a vulnerability knowledge base 210 containing known source code threats aids the static code scanning 216 of the static analysis unit 204 in making these comparisons. Differences and/or similarities from these comparisons comprise a static result 228.

[0094] In certain embodiments, a dynamic analysis unit 206 also processes the smart contract crypto asset 212, but in this instance compiles the source code from the smart contract crypto asset 212 and executes the compiled code 238 in a secure environment. As shown, the first step of dynamic analysis for a smart contract crypto asset 212 is to compile 222 its source code. Compiling source code may typically be performed by a special program (i.e., a "compiler") that turns statements written in a high-level programming language into machine language usable by a computer processor, a well-known process by those skilled in the art. Once compiled, the compiled code 238 undergoes execution in virtual machine 224. A virtual machine environment, or virtual environment, allows for dynamic execution of the compiled code 238 in a secure environment, e.g., without potentially infecting a wider computing environment if executing (e.g., running) the program creates malware, implants viruses and so on. Symbolic execution is also available within the virtual machine secure environment allowing searches for known dynamic patterns such as reentrancy and fuzz testing using malformed inputs. As noted above, the vulnerability knowledge base 210 comprises results (e.g., from the smart contract data source 128) against which the execution in virtual machine 224 may be checked. In one embodiment, the vulnerability knowledge base 210 comprises a reentrancy vulnerability pattern 240, which may be detected by the dynamic analysis unit of the sandbox engine. In another embodiment, the reentrancy vulnerability pattern 240 may be

described as key-value pairs 242. Differences or similarities from this comparison comprise the dynamic result 230, i.e., the output of the dynamic analysis unit 206.

[0095] As shown, an online prediction engine 226 evaluates both the static result 228 and dynamic result 230 of the static analysis and dynamic analysis, respectively, of the digital assets to generate a risk score 232 for each vulnerable asset 234. Details of how the online prediction engine 226 determines this risk score 232, including use of a regression model and an online prediction engine, are described below and illustrated in FIG. 5.

[0096] For the purposes of this description, the online prediction engine 226 passes the risk score 232 to the statistical analysis unit 208 of the sandbox engine. The statistical analysis unit 208 comprises a database of global statistics 214 against which the risk score 232 from the online prediction engine 226 compares against similar digital assets 220. Once the standing of digital assets has been assessed with respect to each digital asset's vulnerabilities within the population of similar digital assets 220, the audit report generation unit 218 of the sandbox engine generates results based on the statistical vulnerability data 236 of the statistical analysis unit 208. FIG. 10 below describes further details and illustrations of the audit report generation unit.

Sandbox engine

[0097] A sandbox engine may comprise development hardware, software, data, tools, interfaces, and policies necessary for starting an enterprise machine learning practice. A sandbox comprises an isolated testing environment enabling users to run programs or execute files without affecting the application, system or platform on which they run. Software developers use sandboxes to test new programming code. Cybersecurity professionals use sandboxes to test potentially malicious software. A sandbox may also be used to safely execute malicious code to avoid harming the device on which the code is running, the network or other connected devices. Using a sandbox to detect malware offers an additional layer of protection against security threats, such as stealthy attacks and exploits that use zero-day vulnerabilities. Exemplary sandbox applications include PDFs and other documents that include executable code, so an application like Adobe Reader runs PDF files in a sandbox, which stops them from escaping the PDF viewer and interfering with the rest of the computer. Microsoft Office also has a sandbox mode to stop unsafe macros from tampering with a system.

Smart contract

[0098] A smart contract may comprise a computer protocol intended to digitally facilitate, verify, or enforce the negotiation or performance of a contract and allow for the performance of credible transactions without third parties. Smart contracts may be used to automate the uninterrupted transfer of crypto-tokens between users, according to agreed-upon conditions. A smart contract may be created in a computer text editor able to create source code and be written in any language consistent with blockchain technology, e.g., solidity, go, java, javascript, C++ , SQL.

Symbolic execution

[0099] Symbolic execution may comprise a software testing technique useful to aid in the generation of test data and in proving program quality. A goal of symbolic execution in this context of software testing is to explore as many different program path conditions as possible in a given amount of time, and for each path to both generate a set of concrete input values exercising that path, and check for the presence of various kinds of errors including security vulnerabilities. One idea behind symbolic execution is using symbolic values, instead of concrete data values, as input and to represent the values of program variables as symbolic expressions over the symbolic input values. As a result, the output values computed by a program are expressed as a function of the symbolic input values. In software testing, symbolic execution is used to generate a test input for each execution path of a program. The ability to generate concrete test inputs is one of the major strengths of symbolic execution: from a test generation perspective, it may allow the creation of high-coverage test suites, while from a vulnerability perspective, it may provide developers with a concrete input that triggers the bug, which can be used to confirm and debug the error independently of the symbolic execution tool that generated it.

Reentrancy

[0100] Reentrancy describes a computer program or routine written so that the same copy in memory can be shared by multiple users. Reentrant code is commonly required in operating systems and applications intended to be shared in multi-use systems. Within a code structure, a reentrant function may be used by more than one task concurrently without fear of data corruption and, most importantly, may be interrupted at any time and resumed at a later time without loss of data. Reentrant functions either use local variables or protect their data when global variables are used.

Fuzz testing

[0101] Fuzzing, fuzzing testing, or fuzz testing comprises an automated software testing technique that provides invalid, unexpected, or random data as inputs to a computer program. The program is then monitored for exceptions such as crashes, failing built-in code assertions, or potential memory leaks. Malformed inputs comprise a common input to fuzz testing.

Malformed inputs

[0102] As noted above, malformed inputs may be deployed in fuzz testing to provide known dynamic patterns. If the failure-inducing input is large and mostly malformed, a developer may have difficulty understanding the cause of testing failure. Given the failure-inducing input, an automated minimization tool may be deployed to remove as many input bytes as possible while still reproducing the original error.

Key-value pairs

[0103] Key-value pairs indicate a set of two linked data items: a key, which is a unique identifier for some item of data, and the value, which is either the data that is identified or a pointer to the location of that data. Key-value pairs are frequently used in lookup tables, hash tables and configuration files.

[0104] Referring to FIG. 3, a dynamic execution in virtual machine 300 according to an embodiment is illustrated. As noted above, dynamic execution takes place within the dynamic analysis unit of sandbox engine, i.e., receiving digital asset source code, compiling the source code into compiled code (a.k.a., bytecode 308), and executing the compiled code within a virtual machine secure environment. Execution comprises dynamically executing the bytecode control flow graph based on the vulnerability knowledge base. When executed in the virtual machine, compiled smart contract bytecode 308 instructions, with the code, storage, balance, and so on, are recorded in the blockchain.

[0105] Detection of smart contract vulnerabilities in the dynamic analysis unit of the sandbox engine comprises detecting patterns in a dynamic control flow graph, as shown in FIG. 3, and recorded in the knowledge base engine 122. The flow represents dynamic execution of the smart contract compiled bytecode 308 in the virtual machine and monitoring by the dynamic analysis unit of the sandbox engine for the storage, memory and function call checking against the knowledge base engine 122 for known vulnerability patterns. An exemplary set of bytecode 308 instructions (e.g., locations and functions with corresponding memory locations) comprise

contract A 302 with corresponding calls into contract B 304. As each call 312 (e.g., for storage, memory, or other computing function) is encountered, a path condition 310 for the execution is obtained before the call 312 executes. The path condition 310 is then checked with updated storage values 306 to determine if the path condition 310 remains valid, i.e., if the call 312 may be executed again with the updated storage values 306. If true, a vulnerability is flagged, as the call 312 receiver may re-execute the call 312 before finishing it.

Path condition

[0106] A path condition comprises a sequence of bytecode 308 commands interpreted by a computer in an environment for executing compiled code. In the context of dynamic execution, e.g., monitoring for each call 312 in a smart contract, encountering a call 312 triggers the retrieval of the executable path condition 310 before the call 312 itself executes, to check if the same path condition with updated storage values 306 still holds (i.e., if the call can be executed again).

[0107] Referring to FIG. 4, machine learning 400 according to an embodiment is illustrated. As previously described, a sandbox engine 402 accepts crypto asset smart contracts in the form of digital assets, e.g., from a financial institution. As shown, the smart contract crypto assets may be either legitimate digital assets or vulnerable digital assets by definition -- that is, they are necessarily one or the other. Each digital asset type therefore describes a sandbox engine 402 input: either a legitimate smart contract crypto asset 410 or a vulnerable smart contract crypto asset 412. By verifying the legitimate smart contract crypto asset 410 against a known crypto assets database and the vulnerable smart contract crypto asset 412 against the knowledge base engine 122, both data sets constantly learn about new vulnerability behavior patterns and threat intelligence, the sandbox engine 402 verifies and outputs both known legitimate digital assets 406 and known vulnerable digital assets 408.

[0108] In one embodiment feature vectors may be formulated as inputs to machine learning model training 404 for both known legitimate digital assets 406 and known vulnerable digital assets 408. The feature vectors may comprise any aspect of the smart contract containing multiple numerical elements, e.g., frequency of occurrence of textual terms such as functions and variables, and taken collectively serve as provided training data for the machine learning model training 404.

[0109] Machine learning model training 404 comprises collecting provided training data, or the aforementioned feature vectors from the known legitimate digital assets 406 and known vulnerable digital assets 408 and using a neural network as discussed below and illustrated in FIG. 15 to train an offline training engine to estimate inputs of a machine learning algorithm.

Machine learning algorithm

[0110] A machine learning algorithm 420, in the context of evaluating smart contracts as digital assets, may be an off-line training engine, i.e., with the ability to estimate all parameters of a regression model with the ability to ingest feature vectors as model inputs and apply a regression model 414 to improve its results over time. Inputs comprising a set of numeric features may be conveniently described by a feature vector. In this context, both the known legitimate digital assets 406 and known vulnerable digital assets 408 may be described by a feature vector, an n-dimensional vector of numerical features that represent each set of assets. Many algorithms in machine learning require a numerical representation of objects, since such representations facilitate processing and statistical analysis. An example of reaching a two-way classification from a feature vector includes calculating the scalar product between the feature vector and a vector of weights, comparing the result with a threshold, and deciding the class based on the comparison. Algorithms for classification from a feature vector include nearest neighbor classification, neural networks, and statistical techniques such as Bayesian approaches. As described below and illustrated in FIG. 15, a neural network is utilized for the machine learning algorithm 420 for classification from the named feature vectors.

Regression model

[0111] A regression model 414 comprises algorithms and analysis. Regression analysis is a set of statistical processes for estimating the relationships among variables. It includes many techniques for modeling and analyzing several variables, when the focus is on the relationship between a dependent variable and one or more independent variables. More specifically, regression analysis helps one understand how the typical value of the dependent variable changes when any one of the independent variables is varied, while the other independent variables are held fixed. Regression analysis additionally estimates the conditional expectation of a dependent variable given the independent variables – that is, the average value of the dependent variable when the independent variables are fixed.

[0112] Regression algorithms are named for their continuous outputs, meaning they may have any value within a range, in this case from the machine learning algorithm within the machine learning model training 404. In this case, the regression model take the form of a function, i.e., $f(X)$, where $f()$ is the regression model and X is input feature vectors from either the known legitimate digital assets 406 or known vulnerable digital assets 408, which serve as provided training data 416 and provided training data 418, respectively. The regression model may be linear or non-linear.

[0113] Referring to FIG. 5, a risk scoring engine 500 according to an embodiment is illustrated. As previously described, a sandbox engine 502 accepts a smart contract crypto asset 506 in the form of digital assets, e.g., from a financial institution. As shown, the smart contract crypto asset 506 may be either a legitimate digital asset or a vulnerable digital asset. Either of these digital asset types therefore displays as a sandbox engine 502 input. By verifying the legitimate digital asset against a known crypto assets database and the vulnerable digital asset against the knowledge base engine 122, both data sets constantly learning about new vulnerability behavior patterns and threat intelligence, the sandbox engine 502 uploads its sandbox engine outcomes 514 to at least one database of digital assets 512 one or more of known vulnerabilities 516, known legitimate digital assets 518, and/or known vulnerable digital assets 520. As noted above, the offline training engine 522, using the feature vectors from both the known legitimate digital assets 518 and known vulnerable digital assets 520, estimates parameters of the machine learning algorithm 524 using the regression model 508 as described.

[0114] The risk scoring engine 504 comprises the online training, using the regression model 508 trained in the offline process, to predict the score 510 for a specific smart contract. The risk scoring engine 504 thereby includes the regression model 508, the offline training engine 522, and access to at least one database of digital assets 512 comprised of known vulnerabilities 516, known legitimate digital assets 518, and known vulnerable digital assets 520.

[0115] To create a score 510 from the risk scoring engine 504, a quality scale (0-100) may be defined with 100 as the highest quality smart contract and 0 (zero) is the most vulnerable. With at least one database of digital assets 512 comprising known legitimate digital assets 518 (score of 100) and known vulnerable digital assets 520 (score of 0), the regression model 508 may be trained by interpolating the score from 0 to 100. The mathematics of the linear regression model 508 may also apply to a non-linear regression model. The function $Y_i = a + b_1 * X_{s1} + b_2$

* X_{d_i} is used to compute Y_i , the output quality score, where X_s represents the output of the static analysis, (e.g. deprecated syntax findings), X_d represents the output of the dynamic analysis, (e.g. Reentrancy vulnerability identification), and a , b_1 , and b_2 are model parameters in the regression model learned from the dataset. The counter "i" for Y , X_s , and X_d , indicates the "i-th" data point in the training dataset, where i ranges from 1 to n . The parameters a , b_1 , and b_2 may be computed using the following linear minimization algorithm:

$$\text{Find } \min_{\alpha, \beta} Q(\alpha, \beta), \quad \text{for } Q(\alpha, \beta) = \sum_{i=1}^n \hat{\epsilon}_i^2 = \sum_{i=1}^n (y_i - \alpha - \beta x_i)^2.$$

Equation 1

[0116] Referring to FIG. 6, a digital asset trading exchange 600 is illustrated. A security control system exists to at least in part, identify, protect, detect, respond to, and recover crypto assets in financial institutions. Exemplary institutions may be Crypto Exchange, Crypto Hedge Fund, Crypto Custody, and so forth. The security control system may automatically suspend trading for a vulnerable crypto asset, triggered by a user's placed trade order, in a crypto asset trading exchange. If a digital asset is determined to be vulnerable by the sandbox, only that digital asset is suspended from trading, e.g., other legitimate digital assets continue to function.

[0117] Embodiments of the security control system include logic, as described below, to identify digital assets, scan digital assets for vulnerabilities and identify those digital assets as vulnerable digital assets. The logic may further protect digital assets including the ability to suspend trading of vulnerable digital assets, and respond to a trade on the digital asset trading exchange 600. The trade may include vulnerable digital assets and legitimate digital assets, where the vulnerable digital assets are suspended and the legitimate digital assets are allowed to proceed.

[0118] In an embodiment, the logic begins at block start 602. At block 604, a user places a trade order of asset_A at a digital asset trading exchange. At block 606, logic queries a sandbox security controller to analyze the order at decision block 608 and determine whether asset_A is

vulnerable. If asset_A is found to be vulnerable, it joins a set of vulnerable digital assets 630 and an action 628 is taken. Asset_A is now an identified vulnerable asset 610 and at block 612, asset_A is suspended from trading. At block 614 asset_A is removed from trading pairs 626 (e.g., for asset_A, USD/Bitcoin). At block 616, an operation team is notified to remediate, and the logic ends at done block 622.

[0119] Returning to decision block 608, if asset_A is determined not to be not vulnerable (e.g., it is legitimate and part of the set of legitimate digital assets 624), at block 618 trading completes, and asset_A is settled on the blockchain at block 620 before completing at done block 622.

[0120] Referring to FIG. 7, a proactive scanning 700 process is described. In a similar manner to the logic in the digital asset trading exchange 600 description above in FIG. 6, proactive scanning 700 represents another mode of a security control system. In this embodiment, the security control system continuously scans all digital assets in the trading pairs 718 for vulnerabilities. Once found, the digital asset is removed from the trading pairs 718 while others continue to function normally.

[0121] The logic to proactively scan digital assets begins at start 702. At decision block 704, if the scanning process is true (e.g., active), for an asset X in trading pairs 718 at block 706, the sandbox security control system is queried at block 708. If asset X is determined not to be vulnerable at decision block 710, the logic returns to check the next trading pair. If asset X is vulnerable at decision block 710, it is proactively removed from trading pairs 718 at block 712, and an operation team is notified to remediate the asset at block 714. At this point the logic is complete at done block 716.

[0122] Referring to FIG. 8, an exemplary smart contract 800 is illustrated. In this example, the syntax of the smart contract program language 802 is identified as solidity in the first line of code. Other notable features of this smart contract include the smart contract name 804 ("AnChainAiVulnerableCryptoAsset") and the body of smart contract code 806.

[0123] Referring to FIG. 9, a directed acyclic graph 900 is shown that may use the protective features described above. A directed acyclic directed acyclic graph 900 (DAG) is a finite directed graph with no directed cycles, i.e., consisting of finitely many nodes 902 and edges 904 with each edge directed from one node to another, such that there is no way to start at any node and follow a consistently-directed sequence of edges that eventually loops back to the initial node again. Equivalently, a directed acyclic graph 900 is a directed graph that has a

topological ordering, a sequence of the nodes such that every edge is directed from earlier to later in the sequence.

[0124] A topological ordering of a directed acyclic graph: every edge goes from earlier in the ordering (upper left) to later in the ordering (lower right). A directed graph is defined as acyclic if and only if it has a topological ordering.

[0125] Referring to FIG. 10, an exemplary audit report 1000 is illustrated. As described above and illustrated in FIG. 2, the risk score serves as an input to the statistical analysis unit of the sandbox engine. With the sandbox engine fully automated, the quality scores of all smart contracts may be computed at scale, while the vulnerability of an individual asset may also be isolated. A visualization, such as a heatmap 1006 drawn from heatmap data 1010 aids understanding of a crypto asset's standing in the entire digital asset population. In one embodiment, the heatmap 1006 illustrates a statistical ranking of Lines of Code vs. Risk score, with each cell indicating a cluster of similar smart contracts. The standing 1008 indicates the current smart contract crypto asset in scope and the code line 1002 indicates the vulnerability location. By using a standard clustering algorithm on sandbox behaviors, the most similar smart contracts may be identified.

[0126] The vulnerability auditing report, as shown, may be generated automatically, and the Risk / Compliance / Operation team in this case may be notified immediately, e.g., in the audit report 1000, the crypto asset indicated has a score of 0 (out of 100), indicating a severe level 1004 of vulnerability. This statistical standing among all 50,000 smart contracts visualized in the heatmap 1006, indicates it as among the worst smart contracts.

[0127] Referring to FIG. 11, in block 1102 a method 1100 of protecting digital assets scans digital assets for vulnerabilities using a sandbox engine. In block 1104, the method 1100 acquires information related to new vulnerabilities and threat intelligence using a knowledge base engine including artificial intelligence. In block 1106, the method 1100 analyzes sandbox engine outcomes and assigns a risk score to each vulnerable asset using a risk scoring engine including machine learning. In block 1108, the method 1100 takes an action on an identified vulnerable asset based on the risk score assigned to the identified vulnerable asset using a security control system.

[0128] Referring to FIG. 12, a blockchain transaction process 1200 is an ever-growing set of data blocks. Each block records a collection of transactions. Blockchains distribute these

transactions across a group of computers. Each computer maintains its own copy of the blockchain transactions.

[0129] A blockchain is a continuously growing list of records, called blocks, which are linked and secured using cryptography. Each block typically comprises a cryptographic hash of the previous block, a timestamp, and transaction data. By design, a blockchain is resistant to modification of the data. Blockchains may implement an open, distributed ledger that can record transactions between two parties efficiently and in a verifiable and permanent way.

[0130] A blockchain is typically managed by multiple parties collectively adhering to a protocol for inter-node communication and validating new blocks. Once recorded, the data in any given block cannot be altered retroactively without alteration of all subsequent blocks, which requires consensus among the operators.

[0131] Cryptography involving mathematical methods of keeping data secret and proving identity is utilized when recording transactions. One digital key ensures only an owner for example on a transaction requesting device 1202, can enter a transaction to the blockchain involving their assets, and another digital key lets other parties confirm it really was the owner who added the transaction.

[0132] Blockchain is resistant to tampering or other changes by utilizing a cryptographic technique called the hash. Hashing reduces data to a sequence of seemingly random characters - for example, the hash of the phrase "the quick brown fox" is "9ECB36561341D18EB65484E833EFEA61EDC74B84CF5E6AE1B81C63533E25FC8F" using a hash method called SHA-256. Tweaking just one letter in the phrase produces a completely different hash, and you can't go backward to figure out the original data from the hash.

[0133] With blockchain, hashes are linked together so any minute change is immediately visible, not just for the block housing it but for all other blocks added later. With red flags that big for changes that small, auditing becomes easier.

[0134] FIG. 13 illustrates an exemplary blockchain formation 1300. The mainchain 1302 (M blocks) comprises the longest series of blocks from the start block 1306 (S block) to the current block. Orphan blocks 1304 (O blocks) exist outside of the main chain.

[0135] Blocks hold batches of valid transactions that are hashed and encoded, for example into a Merkle tree. Each block includes the cryptographic hash of the prior block in the blockchain formation 1300, linking the two. The linked blocks form a

chain. This iterative process confirms the integrity of the previous block, all the way back to the original start block 1306.

[0136] Sometimes separate blocks can be produced concurrently, creating a temporary fork. In addition to a secure hash-based history, the blockchain formation 1300 has a specified algorithm for scoring different versions of the history so that one with a higher value can be selected over others. Blocks not selected for inclusion in the mainchain 1302 are called orphan blocks 1304. Peers supporting the blockchain formation 1300 have different versions of the history from time to time. They keep only the highest-scoring version of the blockchain formation 1300 known to them. Whenever a peer receives a higher-scoring version (usually the old version with a single new block added) they extend or overwrite their local version of the blockchain formation 1300 and retransmit the improvement to their peers. There is never an absolute guarantee that any particular entry will remain in the best version of the history forever. Because blockchains are typically built to add the score of new blocks onto old blocks and because there are incentives to work only on extending with new blocks rather than overwriting old blocks, the probability of an entry becoming superseded goes down exponentially as more blocks are built on top of it, eventually becoming very low. For example, in a blockchain using the proof-of-work system, the chain with the most cumulative proof-of-work is always considered the valid one by the network. There are a number of methods that can be used to demonstrate a sufficient level of computation. Within a blockchain the computation is carried out redundantly rather than in the traditional segregated and parallel manner.

[0137] FIG. 14 illustrates an embodiment of an irreversible transaction blockchain 1400. The blockchain 1400 is a sequence of digitally signed transactions (transaction 1 1402, transaction 2 1404, and transaction 3 1410 etc.). Each transaction includes the current owners public key (block 1 owner public key 1406, block 2 owner public key 1412, and block 3 owner public key 1416 respectively) and the previous owner's signature (O(0) signature 1408, O(1) signature 1414, and O(2) signature 1418) which are generated using a hash function. The owner of a transaction can examine each previous transaction to verify the chain of ownership. Unlike traditional check endorsements, the transactions in the blockchain 1400 are irreversible, which mitigates fraud.

[0138] In some configurations, a machine learning algorithm may be implemented as a deep learning neural network. FIG. 15 illustrates a basic deep neural network 1500 is based on a

collection of connected units or nodes called artificial neurons which loosely model the neurons in a biological brain. Each connection, like the synapses in a biological brain, can transmit a signal from one artificial neuron to another. An artificial neuron that receives a signal can process it and then signal additional artificial neurons connected to it.

[0139] In common implementations, the signal at a connection between artificial neurons is a real number, and the output of each artificial neuron is computed by some non-linear function (the activation function) of the sum of its inputs. The connections between artificial neurons are called 'edges' or axons. Artificial neurons and edges typically have a weight that adjusts as learning proceeds. The weight increases or decreases the strength of the signal at a connection. Artificial neurons may have a threshold (trigger threshold) such that the signal is only sent if the aggregate signal crosses that threshold. Typically, artificial neurons are aggregated into layers. Different layers may perform different kinds of transformations on their inputs. Signals travel from the first layer (the input layer 1502), to the last layer (the output layer 1506), possibly after traversing one or more intermediate layers, called hidden layers 1504.

[0140] Referring to FIG. 16, an artificial neuron 1600 receiving inputs from predecessor neurons consists of the following components:

- inputs x_i ;
- weights w_i applied to the inputs;
- an optional threshold (b), which stays fixed unless changed by a learning function; and
- an activation function 1602 that computes the output from the previous neuron inputs and threshold, if any.

[0141] An input neuron has no predecessor but serves as input interface for the whole network. Similarly an output neuron has no successor and thus serves as output interface of the whole network.

[0142] The network includes connections, each connection transferring the output of a neuron in one layer to the input of a neuron in a next layer. Each connection carries an input x and is assigned a weight w .

[0143] The activation function 1602 often has the form of a sum of products of the weighted values of the inputs of the predecessor neurons.

[0144] The learning rule is a rule or an algorithm which modifies the parameters of the neural network, in order for a given input to the network to produce a favored output. This learning

process typically involves modifying the weights and thresholds of the neurons and connections within the network.

[0145] FIG. 17 illustrates one example of a system architecture and data processing device that may be used to implement one or more illustrative aspects described herein in a standalone and/or networked environment. Various network nodes including data server 1710, web server 1706, computer 1704, and laptop 1702 may be interconnected via a wide area network 1708 (WAN), such as the internet. Other networks may also or alternatively be used, including private intranets, corporate networks, LANs, metropolitan area networks (MANs) wireless networks, personal networks (PANs), and the like. Network 1708 is for illustration purposes and may be replaced with fewer or additional computer networks. A local area network (LAN) may have one or more of any known LAN topology and may use one or more of a variety of different protocols, such as ethernet. Devices including data server 1710, web server 1706, computer 1704, laptop 1702 and other devices (not shown) may be connected to one or more of the networks via twisted pair wires, coaxial cable, fiber optics, radio waves or other communication media.

[0146] The term "network" as used herein and depicted in the drawings refers not only to systems in which remote storage devices are coupled together via one or more communication paths, but also to stand-alone devices that may be coupled, from time to time, to such systems that have storage capability. Consequently, the term "network" includes not only a "physical network" but also a "content network," which is comprised of the data--attributable to a single entity--which resides across all physical networks.

[0147] The components may include data server 1710, web server 1706, and client computer 1704, laptop 1702. Data server 1710 provides overall access, control and administration of databases and control software for performing one or more illustrative aspects described herein. Data server 1710 may be connected to web server 1706 through which users interact with and obtain data as requested. Alternatively, data server 1710 may act as a web server itself and be directly connected to the internet. Data server 1710 may be connected to web server 1706 through the network 1708 (e.g., the internet), via direct or indirect connection, or via some other network. Users may interact with the data server 1710 using remote computer 1704, laptop 1702, e.g., using a web browser to connect to the data server 1710 via one or more externally exposed web sites hosted by web server 1706. Client computer 1704, laptop 1702 may be used in concert with data server 1710 to access data stored therein, or may be used for

other purposes. For example, from client computer 1704, a user may access web server 1706 using an internet browser, as is known in the art, or by executing a software application that communicates with web server 1706 and/or data server 1710 over a computer network (such as the internet).

[0148] Servers and applications may be combined on the same physical machines, and retain separate virtual or logical addresses, or may reside on separate physical machines. FIG. 17 illustrates just one example of a network architecture that may be used, and those of skill in the art will appreciate that the specific network architecture and data processing devices used may vary, and are secondary to the functionality that they provide, as further described herein. For example, services provided by web server 1706 and data server 1710 may be combined on a single server.

[0149] Each component including data server 1710, web server 1706, computer 1704, laptop 1702 may be any type of known computer, server, or data processing device. Data server 1710, e.g., may include a processor 1712 controlling overall operation of the data server 1710. Data server 1710 may further include RAM 1716, ROM 1718, network interface 1714, input/output interfaces 1720 (e.g., keyboard, mouse, display, printer, etc.), and memory 1722. Input/output interfaces 1720 may include a variety of interface units and drives for reading, writing, displaying, and/or printing data or files. Memory 1722 may further store operating system software 1724 for controlling overall operation of the data server 1710, control logic 1726 for instructing data server 110 to perform aspects described herein, and other application software 1728 providing secondary, support, and/or other functionality which may or may not be used in conjunction with aspects described herein. The control logic may also be referred to herein as the data server software control logic 1726. Functionality of the data server software may refer to operations or decisions made automatically based on rules coded into the control logic, made manually by a user providing input into the system, and/or a combination of automatic processing based on user input (e.g., queries, data updates, etc.).

[0150] memory 1722 may also store data used in performance of one or more aspects described herein, including a first database 1730 and a second database 1732. In some embodiments, the first database may include the second database (e.g., as a separate table, report, etc.). That is, the information can be stored in a single database, or separated into different logical, virtual, or physical databases, depending on system design. Web server 1706, computer 1704, laptop 1702 may have similar or different architecture as described with

respect to data server 1710. Those of skill in the art will appreciate that the functionality of data server 1710 (or web server 1706, computer 1704, laptop 1702) as described herein may be spread across multiple data processing devices, for example, to distribute processing load across multiple computers, to segregate transactions based on geographic location, user access level, quality of service (QoS), etc.

[0151] One or more aspects may be embodied in computer-usable or readable data and/or computer-executable instructions, such as in one or more program modules, executed by one or more computers or other devices as described herein. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types when executed by a processor in a computer or other device. The modules may be written in a source code programming language that is subsequently compiled for execution, or may be written in a scripting language such as (but not limited to) HTML or XML. The computer executable instructions may be stored on a computer readable medium such as a nonvolatile storage device. Any suitable computer readable storage media may be utilized, including hard disks, CD-ROMs, optical storage devices, magnetic storage devices, and/or any combination thereof. In addition, various transmission (non-storage) media representing data or events as described herein may be transferred between a source and a destination in the form of electromagnetic waves traveling through signal-conducting media such as metal wires, optical fibers, and/or wireless transmission media (e.g., air and/or space). various aspects described herein may be embodied as a method, a data processing system, or a computer program product. Therefore, various functionalities may be embodied in whole or in part in software, firmware and/or hardware or hardware equivalents such as integrated circuits, field programmable gate arrays (FPGA), and the like. Particular data structures may be used to more effectively implement one or more aspects described herein, and such data structures are contemplated within the scope of computer executable instructions and computer-usable data described herein.

[0152] The term “configured to” is not intended to mean “configurable to.” An unprogrammed FPGA, for example, would not be considered to be “configured to” perform some specific function, although it may be “configurable to” perform that function after programming.

[0153] Reciting in the appended claims that a structure is “configured to” perform one or more tasks is expressly intended not to invoke 35 U.S.C. § 112(f) for that claim

element. Accordingly, claims in this application that do not otherwise include the “means for” [performing a function] construct should not be interpreted under 35 U.S.C § 112(f).

[0154] As used herein, the term “based on” is used to describe one or more factors that affect a determination. This term does not foreclose the possibility that additional factors may affect the determination. That is, a determination may be solely based on specified factors or based on the specified factors as well as other, unspecified factors. Consider the phrase “determine A based on B.” This phrase specifies that B is a factor that is used to determine A or that affects the determination of A. This phrase does not foreclose that the determination of A may also be based on some other factor, such as C. This phrase is also intended to cover an embodiment in which A is determined based solely on B. As used herein, the phrase “based on” is synonymous with the phrase “based at least in part on.”

[0155] As used herein, the phrase “in response to” describes one or more factors that trigger an effect. This phrase does not foreclose the possibility that additional factors may affect or otherwise trigger the effect. That is, an effect may be solely in response to those factors, or may be in response to the specified factors as well as other, unspecified factors. Consider the phrase “perform A in response to B.” This phrase specifies that B is a factor that triggers the performance of A. This phrase does not foreclose that performing A may also be in response to some other factor, such as C. This phrase is also intended to cover an embodiment in which A is performed solely in response to B.

[0156] As used herein, the terms “first,” “second,” etc. are used as labels for nouns that they precede, and do not imply any type of ordering (e.g., spatial, temporal, logical, etc.), unless stated otherwise. For example, in a register file having eight registers, the terms “first register” and “second register” can be used to refer to any two of the eight registers, and not, for example, just logical registers 0 and 1.

[0157] When used in the claims, the term “or” is used as an inclusive or and not as an exclusive or. For example, the phrase “at least one of x, y, or z” means any one of x, y, and z, as well as any combination thereof.

LISTING OF DRAWING ELEMENTS

100 continuous vulnerability management system

102 financial institution

104 continuous vulnerability management system

106 crypto asset smart contract

108 security controller client

110 sandbox engine

112 risk scoring engine

114 scheduler

116 policy manager

118 reporting

120 security control system

122 knowledge base engine

124 crypto asset smart contract database

126 threat intelligence

128 smart contract data source

130 digital asset

132 control signal

134 artificial intelligence

136 new vulnerabilities

200 sandbox engine units

202 digital assets

204 static analysis unit

206 dynamic analysis unit

208 statistical analysis unit

210 vulnerability knowledge base

212 smart contract crypto asset

214 database of global statistics

216 static code scanning

218 audit report generation unit

220 similar digital assets

222 compile

224 execution in virtual machine

226 online prediction engine

228 static result

230 dynamic result

232 risk score

234 vulnerable asset

236 statistical vulnerability data

238 compiled code

240 reentrancy vulnerability pattern

242 key-value pairs

300 dynamic execution in virtual machine

302 contract A

304 contract B

306 storage values

308 bytecode

310 path condition

312 call

400 machine learning

402 sandbox engine

404 machine learning model training

406 known legitimate digital assets

408 known vulnerable digital assets

410 legitimate smart contract crypto asset

412 vulnerable smart contract crypto asset

414 regression model

416 provided training data

418 provided training data

420 machine learning algorithm

500 risk scoring engine

502 sandbox engine

504 risk scoring engine

506 smart contract crypto asset

508 regression model

510 score

512 at least one database of digital assets

514 sandbox engine outcomes

516 known vulnerabilities

518 known legitimate digital assets

520 known vulnerable digital assets

522 offline training engine

524 machine learning algorithm

600 digital asset trading exchange

602 start

604 block

606 block

608 decision block

610 identified vulnerable asset

612 block

614 block

616 block

618 block

620 block

622 done block

624 legitimate digital assets

626 trading pairs

628 action

630 vulnerable digital assets

700 proactive scanning

702 start

704 decision block

706 block

708 block

710 decision block

712 block

714 block

716 done block

718 trading pairs

800 smart contract

802 smart contract program language

804 smart contract name

806 smart contract code

900 directed acyclic graph

902 nodes

904 edges

1000 audit report

1002 code line

1004 severe level

1006 heatmap

1008 standing

1010 heatmap data

1100 method

1102 block

1104 block

1106 block

1108 block

1200 blockchain transaction process

1202 transaction requesting device

1300 blockchain formation

1302 mainchain

1304 orphan blocks

1306 start block

1400 blockchain

1402 transaction 1

1404 transaction 2

1406 block 1 owner public key

1408 $O(0)$ signature

1410 transaction 3

1412 block 2 owner public key

1414 $O(1)$ signature

1416 block 3 owner public key

1418 $O(2)$ signature

1500 basic deep neural network

1502 input layer

1504 hidden layers

1506 output layer

1600 artificial neuron

1602 activation function

1702 laptop

1704 computer

1706 web server

1708 network

1710 data server

1712 processor

1714 network interface

1716 RAM

1718 ROM

1720 input/output interfaces

1722 memory

1724 operating system software

1726 control logic

1728 other application software

1730 first database

1732 second database