(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
01 November 2018 (01.11.2018)

WIPO | PCT

(10) International Publication Number
WO 2018/199912 A1

(51) International Patent Classification:
G05B 19/418 (2006.01) *G06Q 50/04* (2012.01)
G06Q 10/06 (2012.01) *H04L 12/24* (2006.01)

(21) International Application Number:
PCT/US2017/029239

(22) International Filing Date:
25 April 2017 (25.04.2017)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant: **SIEMENS AKTIENGESELLSCHAFT** [DE/DE]; Werner-von-Siemens-Straße 1, 80333 München (DE).

(72) Inventors: **PFLEGER DE AGUIAR, Leandro**; 34 Hawk Road, Lawrenceville, New Jersey 08648 (US). **WORONKA, Stefan**; Brockenblick 23, 38271 Baddeckenstedt (DE).

(74) Agent: **RASHIDI-YAZD, Seyed Kaveh E.**; Siemens Corporation- Intellectual Property Dept., 3501 Quadrangle Blvd. Ste. 230, Orlando, Florida 32817 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC,

(54) Title: PLANT DESIGN BASED INCREMENTAL ASSET DISCOVERY ARCHITECTURE, METHOD AND PROTOCOL

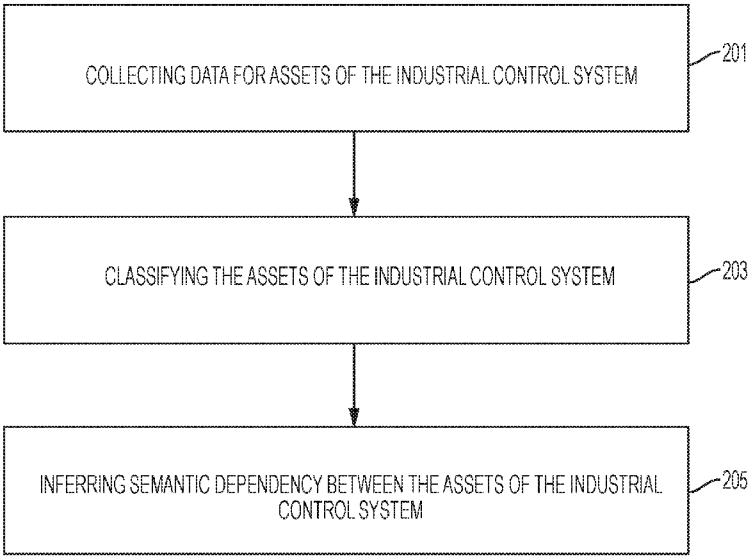

FIG. 2

(57) Abstract: The present embodiments relate to performing an asset inventory for an industrial control system. By way of introduction, the present embodiments described below include systems and methods for asset discovery, asset characterization and semantic analysis. Incremental asset discovery in an industrial control system is provided using semantic analysis of the underlying production process using semantic models, template ontologies, and data analytics (e.g., based on metadata, network data, process data, etc.). Using the semantic analysis, different asset discovery methods are used to target individual assets of the industrial control systems to identify and baseline the assets. An accurate asset inventory is provided without risking disruption of the production process, and the asset inventory includes assets from highly segmented and isolated networks by adopting a combination of discovery methods, intelligent agents, and data analytics.

WO 2018/199912 A1

[Continued on next page]

SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

PLANT DESIGN BASED INCREMENTAL ASSET DISCOVERY ARCHITECTURE, METHOD AND PROTOCOL

Background

[0001] There is increased interest by cyber attackers in attacking critical infrastructure by compromising industrial automation and control systems. Due to vertical integration of the production systems and horizontal integration of the value chain, industrial control systems (ICS) and industrial control networks are often directly or indirectly connected to information technology (IT) networks, such as local office and plant networks and the Internet. This vertical integration may provide an opportunity for cyber attackers to exploit those networks to take advantage of known and newly discovered infrastructure vulnerabilities.

[0002] Unlike computers and other computing devices running on conventional IT networks, ICS products (e.g., programmable logic controllers (PLCs), distributed control systems (DCS), motion controllers, supervisory control and data acquisition (SCADA) systems, and human-machine interfaces (HMIs)) are designed for process control functionalities without intrinsic consideration of cybersecurity. Most process control system networks, including multiple PLCs, DCS devices, motion controllers, SCADA devices and HMIs, are also integrated without consideration for potential cyber threats.

[0003] To successfully manage risks in an industrial environment, each industrial asset is typically identified in an inventory and baselined by an individual profile in order to detect unauthorized deviations from the baseline. As such, asset management is

acknowledged by the security community as one of the most critical processes to ensure cyber security risks are appropriately controlled. For example, being able to identify and/or enumerate assets is typically the first step in defining stages of a cyber kill chain and preventing exploitation of a target. Asset identification and classification is also considered by many international security frameworks (e.g., ISO 27001 and IEC 62443) as a key security control to be implemented as part of any organization's security program.

[0004] Standard IT asset inventory tools identify IT assets by deploying a range of automated methods to discover devices from a central server or network node. These methods typically include active connection methods through, for example, simple network management protocol (SNMP) and Windows Management Instrumentation (WMI) to identify devices supporting the protocols. Standard IT asset inventory tools also include special software agents for Windows and Linux based systems that typically run with local administrative privileges to collect and export detailed inventory information regarding the Windows and Linux based systems. Existing tools directed to asset inventory for operations technology (OT) and industrial control systems often utilize similar active connection methods and special software agents, and also include passive detection options to avoid service disruption. For example, passive devices monitor and collect network traffic in order in order to make inferences about the

assets that originate the traffic. The information gathered may be used to make inferences about vulnerabilities on these protocols and devices.

[0005] Technical challenges currently prevent automatically performing asset inventories of industrial control systems. For example, industrial systems are highly heterogeneous and operate in highly segmented networks typically configured such that centralized automated scanning tools are prevented from successfully reaching all devices. In addition, transient assets, such as maintenance or contractor laptops, are often not detected during the inventorying. Further, centralized automated scanning tools are often severely limited in the ability to read and interpret the content of third party system configuration files. In many cases, the operator of an industrial automation and control system does not have any documentation or knowledge of the implemented third-party assets. As a result, asset inventorying and baselining is often underperformed or simply not executed by asset owners, leading to security risks.

[0006] Existing solutions also lack the ability to distinguish criticality levels of each asset in relation to the production process the assets support. Criticality levels are important for vulnerability based risk assessments where worst case impact levels are determined in order to quantify the overall risk that a given asset is exposed to and to aid the risk decisions by asset owners. Existing network management tools that automate network discovery with topology mapping algorithms are unable to

automatically derive the criticality level of an asset and the relationship of the asset to a production process.

Summary

[0007] The present embodiments relate to performing an asset inventory for an industrial control system. By way of introduction, the present embodiments described below include systems and methods for asset discovery, asset characterization and production process semantic analysis. Incremental asset discovery in an industrial control system is provided using semantic analysis of the underlying production process with semantic models, template ontologies, and data analytics (e.g., based on metadata, network data, process data, etc.). Using the semantic analysis, different asset discovery methods are used to target individual assets of the industrial control systems to identify and baseline the assets. An accurate asset inventory is provided without risking disruption of the production process, and the asset inventory includes assets from highly segmented and isolated networks by adopting a combination of discovery methods, intelligent agents, and data analytics.

[0008] In a first aspect, a method of asset discovery for an industrial control system is provided. The method includes characterizing production processes, production zones and automation packages of the industrial control system. Based on the characterization, the method includes instrumenting assets of the industrial control system for incremental asset discovery and performing incremental asset discovery. The incremental asset discovery includes collecting data for the assets of the industrial control system using one or more instrumentation options, classifying the assets of the

industrial control system based on the collected data and inferring semantic dependency between the assets of the industrial control system.

[0009] In a second aspect, a system for industrial asset discovery is provided. The system includes a plurality of asset discovery methods configured to collect asset and process information from an industrial control system and a server configured to store the collected asset and process information received from the plurality of asset discovery methods. The server is further configured to classify the assets of the industrial control system based on the collected asset and process information received from the server. The server is also configured to infer semantic dependency between the assets of the industrial control system based on the collected asset and process information received from the server.

[0010] In a third aspect, a method for automatically identifying and baselining industrial assets is provided. The method includes characterizing major production processes, production zones, and automation packages of the industrial control system. The method also includes assigning discovery methods for an incremental asset discovery of assets of the industrial control system, executing the incremental asset discovery of the assets of the industrial control system and generating an output of the incremental asset discovery. The output includes displaying an attack tree for the assets of the industrial control system. The output also includes detailed asset information

(e.g., Internet protocol (IP) addresses, operating systems, system types, etc.), dependency relationships of the production process to a given asset, etc.

[0011] The present invention is defined by the following claims, and nothing in this section should be taken as a limitation on those claims. Further aspects and advantages of the invention are discussed below in conjunction with the preferred embodiments and may be later claimed independently or in combination.

Brief Description of the Drawings

[0012] The components and the figures are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the embodiments. Moreover, in the figures, like reference numerals designate corresponding parts throughout the different views.

[0013] Figure 1 illustrates a flowchart diagram of an embodiment of a method of asset discovery for an industrial control system.

[0014] Figure 2 illustrates a flowchart diagram of an embodiment of a method of incremental asset discovery.

[0015] Figure 3 illustrates a flowchart diagram of another embodiment of a method of automatic incremental asset discovery.

[0016] Figure 4 illustrates an example of manually importing a production zone on from a pre-compiled list of production units into a process map.

[0017] Figure 5 illustrates an example of a process map for reporting an asset inventory.

[0018] Figure 6 illustrates an embodiment of a system for asset discovery for an industrial control system.

[0019] Figure 7 illustrates and example of an asset discovery agent configured as a relay agent.

[0020] Figure 8 illustrates a passive industrial perimeter assessment device configured to provide a temporary overlay asset inventory network.

Detailed Description of Exemplary Embodiments

[0021] The present embodiments provide systems and methods for identifying and baselining assets of an industrial control system. A workflow begins by understanding a production process, capturing the essence of how the production process runs in real life. A combination of hardware and software agents are then deployed based on the production process to perform asset discovery. The workflow uses process related data sources and asset information gathered by the hardware and software agents to semantically infer additional information about the assets, such as asset categories, relationships between assets, and potential risks associated with the assets.

[0022] The present embodiments obviate one or more of the drawbacks or limitations in existing information technology (IT) and operations technology (OT) asset inventory tools preventing the existing tools from being successful in identifying and

baselining assets of an industrial control system. For example, the asset inventory process begins by understanding the production process. By understanding the production process prior to identifying production assets, an intrinsically complex and heterogeneous OT environment is understood before connecting and deploying any automatic mapping tools that could negatively impact the process. By understanding the production process prior to performing asset discovery, efficiency and accuracy of the asset inventory process may be improved.

[0023] The process information gathered through understanding the production process is leveraged during asset discovery, allowing for improved risk assessment for OT systems including segregated and/or isolated non-IT related devices. For example, performing asset discovery is guided by the production process, such as by deploying a combination of automatic, manual and user-assisted automatic asset inventory tools. Asset inventory information is automatically assembled using data collected through distributed asset inventory tools, such as from hardware and/or software sensors strategically deployed at different points of the OT network. The deployed asset inventory tools identify and baseline assets in highly segmented and isolated networks by adopting one or more of the hardware and software discovery methods.

[0024] The asset information gathered by the deployed asset inventory tools is processed to infer connections and relationships between assets and production zones/cells. Using the inferences from the asset information, risk assessment may be

facilitated. For example, using the asset inventory, asset baselines and semantic relationships between assets, an attack-tree may be derived with data analytics to facilitate risk assessment. An accurate industrial asset inventory is thus provided including asset information from highly segmented and isolated networks, facilitating risk assessment by inferring semantic relationships and performing data analytics.

[0025] Figure 1 illustrates a flowchart diagram of an embodiment of a method of asset discovery for an industrial control system. The method illustrated in figure 1 provides accurate and efficient industrial asset discovery and analysis. The method is implemented by the system of figure 6 (discussed below) and/or a different system. Additional, different or fewer acts may be provided. For example, the acts of figure 1 may be combined with the acts illustrated in figure 2 (discussed below). The method is provided in the order shown. Other orders may be provided and/or acts may be repeated. For example, acts 101-105 may be repeated for a plurality of assets, or repeated for each production cell/zone. Further, the acts 101-105 may be performed concurrently as parallel acts. Each of acts 101-105 may include sub-acts, such as acts 201-205 illustrated in figure 2 (discussed below).

[0026] At act 101, production processes, production zones and automation packages of the industrial control system are characterized. The main steps of the production process, production cells/zones and/or automation packages may be mapped and characterized based on information derived from multiple sources. For example,

multiple sources of information may be processed by the system, such a data from a process historian (e.g., plant information management system (PIMS)), control system project files (e.g., process control engineering files, totally integrated automation (TIA) portal engineering .S7P project files, etc.), configuration files, direct device reading (e.g., PLC memory reading, running PLC control logic, etc.), etc. Using this information, the production process is mapped into smaller, more manageable production cells/zones. Additionally, the system may be configured to receive user inputs regarding the process, such as manually entering information into the system to establish cells/zones or by importing precompiled lists of common production steps.

[0027] Each of the cells/zones of the production process are mapped, and characteristics of the cells/zones are inferred by the system based on available information. For example, the system age, automation package vendors, typical assets for a process or production step and typical interconnections between process steps are inferred and mapped. By understanding the production process prior to identifying assets, an intrinsically complex and heterogeneous industrial control system is understood and mapped before the connection of any automatic mapping tools that could potentially generate negative impacts to the industrial control system, such as when the system performs a continuous production process or a process with limited downtime.

[0028] At act 103, assets of the industrial control system are instrumented for incremental asset discovery. The industrial control system is instrumented based on the mapped and characterized production process. For example, the mapped production process includes a SCADA server coupled to unknown assets of the industrial control system (e.g., various industrial devices, such as controllers, PLCs, etc.). Based on the mapped production process, the system instruments to SCADA server to automatically discover the assets coupled the SCADA server. Instrumenting the industrial control system includes deploying asset inventory tools (e.g., asset discovery agents) for the different assets of the industrial control system. The industrial control system network is instrumented to perform an incremental industrial asset discovery using a combination of deployed hardware and software discovery agents (e.g., quickly and easily deployable sensors) configured to employ a mixture of active, passive, intrusive, and/or non-intrusive discovery methods. The instrumentation is installed at the control cell where the control equipment resides (e.g., a ruggedized industrial personal computer (IPC) installed in a network of a control cell).

[0029] For example, a combination of hardware and software asset discovery agents are deployed for each of the cells/zones of the production process. Using information gathered for each of the cells/zones, specific hardware and software asset discovery agents are selected and deployed to perform asset discovery and to baseline the discovered assets. Hardware asset discovery agents are connected to an asset and/or

are installed in the network of the industrial control system to monitor and collect communication, activity, and other information about one or more assets of the industrial control system. Software asset discovery agents are installed on and executed by an asset to monitor and collect communications, activity, and other information about one or more assets.

[0030] In addition to deploying a combination of hardware and software discovery agents, the deployed discovery agents may be a combination of active and passive discovery agents. Active discovery agents actively interact with an asset, such as by sending a command to the asset and receiving a response. Based upon the response to a command, the active discovery agent gathers information for classifying and baselining assets. Passive discovery agents do not actively interact with the asset. The passive discovery agents monitor action by the asset, such as by passively monitoring communication and activity of the asset (e.g., read-only). For example, based on the type, frequency, and other characteristics of communication and activity, information gathered by a passive discovery agent may be used for classifying and baselining assets. Further, classifying assets may be performed based on data analytics (e.g., asset signatures) pushed from the cloud database.

[0031] Additionally, the deployed discovery agents may be a combination of intrusive and non-intrusive discovery agents. Non-intrusive discovery agents do not impact the production process as the agents monitor assets and/or collect information. For

example, hardware discovery agents configured to monitor communication in a read-only configuration may not impact the production process as data is collected. Intrusive discovery agents impact the production process as discovery agents monitor and/or collect information about one or more assets of the industrial control system. For example, for a software agent installed and executed by an asset, the software agent may utilize computational bandwidth of the installed asset, impacting the performance of the installed asset (e.g., by slowing down the asset). In some embodiments, the intrusive discovery agents are deployed under strict monitoring for potential adverse impacts to the production process, and may be configured to prevent the adverse effects during production. For example, the collective network bandwidth used by all agents during active asset inventory operations may be limited to a given amount of Mbps and/or during a given timeframe where dependency on systems is less critical to the production. Local impact by deployed software agents (e.g., overall CPU percent usage, physical memory consumption, etc.) may also be controlled. Such constraints may be learned and derived automatically (e.g., via machine learning) based on processing configured data sources. Alternatively, some assets cannot be discovered without adversely impacting or stopping the production process. In this example, the intrusive discovery agent may only be deployed during an idle phase of the production process.

[0032] Instrumenting assets of the industrial control system may also include deploying temporary or permanent networking devices for asset discovery and monitoring in the industrial control system. For example, various cells/zones of the industrial control system may be isolated or segregated, such as air-gapped or firewalled assets or networks of assets. To collect information regarding the isolated or segregated assets, an overlay network is deployed for connecting discovery agents collecting information for assets within the isolated or segregated cells/zones of the industrial control system. For example, a hardware device is connected to an asset or network device for monitoring the asset or network. The overlay network facilitates one-way (e.g., unidirectional) communication from the control system over the overlay network. The hardware device is also configured to transmit collected data to a server using the overlay network. An overlay network is a network (e.g., a wireless network) deployed in addition to and without affecting existing networks of the industrial control system. Using the overlay network, the asset information may be provided for risk assessment, such as uploaded to a central server for analysis.

[0033] In an example, if one-way connectivity to the overlay network is not feasible, mobile USB or Ethernet sniffing dongles (e.g., USB storage with the one-way effect) may be used. In this example, the mobile collector is installed (e.g., plugged-in) for a period of time (e.g., two weeks, etc.) and removed for manual upload.

[0034] At act 105, incremental asset discovery is performed. As discussed above, the discovery agents are deployed and monitored for potential adverse impacts to the production process. Using the deployed discovery agents, asset information is gathered and used to classify the discovered assets by type. Using the information collected, and based on the mapped network topology and inferred production zones, process steps, and automation packages, semantic dependency between assets within the production process may be inferred. By classifying the discovered assets and inferring semantic dependency, risk assessment may be accurately performed.

[0035] Figure 2 illustrates a flowchart diagram of an embodiment of a method of incremental asset discovery. For example, the method illustrated in figure 2 may perform incremental asset discovery executed in act 105 of figure 1. The method is implemented by the system of figure 6 (discussed below) and/or a different system. Additional, different or fewer acts may be provided. The method is provided in the order shown. Other orders may be provided and/or acts may be repeated. For example, act 205 may be performed before act 203, or acts 201-205 may be repeated for a plurality of assets, and/or repeated for each production cell/zone. Further, the acts 201-205 may be performed concurrently as parallel acts. Act 205 may be omitted in some embodiments.

[0036] At act 201, data is collected for assets of the industrial control system. Data is collected using the deployed instrumentation of act 103 (e.g., by each asset discovery

agent). The data is stored locally or uploaded to a central server by each discovery agent. For example, data for the assets of the industrial control system is collected by monitoring asset activity and communication between assets. Additionally, data is collected by scanning for assets coupled to each software or hardware discovery agent. In an example, a software agent is installed on a SCADA server coupled to multiple PLCs and other assets. The software agent scans and collects data for each PLC and uploads the data to a central server configured to store the data in an asset database. In another example, a hardware agent is installed for monitoring communication from an asset. The hardware agent is deployed for a period of time (e.g., two weeks) and the communication data is automatically or manually uploaded to the asset database stored on the central server. Further, a discovery agent may send commands and other communications to an asset and receiving a response or monitoring data from the asset. Different discovery assets are deployed to gather data for each asset of the industrial control system.

[0037] At act 203, the assets of the industrial control system are classified. Using the data collected for each asset, an inference is made classifying each asset of the industrial control system. For example, by monitoring communication and activity of the assets, inferences are automatically generated classifying the assets. Based on the type, duration, frequency, etc. of each communication and/or activity, the type of asset is determined. Additionally, using asset data collected in act 201, an asset age, asset

vendor (e.g., for third-party assets), typical characteristics, and other asset characteristics is also inferred for each asset. Classifying the assets may include identifying the vendor of a given system and the role of the asset in the network. Additional data may be provided by a user (e.g., semi-supervised machine learning) and user provided asset signatures may be used to automatically classify similar devices in future.

[0038] Characterizing each asset may also include baselining each asset to determine activities and communications that are typical of the asset during different production processes, etc. For example, using the data collected for each asset, abnormal activity by one or more assets may be easily identified (e.g., for risk assessment and/or deriving an attack-tree) based on deviation from the recorded norm. Abnormal activities may also be used then to detect subsequent changes made to these inventoried assets.

[0039] In an embodiment, classifying assets also includes supervised or semi-supervised machine learning based on user input. For example, similar unclassified assets (e.g., third-party assets) are classified based on a user input. For example, similar unclassified assets are identified from monitored communications and other collected data. An unclassified asset representative of the unclassified assets is presented to a user for manual classification. Based on a user input, the presented asset and similar unclassified assets may be classified and/or characterized. Alternatively, a representative asset is classified and presented to a user to confirm the classification.

Based on a user input, the representative and similar third-party assets are classified and/or characterized. The related third-party assets may be classified utilizing the user input and further inferences from collected data to classify each asset. For example, based on monitored communication between the assets, related dissimilar assets are identified and classified using a user input related to one of the assets. For example, classifying one asset of a third-party automation platform is provided based on a user input of one asset of a plurality of assets making up the automation platform.

[0040] At act 205, semantic dependency is inferred between the assets of the industrial control system. For example, using asset data collected in act 201, asset classifications and characteristics provided in act 203, and/or the characterized processes, production zones and automation packages of the industrial control system of act 101, semantic relationships and an associated criticality is inferred for each identified assed. For example, based on the frequency, average packet size, traffic direction, and protocols uses, roles of each asset (e.g., a human-machine interface (HMI), PLC, etc.) may be identified. Based on the discovered topology of connected assets, additional computer systems that belong to the same control cell are identified, and based on the grouping of such devices, groups of control cells are mapped. The activity observed at the application layer of the communication (e.g., process control and monitoring commands) and asset and user activity may be used to infer process sequences. For example, in a discrete manufacturing process, such as a car assembling

line, inferences are automatically generated about the process sequence based on the observed network communications between nodes and content exchanged. Data provenance techniques may be used to track raw data being transmitted and transformed across components (e.g., raw sensor data is feed into a PLC, read from the SCADA server, displayed at the HMI, and recorded at the process historian). Such sequence reveals dependency is inferred to derive a hierarchy of components in a growing level of criticality. Semantic relationships allow the system to infer interdependency and other relationships between assets. Supervised or semi-supervised machine learning may also be performed with user input to provide additional semantic relationships, or to confirm inferences of semantic relationships between assets.

[0041] The semantic information may be used to identify abnormal activity by one or more assets (e.g., for risk assessment and/or deriving an attack-tree). For example, abnormal activity may include one or more assets performing activities during an idle phase of the process for the asset(s), one or more assets idling during an active phase of the process for the asset(s), higher than normal computational loads for one or more of the assets, communications or other signals from one or more assets that are inconsistent with the process, one or more of the assets being unresponsive to a communication from another asset, etc. Other abnormal activities may also be

detected, such as when an asset is active or inactive when other related assets are active or are not active.

[0042] Further, using the semantic relationships, an associated criticality may be inferred. For example, some assets of the industrial control system have a larger impact on the production process than others. In this example, a SCADA server connecting multiple PLCs and other assets to the industrial control system may have a larger impact on the production process than any one of the PLCs or other assets of the system. As such, a higher criticality level may be inferred for the SCADA server than for each of the assets connected to the SCADA server. Additionally, the SCADA server may also be at a higher risk because of being coupled to and able to communicate with other devices accessible to an outside network.

[0043] Figure 3 illustrates a flowchart diagram of an embodiment of a method of automatic incremental asset discovery for OT systems, industrial control systems and other systems that include segregated and non-IT related devices. For example, the main acts of the automated asset inventory process are provided for identifying and baselining industrial assets. The method is implemented by the system of figure 6 (discussed below) and/or a different system.

[0044] At 301, the automated asset inventory process performs an understand act. In the understand act, data from multiple sources is processed to derive and infer production steps, production zones, and to classify different automation packages. For

example, automation packages to be classified include a Siemens PCS7 control system used to control a water pump station, an Allen Bradley ControlLogix control system used to control a gas turbine, or any other group of automation hardware and software provided and commissioned as a "turn-key" package. For example, the multiple sources of information may include process historian information, control system project files, configuration files, direct device information (e.g., PLC memory reading), etc. Other types of information may be processed to understand the production process. Using the processed information, production steps and zones are automatically inferred by a semantic mapper of a software platform. Production steps and zones may also be suggested to a user for confirmation, or manually entered by the user.

[0045] Production steps and zones may be imported manually based on a pre-defined list of standard production units. For example, figure 4 illustrates manually importing a production zone from a pre-compiled list of production units into a process map. As illustrated in figure 4, a production zone A may be imported into a semantic mapper using drag and drop functionality from the available process blocks. In this example, the semantic mapper adopts industry specific vertically provided ontology packets as process blocks. Ontology packets are industry specific (e.g., Oil & Gas, Metals & Mining, etc.) production units that define common production steps and assets based on the industry and/or process. The semantic mapper processes and transforms the multiple sources of data into data elements and namespaces in process map.

[0046] The process map includes a combination of automatically and manually defined process zones. In addition to semantically mapping the production process, the understand act also includes mapping network topologies of the industrial control system. The various networks employed by the production process are understood and mapped. For example, the production process may utilize connections to the Internet, a local intranet, segregated networks, firewalled networks and/or isolated networks. Other types of networks may be mapped. The understand act identifies the different types of networks and maps each topology to understand connectivity and interdependency characteristics of the production zones.

[0047] The understand act may additionally classify automation packages of the industrial control system and different asset types. For example, based on the industry specific production units, common automation packages may be mapped. For example, in a production zone of an Oil & Gas process, a batch or continuous automation process may include a package of hardware and software specifically designed and implemented for the zone. The automation package may include controllers, PLCs, sensors, motors, valves, actuators, etc. Further, by understanding automation packages and other characteristics of the production zones, the understand act may also identify and classify asset types (e.g., switches, PLCs, etc.) and characteristics (e.g., new or legacy, model, vendor, etc.) of each zone prior to performing asset discovery, simplifying the asset discovery process and providing additional information to draw from to infer

semantic relationships. For example, based on the classified asset types and characteristics, asset discovery methods may be assigned based on the types and characteristics to discovery unknown assets in the production zones and automation packages.

[0048] At 303, the automated asset inventory process performs a plan act. Using the process map and information gathered during the understand act, incremental asset discovery is planned for each production zone, automation package, etc. Different discovery methods (e.g., asset discovery agents) are suggested and assigned based on information and inferences about existing controls systems and assets (e.g., the system age, automation package vendor, etc.). Further, centralized passive discovery may be used in the planning act to provide additional insights about assets to be mapped. The centralized passive discovery gathers asset information without deploying additional discovery agents and discovery methods. The planning act uses this asset information to better plan for the deployment of additional discovery agents and discovery methods for undiscovered and non-baselined assets.

[0049] In an embodiment, each mapped production zone or network zone is assigned a combination of different asset discovery agents. For example, zones that do not tolerate active intrusive discovery methods (e.g., due to risks of running such active methods on legacy equipment) are configured with passive methods only, and zones with newer control equipment are be configured with more intrusive methods (e.g.,

providing more detailed results). The central server controls which method will be triggered at which time and under specific constraints for each zone. For example, the asset discovery relay agents may be hardware, software or a combination thereof that gathers asset information and relays the information to a central information server. Each asset discovery agent is also specified as having one or more discovery methods, including active, passive, intrusive, non-intrusive and/or a combination thereof. Any combination of the asset discovery agents and asset discovery methods may be specified based on the inferences derived for the production process, production and network zones, automation packages, learned asset information for different assets, etc. The specified asset discovery agents may be confirmed by a user during planning, and the asset inventory plan may be refined and confirmed.

[0050] The asset discovery agents and discovery methods are assigned to collect asset and process information from an industrial control system with minimal impact to the production process. For example, in some applications, specific hardware or software discovery agents may have a greater or lesser impact on the production process, and are selected accordingly. For example, a discovery agent employing a passive discovery method (e.g., passively monitoring network communications, etc.) will impact a process less than a discovery agent employing an active discovery method (e.g., actively sending a command to an asset, etc.). A discovery agent employing a non-intrusive discovery method (e.g., passively scanning asset attributes, versions, etc.) will

likewise impact a process more than an intrusive discovery method (e.g., actively scanning an asset resulting in a fault state of the asset). Where asset discovery may use different types of agents, the agent with the lesser impact is selected.

[0051] The planning act may designate a schedule for the incremental asset discovery. For example, an active and/or intrusive discovery agent may be scheduled to perform asset discovery during idle phases of the production process. The inventory process may also be scheduled for different zones according to the production schedule, maintenance schedule, planned outages, etc.

[0052] Further, some asset discovery agents are specified as customized intelligent discovery agents. For example, in addition to being customized to a specific process or network zone, some discovery agents monitor the process to minimize computational overhead of the discovery agent during a computationally intensive phase of the process. Intelligent discovery agents are equipped with real-time, heart-beat monitoring functionality to quickly detect a disruption or other impact on any asset during incremental asset discovery. For example, some performance indicators for monitoring asset disruption include processor queue length, percent CPU process time, throughput, memory load, log size, overall CPU load, cache size, etc. Other performance indicators may be selected and used. Intelligent discovery agents are configured with upper limits of allowed system and/or network overhead for asset discovery, and may prevent the asset discovery from disrupting or negatively impacting the production process.

[0053] At 305, the automated asset inventory process performs an execute act. In the execute act, the automatic data collection hardware or software asset discovery agents are deployed. The asset discovery agents are deployed based on the planned discovery methods assigned based on the process map. The discovery agents are deployed to collect asset and process information for the different process and network zones of the industrial control system. Assets of the industrial control system and assets of the underlying automation network topology are identified and baselined, suggesting asset categories, semantic relationships and associated criticality for each identified assed. The asset discovery agents utilize existing networks to upload asset information to a central server. Some asset discovery agents are manually deployed isolated and segregated assets, such as by deploying a hardware agent establishing a temporary overlay network for air-gapped assets.

[0054] During deployment, performance of the assets are monitored, such as by monitoring the actual performance impact of the asset inventory process in relation to estimated and/or allowed impact determined in the planning act. The plurality of asset discovery agents are configured to monitor the computational overhead of the asset discovery agents to minimize disruption or other impacts on the assets during the production process. The asset discovery agents are configured with self-modifying behavior. For example, the asset discovery agents may be configured with upper limits (e.g., thresholds) of allowed system and/or network overhead for the asset discovery.

The asset discovery agents may be configured with real-time, heart-beat monitoring functionality to quickly detect any availability disruptions or other impacts on the asset (e.g., such as monitoring rata rate, throughput, memory load, log size, CPU load, cache size, etc.). Based on heart-beat monitoring, the asset discovery agents may be modified to prevent disruption of the process, such as by reducing computational bandwidth utilization, pausing asset discovery, etc. An accurate industrial asset inventory is performed without risking disruption of the production process.

[0055] The execute act may also display and coordinate user activities during the automated asset inventory process. For example, the system may provide a display allowing the user to confirm correlations, inferences and asset criticalities automatically suggested by the system to confirm semantic relationships between industrial and network assets, and between different production zones. Supervised and semi-supervised machine learning may also be performed based on user input to confirm assumptions and inferences of semantic relationships between assets. Correlations, inferences and asset criticalities are automatically suggested utilizing data analytics. By analyzing the timing and frequency of asset activities and communications, inferences are made including categorizing and grouping assets, inferring connections between assets and determining which assets are used in each production step.

[0056] For example, asset criticality in a process may also be inferred utilizing data analytics based on frequency of use count in an entropy analysis from system access.

The entropy analysis may plot binary and other data acquired from an asset (e.g., binary data from a SCADA server, PLC, etc.). The data may not be designed to be plotted. Hidden keys, such as increased and/or abnormal activity or other abnormalities in the data, are reflected in plotted binary data. Using the hidden keys, inferences are made as to the criticality and/or relationship of an asset to other assets of a production process. Alternatively, the hidden keys may identify a time frame of activity to inquire additional information from a user to characterize assets of the industrial control system.

[0057] The automated asset inventory platform may also utilize data from outside production processes (e.g., from disparate client facilities) within the same industry to generate ontology packets based on patterns in communications between the facilities. For example, patterns arise regarding communications for particular vendors, for particular types of assets and/or for particular production processes. Inferences drawn from different facilities may be used to characterize communications and activities of assets within an industry. The patterns are recognized and utilized in providing standardized ontology packets for use by a user in a particular industry, leveraging platform experience from the disparate client facilities.

[0058] At act 307, the automated asset inventory process performs a report act. The automatic and manually collected asset information from individual production zones is assembled and combined to infer, generate and/or confirm the overall communication topology (e.g., a map and/or graph). For example, networked discovery assets relay

information to a central server and offline discovery assets are connected to a network for uploading to the central server. The asset information may also be uploaded to a cloud-based server accessible by a cloud-based software platform.

[0059] Figure 5 illustrates an embodiment of a process map for reporting an asset inventory. The process map may be presented as an asset inventory operation control dashboard. The process map presents a network topology mapping, including segregated, isolated and firewalled assets and networks of assets, to a user. Optionally, the reassembling and combining asset information supports additional analytics using the available semantic and mapped asset characteristics and a zone-data-matrix generated from collected process data at different network points. The additional analytics provides data paths throughout the different network subnets and tracking (e.g., based on the observation of repeated data streams in different networks or zones).

[0060] The asset inventory results are thus consolidated for display to a user. Inferred semantic relationships between assets of the industrial control system may be presented to a user and confirmed via a user input. For example, physical connections and dependency relationships are inferred based on collected data and network traffic. The semantic relationships are displayed for confirmation by a user (user assisted). In addition to providing the output of the incremental asset discovery, the reporting act may also provide an attack tree for the assets of the industrial control system. The

reporting act may compute and display an attack tree depicting points of entry and vulnerable assets in the event of a cyber-attack. The attack-tree is derived from the asset inventory information utilizing data analytics. The attack-tree and asset baselines are provided to aid future risk assessments.

[0061] Figure 6 illustrates system an embodiment of a system for asset discovery for an industrial control system. For example, system 600 includes a plurality of networked, segregated, isolated and/or firewalled hardware and software components for performing incremental asset discovery according to the methods of figure 1, figure 2, figure 3 or another method. The system 600 includes a computing platform 601 coupled to a server 605 and workstation 607 via the internet 603. Computing platform 601 may be implemented as a cloud computing platform, with a cloud server 601A in addition to or in place of server 605. Alternatively, computing platform 601 may be implemented locally as part of server 605 and workstation 607. The computing platform 601 is configured to store asset characterizations (e.g., signatures) to cloud server 601 for use in classifying other assets in the same or similar industrial control systems.

[0062] The system 600 includes a plurality of asset discovery agents 611, 613, 615, 617 and 621. The plurality of asset discovery agents are configured to collect asset and process information from an industrial control system. The asset discovery agents may be hardware agents 613, software agents 611, or combination hardware and software agents (e.g., combining PC agent 615 with SW agent 611). The software asset discovery

agents 611 are configured for installation on a process device or a networking device, such as using software or firmware running on the device. The software asset discovery agents 611 may also be pre-loaded on the process device or networking device. The asset discovery agents may include PC agents 615 implemented as an industrial computing environment (e.g., industrial personal computer (IPC), ruggedized person computer (PC), industrial server computer, industrial controller, etc.). The PC agents 615 may include software agents 611 deployed on the industrial computing environment. Each of the asset discovery agents are coupled to or executed by an asset (e.g., a process device or a networking device) to collect asset information.

[0063] The asset discovery agents may be configured as relay agents. For example, relay agents are configured to upload asset and process information to one or more central servers, such as server 605 via intranet 609, overlay network 619 or another network. The networks, such as intranet 609 and overlay network 619, may be provided with known or future networking technology (e.g., Ethernet, wireless, cellular, optical, NFC, etc.). Alternatively, the asset discovery agent may be a standalone agent 617 (e.g., not connected to networks 609, 619, etc.) configured to capture and store information for manual upload to server 605.

[0064] Figure 7 illustrates and example of an asset discovery agent configured as a relay agent. For example, during the discovery process, specially configured asset discovery agents (e.g., software, hardware or firmware based) act as relays for exporting

data from assets and/or other discovery agents. For example, software asset discovery agent 611 installed on SCADA server 711 is implemented on a SCADA server connected to one or more assets, such as PLC device 723. The software relay agent 611 installed on SCADA server 711 is configured to retrieve information from one type of network (e.g., a layer 1/2 network card) and to transmit the information over another type of network (e.g., a layer 2/3 network). As such, the software relay agent 611 installed on SCADA server 711 captures PLC information and utilizes switch 709 to transmit the information to central server 705. The relay agents allow for various networks (e.g., complex mesh network topologies) to reach the central asset inventory consolidating node (e.g., server 705) via the secure connection (e.g., switch 709 and intranet 609). Relay agents are used for instrumenting non-routable control devices have to be discovered. For example, non-routable control devices cannot be reached directly from the central server via a network, requiring a relay device between two disparate networks. Using relay agents may reduce or minimize the number of deployed agents necessary. For example, local relay agents discover devices in a network segment and retransmit the device information to the main server.

[0065] In an embodiment, the software agent 611 is executed as an autonomous installing software agent. For example, referring to figure 7, the software agent 611 is autonomously installed on the SCADA server 711. The autonomous installing software agent executes self-replicating code configured to analyze a process or networking

device for code installation as software agent 611. Based on the analysis, the software agent code is customized and installed on the device. The software agent code is executed to gather asset and process information, then the code is uninstalled returning the device to an uninstalled state.

[0066] The asset discovery autonomous installing agent (e.g., a melting agent) employs intelligent agent droppers (e.g., installers) to obtain and deploy a non-disruptive deployment strategy. The executable software has characteristics and actions often found in self-replicating code or malware, without the detrimental effects of malware code, in order to achieve silent, seamless, and zero-downtime installation, execution, and removal of the agent code. For example, the code may include features for delaying the inclusion of the code, executing the code, run-time loading and/or generating additional code at a time when the device is not performing critical actions and/or has computational bandwidth free for the code operations. The self-mutating and/or replicating code may install the agent code on neighbor peer devices. Peer devices that do not support the software agent code may be noted and tracked for future asset discovery actions. The executable software code includes performance control that is configured to avoid undesired performance effects on the target devices, including constant monitoring OS interactions and critical system variables and including sleep calls intertwined with execution operations. The executable software includes secure unpacking and encryption, generating and transmitted data encrypted at the

source. The executable software also includes clean and self-software removal. For example, after the data is collected, the agent optionally self-removes the installed program code automatically.

[0067] In an embodiment, some asset discovery relay agents are provided as hardware agents. For example, some assets of the industrial control system are identified and baselined by installing and deploying a hardware agent to gather asset information. The hardware agents may be customized intelligent discovery agents for a specific production or network zone. A temporary overlay asset inventory network may be deployed. Using a temporary overlay asset inventory network, a temporary secure wireless network allows for quick and seamless asset inventory without changing the target network topology.

[0068] Figure 8 illustrates a passive industrial perimeter assessment (PIPA) secure device configured to provide a temporary overlay asset inventory network. For example, a PIPA secure device 821 is configured as a passive scanner and collector of network and process data. The PIPA passive scanning device is configured to provide a network connection to a central server, such as using Zigbee wireless communication, facilitating the asset inventory temporary overlay network. In this embodiment, the PIPA passive scanning device is provided with an embedded opto-coupler 825 to create the short-term, temporary overlay asset inventory network (e.g., using wireless, ZigBee, a personal area network (PAN) wireless communication standard, etc.) for use only during

the incremental asset discovery and inventory. Deployed PIPA secure devices 821 at different points throughout the facility collect asset inventory relevant information and communicate in a wireless mesh network topology via network 619. An internal, uni-directional communication device of the PIPA secure device 821 avoids potential contamination in case malware is present. The PIPA secure device 821 may be toggled as "read-only" or "write-only," depending on the desired functionality.

[0069] As discussed above, a central server 605 is configured to store the collected asset and process information received from the plurality of asset discovery agents. The server 605 may receive the data according to an asset discovery and reporting protocol. Communications from agent nodes 611, 613, 615, 617 to the central server 605 via Internet 609 and from agent 621 via the temporary overlay network 619 are provided using a specific protocol to facilitate the asset discovery and reporting. For example, the asset discovery and reporting protocol includes source and destination authentication and content encryption with asymmetric key encryption. Authentication and encryption provide for secure communication of asset information to the server 605 and/or the computing platform 601. Each agent may self-select a transmission mode (e.g., connection oriented or user datagram protocol (UDP)). For example, agents may self-select available paths and/or ports of the network, based on locally collected information from a note communicating on the network. Using an inference based on

the available communication paths, smart agents change communication ports and transmission modes in a hop-to-hop discovery mode.

[0070] The workstation 807 is coupled to server 605 and/or computing platform 601 and is configured to display a portal interface to the user. The computing platform 601 is configured to classify the assets of the industrial control system. For example, a classifier 601C classifies assets based on the collected asset and process information received from the server 605. An analyzer 601D infers semantic dependency between the assets of the industrial control system based on the collected asset and process information received from the server 605 and the classifications provided by the classifier 601C. The classifications and semantic dependencies may be stored on server 601A and/or 605, and are displayed to a user via workstation 607.

[0071] Various improvements described herein may be used together or separately. Although illustrative embodiments of the present invention have been described herein with reference to the accompanying drawings, it is to be understood that the invention is not limited to those precise embodiments, and that various other changes and modifications may be affected therein by one skilled in the art without departing from the scope or spirit of the invention.

WE CLAIM:

1. A method of asset discovery for an industrial control system, the method comprising:

characterizing (101) production processes of the industrial control system, production zones of the industrial control system, and automation packages of the industrial control system;

instrumenting (103), based on the characterization, assets of the industrial control system for incremental asset discovery; and

performing (105) incremental asset discovery, wherein incremental asset discovery comprises:

collecting (201), using the instrumentation, data for the assets of the industrial control system;

classifying (203), based on the collected data, the assets of the industrial control system; and

inferring (205) semantic dependency between the assets of the industrial control system.

2. The method of claim 1, wherein instrumenting (103) assets comprises deploying a combination of hardware and software asset discovery agents.

3. The method of claim 2, wherein the asset discovery agents comprise a combination of active and passive discovery agents and a combination of intrusive and non-intrusive discovery agents.

4. The method of claim 1, wherein instrumenting (103) assets of the industrial control system comprises providing a unidirectional overlay network for connecting to an isolated or segregated network.

5. The method of claim 4, wherein the isolated or segregated network comprises a network of air-gapped or firewalled assets.

6. The method of claim 2, wherein a software asset discovery agent comprises a relay software agent configured to retrieve information from a layer 1/2 network card and transmit the information over a layer 2/3 network.

7. The method of claim 1, wherein classifying (203) assets comprises semi-supervised machine learning, the semi-supervised machine learning comprises:

grouping similar third-party assets;

classifying the similar third-party assets based on monitored communication between the assets; and

presenting the similar third-party assets to a user to confirm classification of the similar third-party assets.

8. The method of claim 1, wherein collecting (201) data for the assets of the industrial control system comprises monitoring asset activity and communication.

9. A system for industrial asset discovery, the system comprising:

a plurality of asset discovery agents (611, 613, 615, 617, 621) configured to collect asset information from an industrial control system;

a server (605) configured to:

store the collected asset and process information received from the plurality of asset discovery agents;

classify, based on the collected asset and process information received from the server (605), the assets of the industrial control system; and

infer, based on the collected asset and process information received from the server (605), semantic dependency between the assets of the industrial control system.

10. The system of claim 9, further comprising:

a workstation (607) configured to display a user interface to the server (605).

11. The system of claim 9, wherein the plurality of asset discovery agents comprise a software agent (611) configured for installation on a process device or a networking device.

12. The system of claim 11, wherein the software agent (611) is an autonomous installing software agent, wherein the autonomous installing software agent comprises self-replicating code configured to:

analyze the process device or the networking device for installation of software agent code;

install the software agent code based on the process device or the networking device;

execute the software agent code to gather asset and process information; and

uninstall the software agent code.

13. The system of claim 11, wherein the plurality of asset discovery agents comprise an industrial computer (615), wherein the software agent (611) is deployed on the industrial computer (615).

14. The system of claim 9, wherein the plurality of asset discovery agents comprise a hardware agent (613) coupled to a process device or a networking device.

15. A method for automatically identifying and baselining industrial assets, the method comprising:

characterizing (301) production processes, production zones and automation packages of the industrial control system;

assigning (303) discovery methods to assets of the industrial control system for an incremental asset discovery;

executing (305) the incremental asset discovery of the assets of the industrial control system; and

generating (307) an output of the incremental asset discovery of the assets of the industrial control system, wherein reporting comprises displaying an attack tree for the assets of the industrial control system.

16. The method of claim 15, wherein characterizing (301) the industrial control system comprises:

mapping the production process of the industrial control system;

mapping production zones of the industrial control system;

mapping network topologies of the industrial control system; and

classifying automation packages of the industrial control system.

17. The method of claim 15, wherein assigning (303) comprises specifying a plurality of asset discovery relay agents configured to collect asset and process information from an industrial control system.

18. The method of claim 15, wherein executing (305) comprises deploying a plurality of asset discovery relay agents configured to collect asset and process information from an industrial control system.

19. The method of claim 18, wherein the plurality of asset discovery relay agents are configured to monitor asset discovery to minimize a disruption impact to the assets.

20. The method of claim 15, wherein generating (307) comprises:

inferring semantic relationships between assets of the of the industrial control system;

confirming, based on a user input, the semantic relationships.

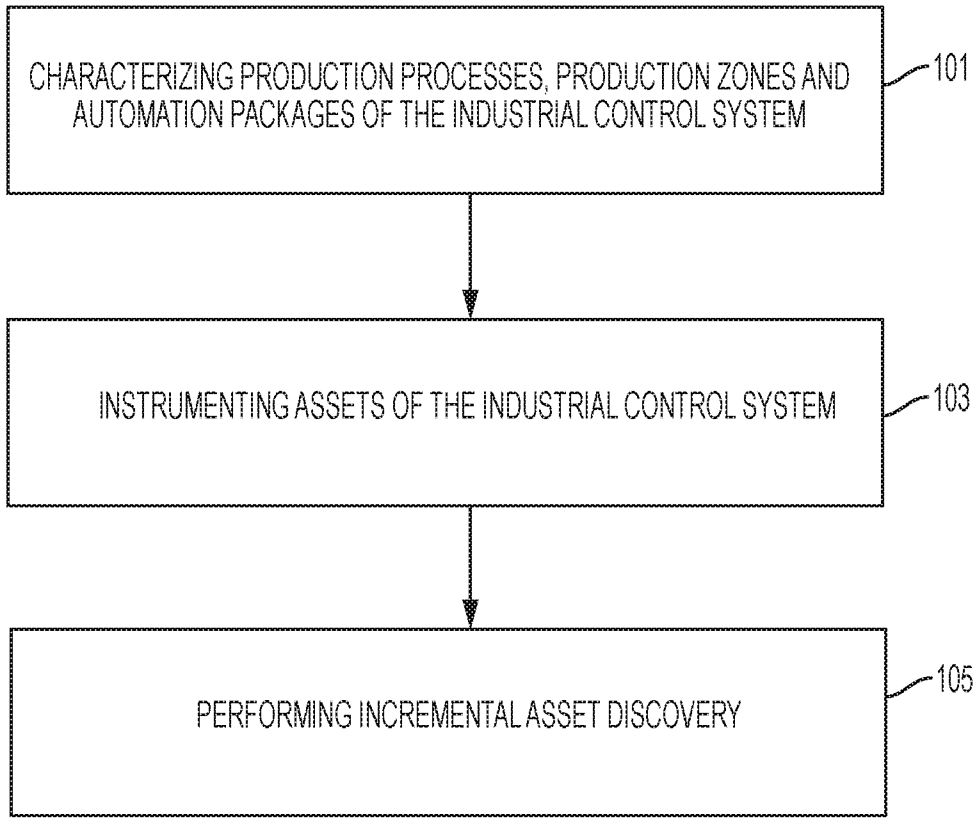
CHARACTERIZING PRODUCTION PROCESSES, PRODUCTION ZONES AND AUTOMATION PACKAGES OF THE INDUSTRIAL CONTROL SYSTEM
101
INSTRUMENTING ASSETS OF THE INDUSTRIAL CONTROL SYSTEM
103
PERFORMING INCREMENTAL ASSET DISCOVERY
105

FIG. 1

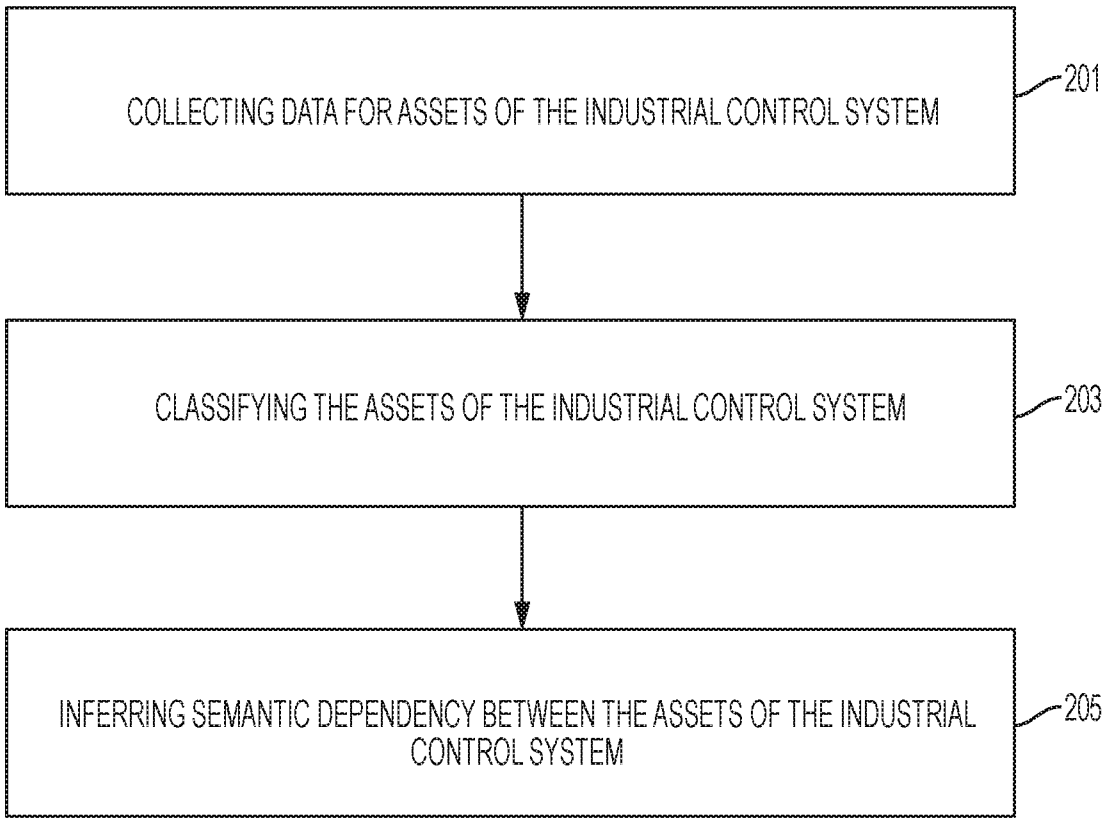
COLLECTING DATA FOR ASSETS OF THE INDUSTRIAL CONTROL SYSTEM
201
CLASSIFYING THE ASSETS OF THE INDUSTRIAL CONTROL SYSTEM
203
INFERRING SEMANTIC DEPENDENCY BETWEEN THE ASSETS OF THE INDUSTRIAL CONTROL SYSTEM
205

FIG. 2

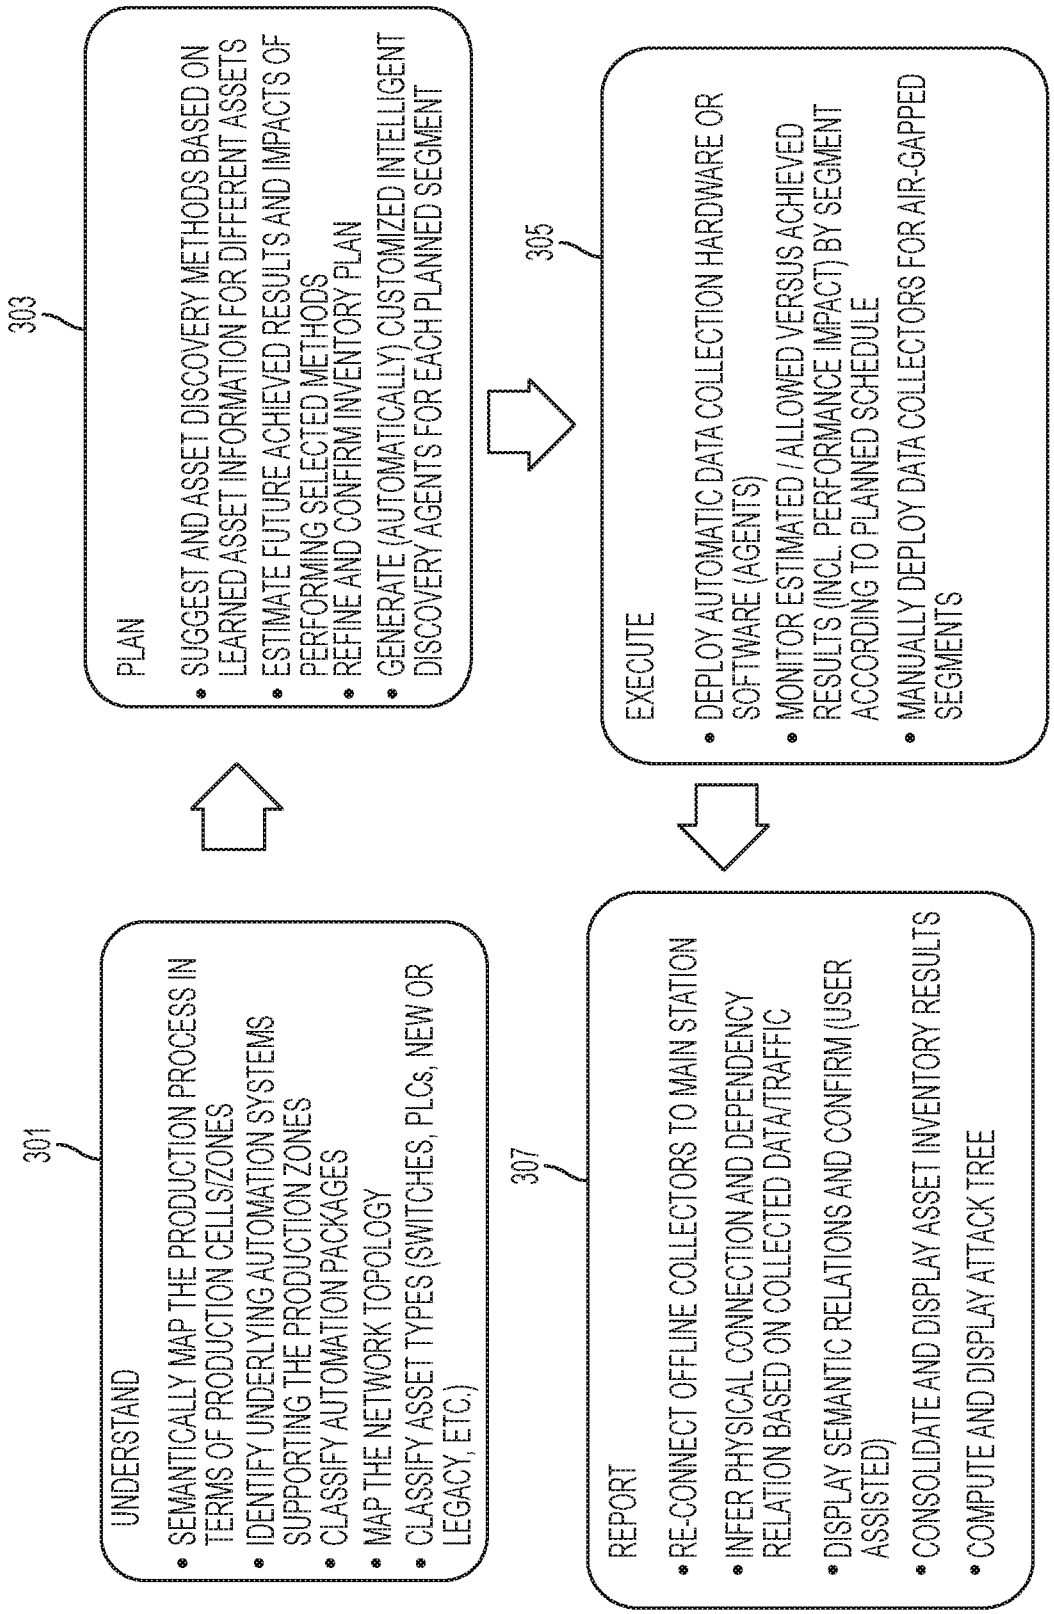
301
UNDERSTAND
• SEMANTICALLY MAP THE PRODUCTION PROCESS IN TERMS OF PRODUCTION CELLS/ZONES
• IDENTIFY UNDERLYING AUTOMATION SYSTEMS SUPPORTING THE PRODUCTION ZONES
• CLASSIFY AUTOMATION PACKAGES
• MAP THE NETWORK TOPOLOGY
• CLASSIFY ASSET TYPES (SWITCHES, PLCs, NEW OR LEGACY, ETC.)
303
PLAN
• SUGGEST AND ASSET DISCOVERY METHODS BASED ON LEARNED ASSET INFORMATION FOR DIFFERENT ASSETS
• ESTIMATE FUTURE ACHIEVED RESULTS AND IMPACTS OF PERFORMING SELECTED METHODS
• REFINE AND CONFIRM INVENTORY PLAN
• GENERATE (AUTOMATICALLY) CUSTOMIZED INTELLIGENT DISCOVERY AGENTS FOR EACH PLANNED SEGMENT
305
EXECUTE
• DEPLOY AUTOMATIC DATA COLLECTION HARDWARE OR SOFTWARE (AGENTS)
• MONITOR ESTIMATED / ALLOWED VERSUS ACHIEVED RESULTS (INCL. PERFORMANCE IMPACT) BY SEGMENT ACCORDING TO PLANNED SCHEDULE
• MANUALLY DEPLOY DATA COLLECTORS FOR AIR-GAPPED SEGMENTS
307
REPORT
• RE-CONNECT OFFLINE COLLECTORS TO MAIN STATION
• INFER PHYSICAL CONNECTION AND DEPENDENCY RELATION BASED ON COLLECTED DATA/TRAFFIC
• DISPLAY SEMANTIC RELATIONS AND CONFIRM (USER ASSISTED)
• CONSOLIDATE AND DISPLAY ASSET INVENTORY RESULTS
• COMPUTE AND DISPLAY ATTACK TREE

FIG. 3

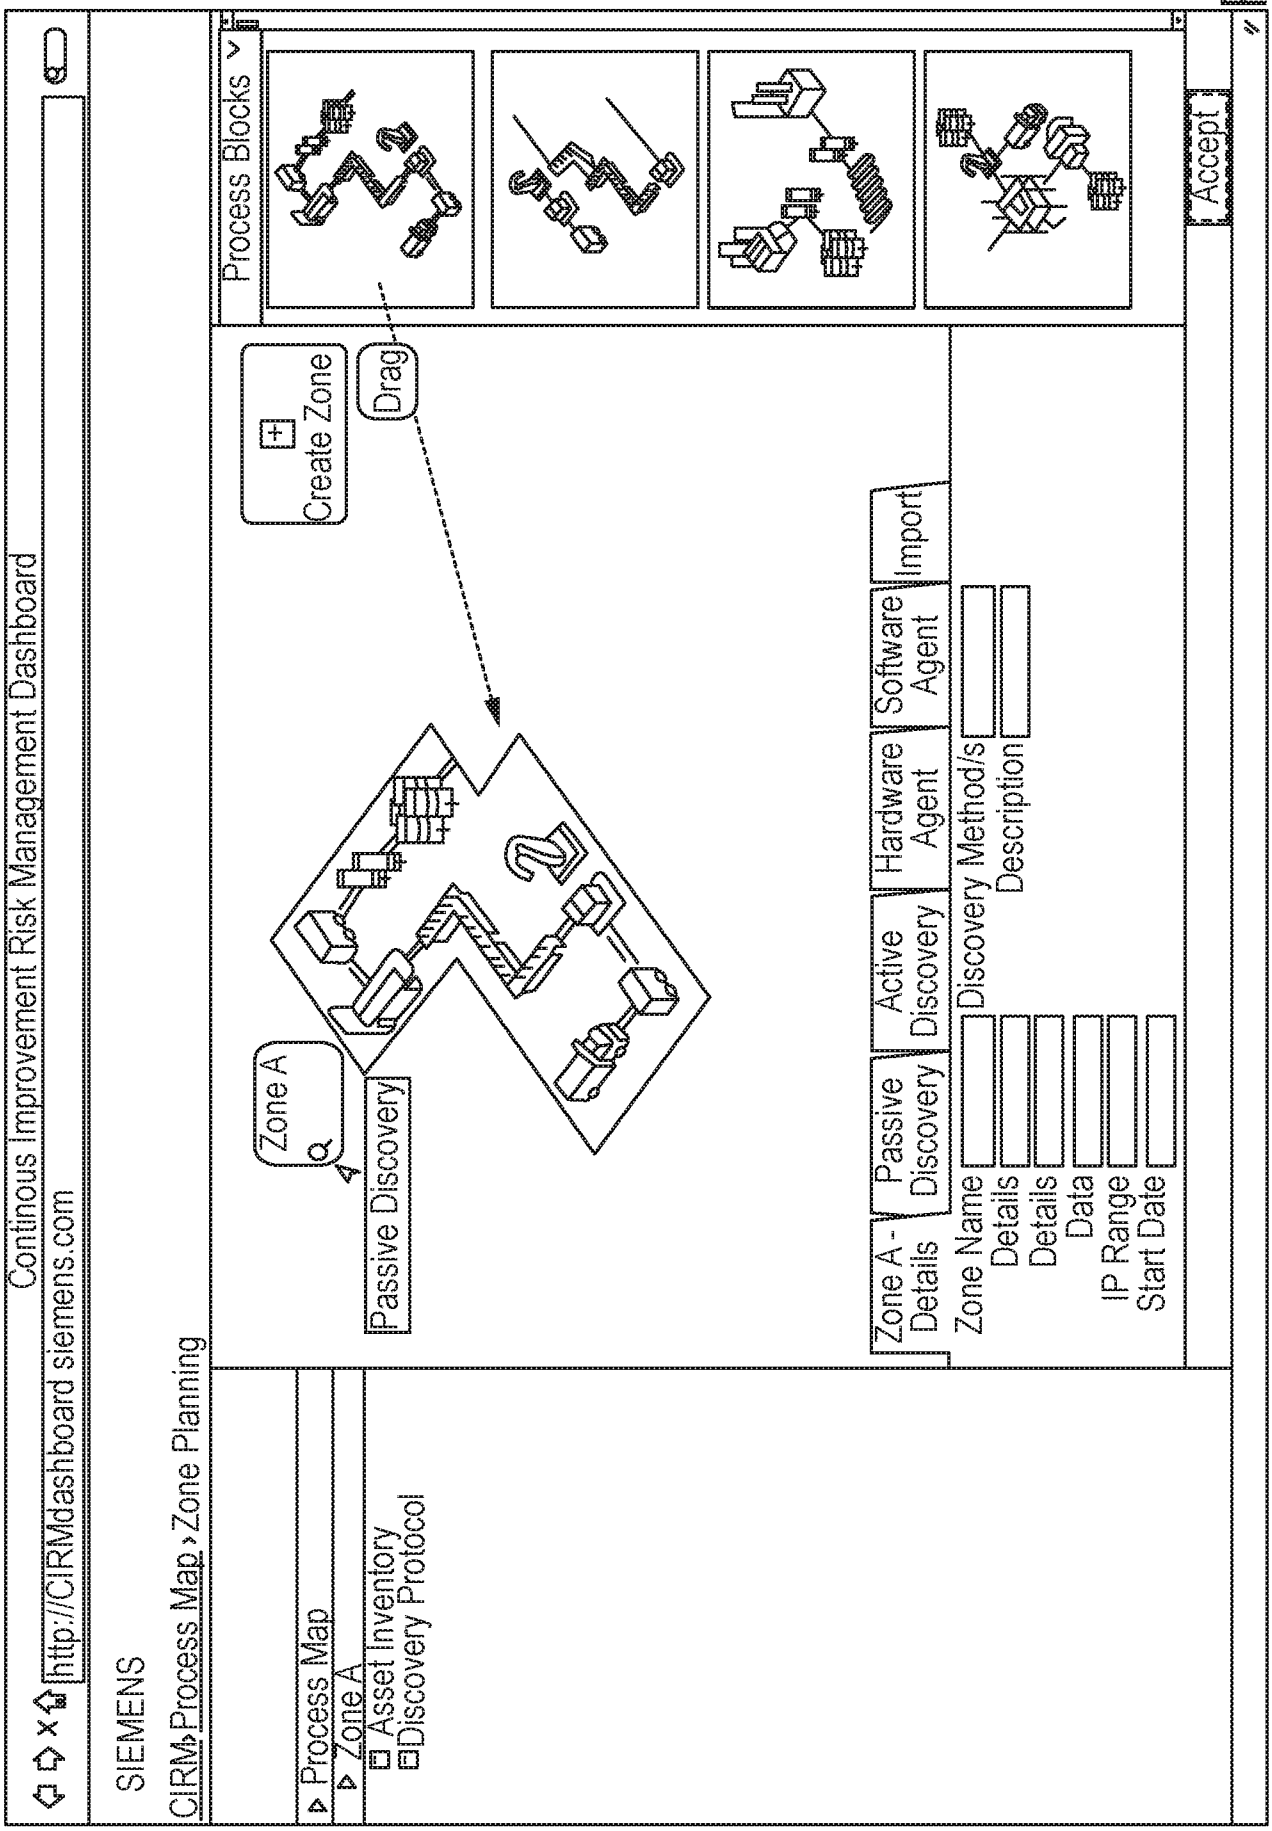
Continous Improvement Risk Management Dashboard
http://CIRMdashboard.siemens.com
SIEMENS
CIRM Process Map › Zone Planning
Process Map
Zone A
Asset Inventory
Discovery Protocol
Process Blocks
Create Zone
Drag
Zone A
Passive Discovery
Zone A - Details
Passive Discovery
Active Discovery
Hardware Agent
Software Agent
Import
Zone Name
Details
Details
Data
IP Range
Start Date
Discovery Method/s
Description
Accept

FIG. 4

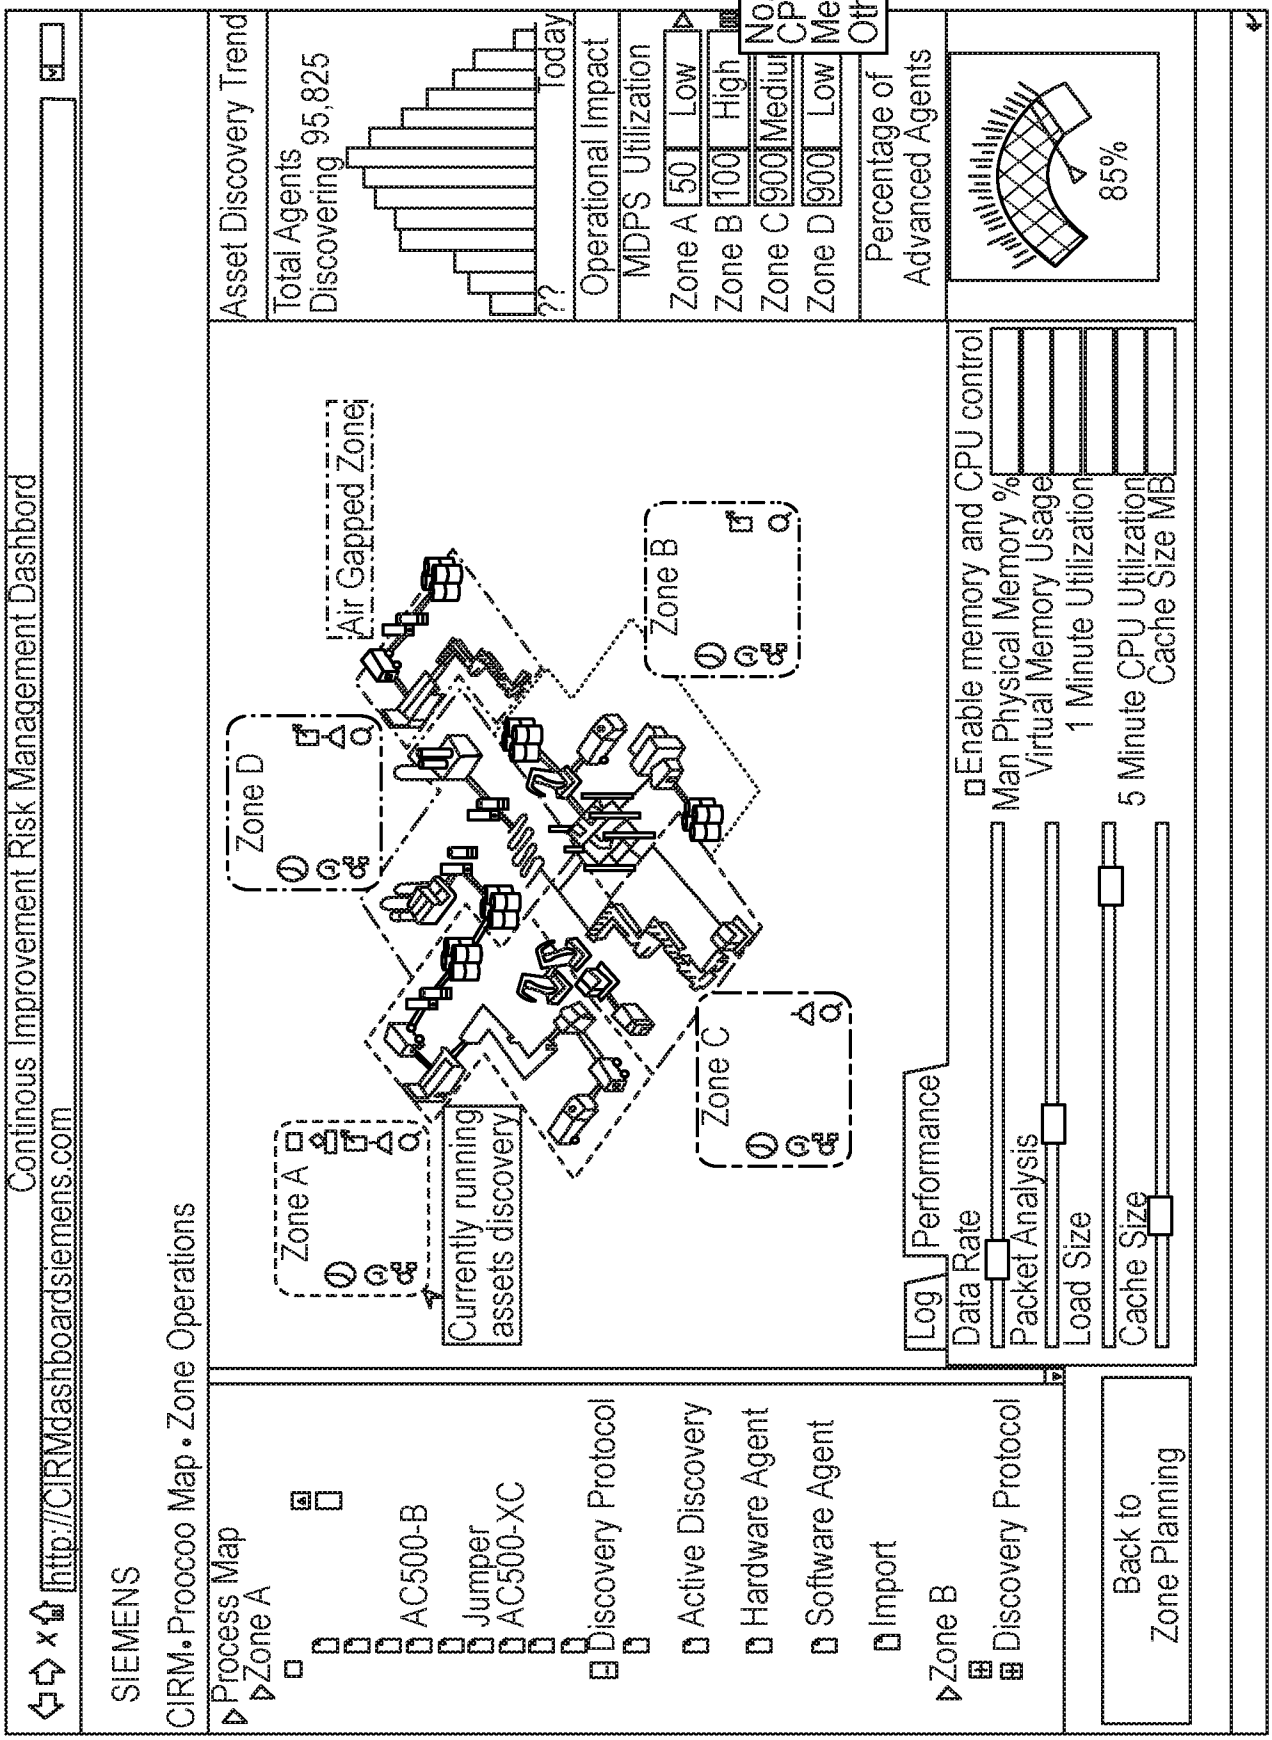
Continous Improvement Risk Management Dashbord
http://CIRMdashboardsiemens.com
SIEMENS
CIRM•Proocoo Map • Zone Operations
Process Map
Zone A
AC500-B
Jumper
AC500-XC
Discovery Protocol
Active Discovery
Hardware Agent
Software Agent
Import
Zone B
Discovery Protocol
Back to
Zone Planning
Asset Discovery Trend
Total Agents
Discovering 95,825
Today
??
Operational Impact
MDPS Utilization
Zone A 50 Low
Zone B 100 High
Zone C 900 Mediu
Zone D 900 Low
No. of Agents 2
CPU Usage 60%
Memory 40%
Other agents 40%
Percentage of
Advanced Agents
85%
Zone D
Air Gapped Zone
Zone A
Zone B
Zone C
Currently running
assets discovery
Log
Performance
Data Rate
Packet Analysis
Load Size
Cache Size
Enable memory and CPU control
Man Physical Memory %
Virtual Memory Usage
1 Minute Utilization
5 Minute CPU Utilization
Cache Size MB

FIG. 5

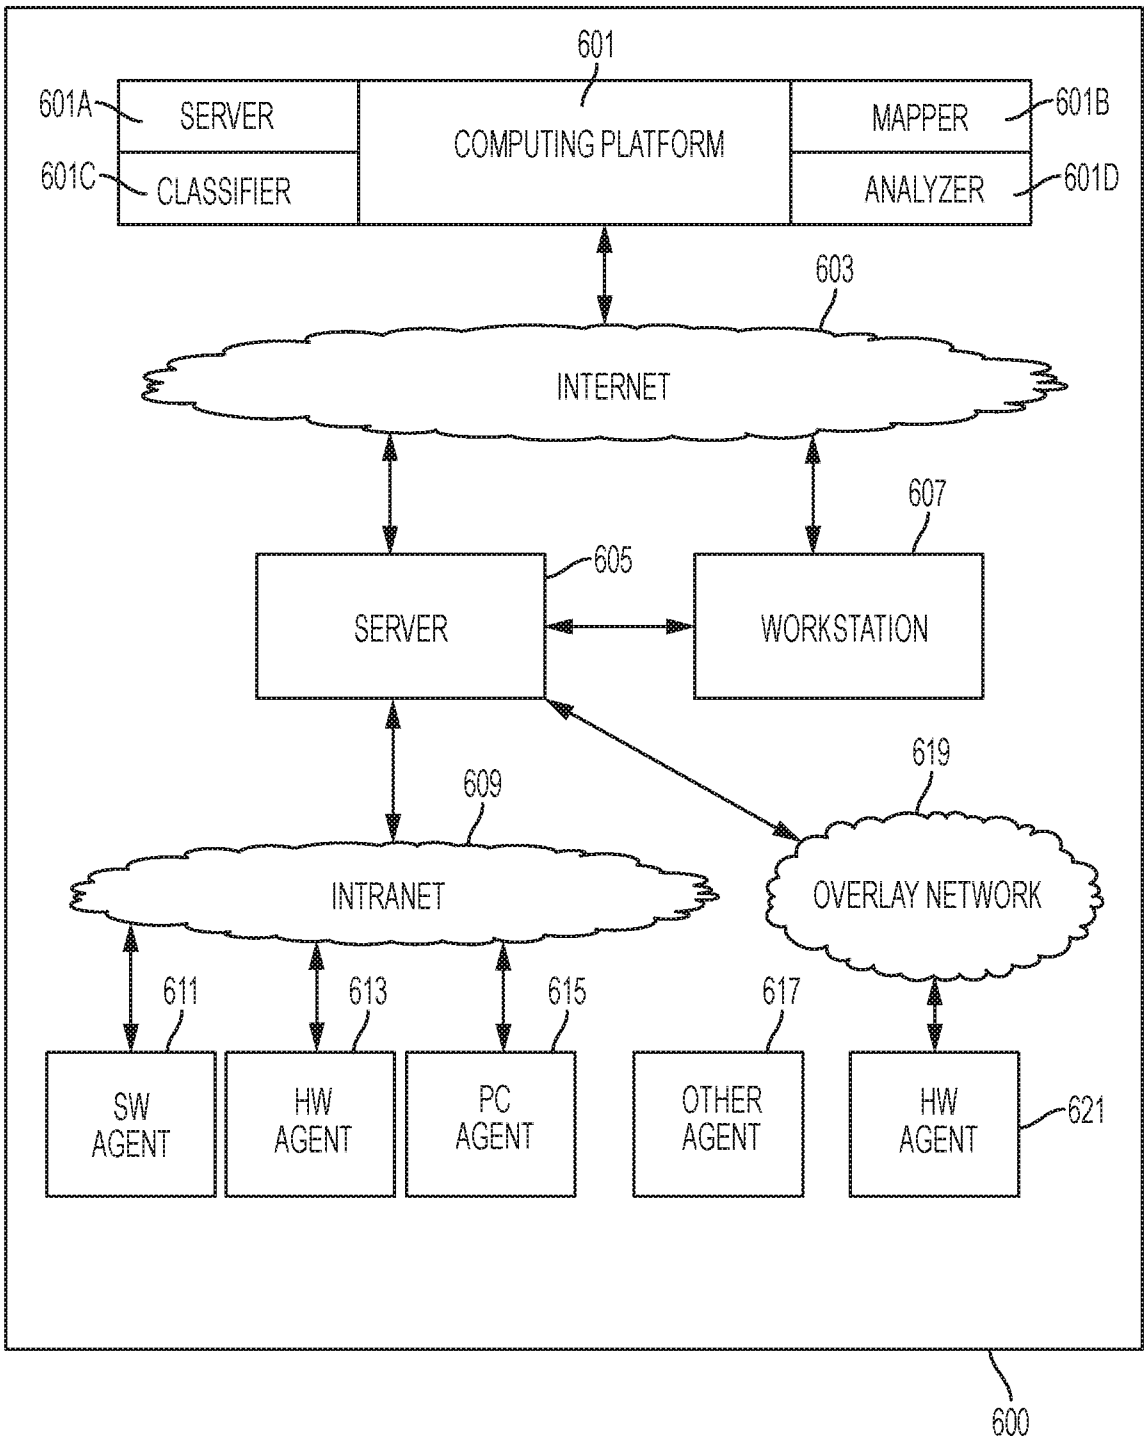
601
601A
SERVER
COMPUTING PLATFORM
MAPPER
601B
601C
CLASSIFIER
ANALYZER
601D
603
INTERNET
607
605
SERVER
WORKSTATION
619
609
INTRANET
OVERLAY NETWORK
611
613
615
617
SW AGENT
HW AGENT
PC AGENT
OTHER AGENT
HW AGENT
621
600

FIG. 6

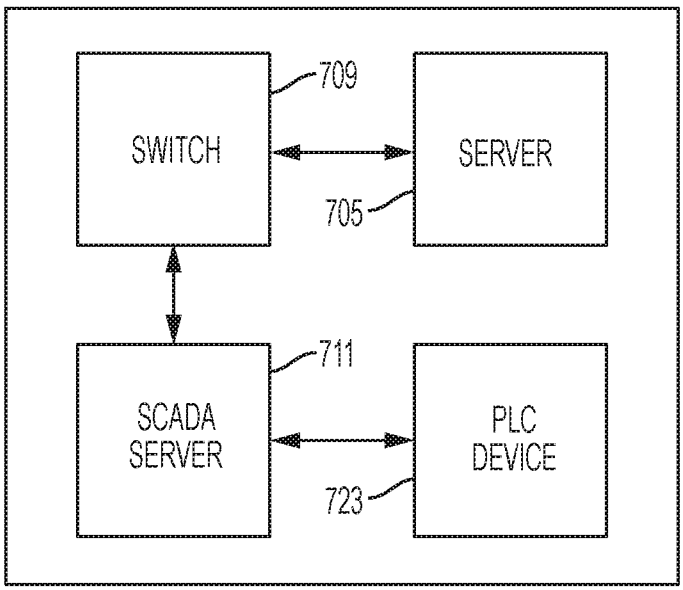

FIG. 7

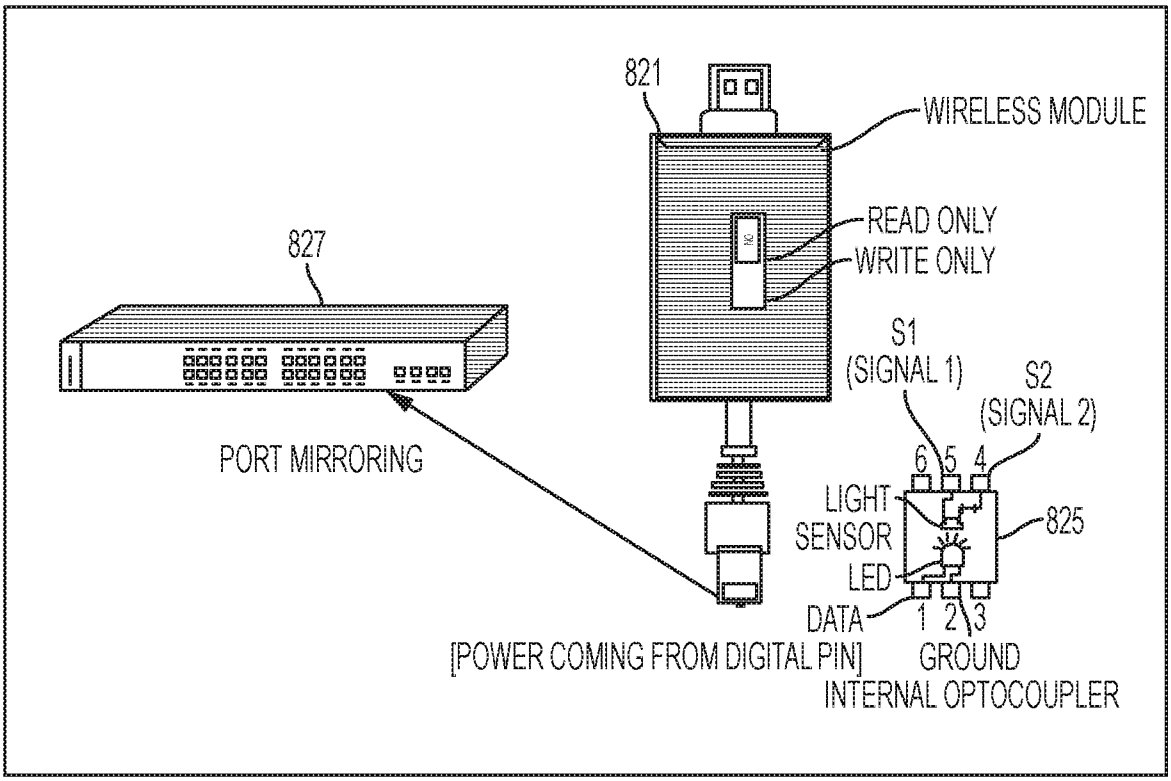
821
WIRELESS MODULE
READ ONLY
WRITE ONLY
827
S1
(SIGNAL 1)
S2
(SIGNAL 2)
6 5 4
LIGHT
SENSOR
LED
825
DATA
1 2 3
PORT MIRRORING
[POWER COMING FROM DIGITAL PIN]
GROUND
INTERNAL OPTOCOUPLER

FIG. 8

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/029239

A. CLASSIFICATION OF SUBJECT MATTER
INV. G05B19/418 G06Q10/06 G06Q50/04 H04L12/24
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G05B G06Q H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 3 070 550 A1 (ROCKWELL AUTOMATION TECH INC [US]) 21 September 2016 (2016-09-21) paragraph [0014] - paragraph [0030] paragraph [0508] paragraph [0103] - paragraph [0112] -----	1-20
A	WO 2016/126573 A1 (HONEYWELL INT INC [US]) 11 August 2016 (2016-08-11) paragraph [0004] paragraph [0013] - paragraph [0027] paragraph [0031] - paragraph [0049] -----	1-20

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
15 January 2018	23/01/2018

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Bassi, Luca

1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2017/029239

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 3070550 A1	21-09-2016	CN 105988367 A	05-10-2016
		EP 3070550 A1	21-09-2016
		US 2016274553 A1	22-09-2016
WO 2016126573 A1	11-08-2016	AU 2016215566 A1	17-08-2017
		CN 107409140 A	28-11-2017
		EP 3254411 A1	13-12-2017
		US 2016234241 A1	11-08-2016
		WO 2016126573 A1	11-08-2016