

(12) STANDARD PATENT
(19) AUSTRALIAN PATENT OFFICE

(11) Application No. **AU 2011227319 B2**

(54) Title
Intelligent network

(51) International Patent Classification(s)
G05B 19/418 (2006.01) **H04L 12/24** (2006.01)

(21) Application No: **2011227319** (22) Date of Filing: **2011.03.16**

(87) WIPO No: **WO11/116074**

(30) Priority Data

(31) Number	(32) Date	(33) Country
61/315,897	2010.03.19	US
12/830,053	2010.07.02	US

(43) Publication Date: **2011.09.22**

(44) Accepted Journal Date: **2015.01.22**

(71) Applicant(s)
ACCENTURE GLOBAL SERVICES LIMITED

(72) Inventor(s)
Dorn, John; Taft, Jeffrey D.

(74) Agent / Attorney
Murray Trento & Associates Pty Ltd, Suite 4 1175 Toorak RD, Camberwell, VIC, 3124

(56) Related Art
EP2045680
US6553336
WO2001/351190

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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
22 September 2011 (22.09.2011)



(10) International Publication Number
WO 2011/116074 A3

(51) International Patent Classification:

G05B 19/418 (2006.01) H04L 12/24 (2006.01)

(21) International Application Number:

PCT/US2011/028641

(22) International Filing Date:

16 March 2011 (16.03.2011)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

61/315,897	19 March 2010 (19.03.2010)	US
12/830,053	2 July 2010 (02.07.2010)	US

(71) Applicant (for all designated States except US): **AC-CENTURE GLOBAL SERVICES LIMITED [IE/IE]**;
3 Grand Canal Plaza, Grand Canal Street Upper, Dublin 4 (IE).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **DORN, John** [US/US]; 843 River Road, Erwinna, PA 18920 (US).
TAFT, Jeffrey, D. [US/US]; 119 W. College, Apt. 1, Canonsburg, PA 15317 (US).

(74) Agent: **GREENE, Nathan, O.**; Brinks Hofer Gilson & Lione, P.O. Box 10087, Chicago, Illinois 60610 (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, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, 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, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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, ML, MR, NE, SN, TD, TG).

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

(88) Date of publication of the international search report:

10 November 2011

(54) Title: INTELLIGENT NETWORK




FIG. 20

(57) Abstract: A network intelligence system may include a plurality of sensors located throughout and industry system. The sensors may obtain data related to various aspects of the industry network. The network intelligence system may include system endpoint intelligence and system infrastructure intelligence. The system endpoint and system infrastructure intelligence may provide distributed intelligence allowing localized decision-making to be made within the industry system based in response to system operation and occurrences. The network intelligence may include a centralized intelligence portion to communicate with endpoint and infrastructure intelligence. The centralized intelligence portion may provide responses on a localized level of the system or on a system-wide level.

WO 2011/116074 A3

INTELLIGENT NETWORK

CLAIM OF PRIORITY

[0001] This application claims the benefit of priority of United States Patent Application Ser. No. 12/830,053 filed July 2, 2010 and United States Provisional Patent Application Ser. No. 61/315,897 filed on March 19, 2010, both hereby incorporated by reference.

BACKGROUND

1. Field of the Invention

[0002] The present invention relates generally to a system and method for managing an industry network, and more particularly to a system and method for collecting data at different sections of the industry network and analyzing the collected data in order to manage the industry network.

2. Related Art

[0003] Various industries have networks associated with them. The industries may include utilities, telecommunication, vehicle travel (such as air travel, rail travel, automobile travel, bus travel, etc.), and energy exploration (such as oil wells, natural gas wells, etc.).

[0004] One such industry is the utility industry that manages a power grid. The power grid may include one or all of the following: electricity generation, electric power transmission, and electricity distribution. Electricity may be generated using generating stations, such as a coal fire power plant, a nuclear power plant, etc. For efficiency purposes, the generated electrical power is stepped up to a very high voltage (such as 345K Volts) and transmitted over transmission lines. The transmission lines may transmit the power long distances, such as across state lines or across international boundaries, until it reaches its wholesale customer, which may be a company that owns the local distribution network. The transmission lines may terminate at a transmission substation, which may step down the very high voltage to an intermediate voltage (such as 138K Volts). From a transmission substation, smaller transmission lines (such as sub-transmission lines) transmit the intermediate voltage to distribution substations. At the distribution substations, the intermediate voltage may be again stepped down to a “medium voltage” (such as from 4K Volts to 23K Volts). One or more feeder circuits may emanate from the distribution substations. For example, four to tens of feeder circuits may emanate from the distribution substation. The feeder circuit is a 3-phase

circuit comprising 4 wires (three wires for each of the 3 phases and one wire for neutral). Feeder circuits may be routed either above ground (on poles) or underground. The voltage on the feeder circuits may be tapped off periodically using distribution transformers, which step down the voltage from “medium voltage” to the consumer voltage (such as 120V). The consumer voltage may then be used by the consumer.

[0005] One or more power companies may manage the power grid, including managing faults, maintenance, and upgrades related to the power grid. However, the management of the power grid is often inefficient and costly. For example, a power company that manages the local distribution network may manage faults that may occur in the feeder circuits or on circuits, called lateral circuits, which branch from the feeder circuits. The management of the local distribution network often relies on telephone calls from consumers when an outage occurs or relies on field workers analyzing the local distribution network.

[0006] Power companies have attempted to upgrade the power grid using digital technology, sometimes called a “smart grid.” For example, more intelligent meters (sometimes called “smart meters”) are a type of advanced meter that identifies consumption in more detail than a conventional meter. The smart meter may then communicate that information via some network back to the local utility for monitoring and billing purposes (telemetry). While these recent advances in upgrading the power grid are beneficial, more advances are necessary. It has been reported that in the United States alone, half of generation capacity is unused, half the long distance transmission network capacity is unused, and two thirds of its local distribution is unused. Therefore, a need clearly exists to improve the management of the power grid.

[0007] Another such industry is the vehicle travel industry. The vehicle travel industry generally relates to the management of the movement of one or more types of means of transportation, such as an airplane, train, automobile, bus, etc. For example, the train industry includes rail lines, trains that run on the rail lines, a central control, and a network to control the rail lines/trains. The network may include the sensors to sense the various parts of the rail lines, the means by which to communicate to/from the central control, and the means by which to control the rail lines. Typically, the network for the rail industry is primitive. Specifically, the network limits the type of sensors used, the means by which to communicate to/from the central control, and the ability to control the rail lines. Therefore, a need clearly exists to improve the management of the rail lines.

SUMMARY OF THE INVENTION

- [0008]** An integration framework for facilitating communication with a central authority that manages an industry network is provided. The integration framework may be customizable and applied to the one or
5 more industries. Examples include applications to the utility industry and vehicle travel industry (such as air travel network, rail travel network, automobile travel network, bus travel network, etc.). The integration framework may also be customized and applied to a telecommunication network and to energy exploration.
- [0009]** In one aspect, the present invention provides an integration
10 framework for facilitating communication with a central authority that manages an industry network, the integration framework including a first bus for communicating operational data to the central authority, the operational data including a real-time measurement for at least a part of
15 the industry network, and a second bus for communicating event data to the central authority, the second bus being separate from the first bus, the event data being distinct from and derived from the real-time measurement and including at least one analytical determination based on the at least one real time measurement, wherein the operational data
20 is communicated via the first bus and not communicated via the second bus, and wherein the event data is communicated via the second bus and not communicated via the first bus.
- [00010]** In another aspect, the present invention provides an integration
25 framework for facilitating communication with a central authority that manages an industry network, the integration framework including a system infrastructure including a plurality of stationary sensors configured to detect at least one aspect regarding infrastructure of the industry network and generate infrastructure data indicative of the at least one aspect, and an infrastructure analysis module configured to receive the infrastructure data
30 and to receive endpoint data from one or more mobile sensors attached to mobile units of the industry network, the infrastructure analysis module

configured to generate event data based on the infrastructure and endpoint data, an operational bus in communication with the plurality of stationary sensors and the infrastructure analysis module, the operational bus configured to receive the infrastructure and endpoint data and to
5 communicate the infrastructure and endpoint data to the central authority, the infrastructure and endpoint data including at least one real-time measurement for at least a part of the industry network, and an event bus in communication with the infrastructure analysis module, the event bus configured to receive the event data and to communicate the event data to
10 the central authority, the event bus being separate from the operational bus, the event data being distinct from and derived from the real-time measurement and including at least one analytical determination as to operation of the industry network based on the at least one real time measurement, and a network core configured to receive the event data
15 exclusively through the event bus and to generate a command based on the event data, wherein events within the event data include undesired or abnormal conditions occurring within the industry system, wherein the command is sent to an endpoint that includes at least one of the mobile sensors that generated the endpoint data, for execution of the command to
20 control a mobile unit of the industry network in response to the undesired or abnormal conditions.

[0011] In yet another aspect, the present invention provides a computer-implemented method for communicating data to a central authority that manages an industry network, the method including
25 communicating, at least partly wirelessly, operational data to the central authority on a first bus, the operational data including at least one real-time measurement of at least a part of the industry network, and communicating, at least partly wirelessly, event data to the central authority on a second bus, the second bus being separate from the first
30 bus, the event data being distinct from and derived from the real-time measurement and including at least one analytical determination as to operation of the industry network based on the at least one real time

measurement, wherein the operational data is communicated via the first bus and not communicated via the second bus, and wherein the event data is communicated via the second bus and not communicated via the first bus.

- 5 **[0012]** Other systems, methods and features will be, or will become, apparent to one with skill in the art upon examination of the following figures and detailed description. It is intended that all such additional systems, methods and features included within this description be protected by the following claims.

10 **BRIEF DESCRIPTION OF THE DRAWINGS**

[0013] Figures 1A-C are block diagrams of one example of the overall architecture for a power grid.

-5a-

[0014] Figure 2 is a block diagram of the Intelligent Network Data Enterprise (INDE) CORE depicted in Figure 1.

[0015] Figures 3A-C are block diagrams of another example of the overall architecture for a power grid.

5 **[0016]** Figure 4 is a block diagram of the INDE SUBSTATION depicted in Figures 1 and 3.

[0017] Figures 5A-B are block diagrams of the INDE DEVICE depicted in Figures 1A-C and 3A-C.

10 **[0018]** Figure 6 is a block diagram of still another example of the overall architecture for a power grid.

[0019] Figure 7 is a block diagram of still another example of the overall architecture for a power grid.

[0020] Figure 8 is a block diagram including a listing of some examples of the observability processes.

15 **[0021]** Figures 9A-B illustrate flow diagrams of the Grid State Measurement & Operations Processes.

[0022] Figure 10 illustrates a flow diagram of the Non-Operational Data processes.

20 **[0023]** Figure 11 illustrates a flow diagram of the Event Management processes.

[0024] Figures 12A-C illustrate flow diagrams of the Demand Response (DR) Signaling processes.

[0025] Figures 13A-B illustrate flow diagrams of the Outage Intelligence Processes.

25 **[0026]** Figures 14A-C illustrate flow diagrams of the Fault Intelligence processes.

[0027] Figures 15A-B illustrate flow diagrams of the Meta-data Management Processes.

30

-5b-

[0028] Figure 16 illustrates a flow diagram of the Notification Agent processes.

[0029] Figure 17 illustrates a flow diagram of the Collecting Meter Data (AMI) processes.

5 **[0030]** Figures 18A-D are an example of an entity relationship diagram, which may be used to represent the baseline connectivity database.

[0031] Figures 19A-B illustrate an example of a blueprint progress flow graphic.

10 **[0032]** Figure 20 is block diagram of an example intelligent network.

[0033] Figures 21A-21C is a block diagram of one example of the overall architecture for INDE architecture.

[0034] Figure 22 is a block diagram of the INDE CORE depicted in Figure 21.

15 **[0035]** Figures 23A-23C are block diagrams of another example of the overall INDE architecture.

[0036] Figures 24A-24C are block diagrams of an example of the INDE architecture implemented in a rail network.

20 **[0037]** Figure 25 are block diagrams of an example train in the INDE architecture of Figures 24A-24C.

[0038] Figures 26A-26C are block diagrams of an example of an example of the INDE architecture implemented in an electric rail network.

25 **[0039]** Figures 27A-27C are block diagrams of an example of the INDE architecture implemented in a trucking network.

[0040] Figures 28A-28C are block diagrams of an example of the INDE architecture implemented in an automobile network.

-5c-

[0041] Figure 29 is an example operational flow diagram of the INDE architecture of Figure 20.

[0042] Figure 30 is a block diagram of an example of multiple INDE architectures being used with one another.

5 DETAILED DESCRIPTION OF EMBODIMENTS OF THE INVENTION

[0043] By way of overview, the preferred embodiments described below relate to a method and system for managing an industry network. Applicants provide examples below related to various industry networks, such as utility and vehicle travel networks (such as air travel network, 10 rail travel network, automobile travel network, bus travel network, etc.). However, other industry networks may be used including a telecommunication network, and an energy exploration network (such as a network of oil wells, a network of natural gas wells, etc.).

[0044] As discussed in more detail below, certain aspects relate to a utility network, such as the power grid itself (include hardware and software in the electric power transmission and/or the electricity distribution) or the vehicle travel network. Further, certain aspects relate to the functional capabilities of the central management of the utility network, such as the central management of the power grid and the central management of the vehicle travel network. These functional capabilities may be grouped into two categories, operation and application. The operations services enable the utilities to monitor and manage the utility network infrastructure (such as applications, network, servers, sensors, etc).

[0045] In one of the examples discussed below, the application capabilities may relate to the measurement and control of the utility network itself (such as the power grid or vehicle travel network). Specifically, the application services enable the functionality that may be important to utility network, and may include: (1) data collection processes; (2) data categorization and persistence processes; and (3) observability processes. As discussed in more detail below, using these processes allows one to “observe” the utility network, analyze the data and derive information about the utility network.

[0046] Referring now to Figure 20, a block diagram illustrating an example Intelligent Network Data Enterprise (INDE) architecture 2000 that may be applied to industry systems of various industries is shown. In one example, the INDE architecture may include a network core 2002. The network core 2002 may receive various types of information and/or data based the particular industry of use. Data and information for a particular industry may originate at a system endpoint 2006, which may represent various points with an industry system. Each system endpoint 2006 may include a number of endpoint sensors 2014 that may detect various conditions associated with an industry system. For example, the endpoint sensors 2014 may be dedicated to detecting power line flow in a utility grid or arrival/departure issues of an airline. Each of the system endpoints 2006 may include one or more processors and memory devices allowing localized analytics to be performed. In once example endpoint analytics 2016 may determine various events based on data received from the endpoint sensors 2006.

[0047] The INDE architecture 2000 may also include a system infrastructure 2008, which may support the system endpoints 2006 throughout the industry system. The system infrastructure 2008 may include infrastructure sensors 2022 distributed throughout the industry system to detect conditions associated with the industry system. In one example,

the system infrastructure 2008 may include infrastructure analytics 2020 allowing the system infrastructure to analyze the data received from the infrastructure sensors 2022.

[0048] The network core 2002 may receive information from the system endpoints 2006 and the system infrastructure 2008. In one example, the INDE architecture 2000 may include a number of buses such as an operational/non-operational bus 2010 and an event bus 2012. The operational/non-operational bus 2010 may be used to communicate both operational and non-operational data. In one example, operational data may refer to data associated the various operations of a particular industry system implementing the INDE architecture 2000. The non-operational data may refer to data in the industry associated with aspects concerning the particular industry system itself. The event bus 2012 may receive data related to various events occurring in the industry system. Events may refer to any occurrence of interest in the industry system. Thus, events may include undesired or abnormal conditions occurring in the industry system.

[0049] The INDE architecture 2000 may implement distributed intelligence in that various components of the architecture may be used to process data and determine an appropriate output. In one example, the endpoint analytics 2006 may include one or more processors, memory devices, and communication modules to allow processing to be performed based on data that is received by the endpoint sensors 2006. For example, the endpoint analytics 2016 may receive data from the endpoint sensors 2014 related to an event and may determine that the particular event is occurring based on the data. The endpoint analytics 2016 may generate an appropriate response based on the event.

[0050] The infrastructure analytics 2020 may similarly include one or more processors, memory devices, and communication modules to allow processing to be performed based on data that is received by the infrastructure sensors 2022. The system infrastructure 2008 may communicate with system endpoints 2006 allowing the system infrastructure 2008 to utilize the infrastructure analytics 2020 to evaluate and process the event data, as well as operational/non-operational data from the system endpoints 2014 and infrastructure sensors 2022.

[0051] Data may also be evaluated by the network core 2002 provided by the buses 2010 and 2012. In one example, the network core 2002 may include system analytics 2024 that includes sensor analytics 2026 and event analytics 2028. The analytics 2026 and 2028 may include one or more processors and memory devices allowing event data and operational/non-operational data to be analyzed. In one example, the sensor analytics

2024 may evaluate sensor data from endpoint sensors 2014 and infrastructure sensors 2022. In event analytics 2028 may be used to process and evaluate event data.

[0052] The network core 2002 may also include a data collection 2030. The data collection 2030 may include various data warehouses 2032 used to store raw and processed data allowing historical data to be retrieved as necessary allowing future analytics to be based on historical data.

[0053] The network core 2002 may also include system controls 2034. The system controls 2034 may be responsible for actions taken within in the industry system. For example, the system controls 2002 may include automatic controls 2036 that

automatically control various aspects of the industry system based on event data and/or operational/non-operational data. The network core 2002 may also include user controls 2038 allowing human control over the industry system which may or may not be based on event data and/or operational/non-operational bus.

[0054] An enterprise system 2004 may include various large-scale software packages for the industry. The enterprise system 2004 may receive and transmit data to the network core 2002 for use in such features such as information technology (IT) or other aspects related to the industry. In alternative examples, the buses 2010 and 2012 may be integrated into a single bus or may include additional buses. Alternative examples may also include a system infrastructure 2008 including various sub-systems.

[0055] Referring now to Figure 29, an example operational diagram of the INDE architecture 2000 is shown. In one example, a system endpoint (SE1) 2006 may determine an occurrence of an event E1. Another system endpoint (SE2) 2006 may determine an occurrence of an event E2. Each system endpoint 2006 may report the events E1 and E2 via event data to the system infrastructure 2008. The system infrastructure 2008 may analyze the event data and generate a decision D1 that may be transmitted to the system endpoints SE1 and SE2 allowing the system endpoints to implement the response.

[0056] In another example, an event E3 may be determined by the system endpoint SE1. The event data reporting the event E3 may be transmitted to the network core 2002 allowing the network core 2002 to implement system analytics 2024 and generate a decision D2 via the system controls 2034. The decision D2 may be provided to the system endpoint SE1.

[0057] In another example, the system endpoint SE1 may determine occurrence of an event E4 and notify the network core 2002 of the event E4 via event data. The network core 2002 may generate a decision D3 and provide it the system endpoint SE1 for implementation, while providing information regarding the decision D3 to the enterprise system 2004.

[0058] In another example, the system endpoint SE1 may determine occurrence of an event E5. The system endpoint SE1 may implement the endpoint analytics 2016 to determine and subsequently implement a decision D4. The decision D4 may be provided to the system infrastructure 2008 and the network core 2002 for notification and storage purposes. The examples regarding Figure 29 are illustrative and other events, operational data, and non-operational data may be communicated through the INDE system 2002.

INDE High Level Architecture Description

Overall Architecture

[0059] Turning to the drawings, wherein like reference numerals refer to like elements, Figures 1A-C illustrate one example of the overall architecture for INDE. This architecture may serve as a reference model that provides for end to end collection, transport, storage, and management of utility network data (such as smart grid data); it may also provide analytics and analytics management, as well as integration of the forgoing into utility processes and systems. Hence, it may be viewed as an enterprise-wide architecture. Certain elements, such as operational management and aspects of the utility network itself, are discussed in more detail below.

[0060] The architecture depicted in Figures 1A-C may include up to four data and integration buses: (1) a high speed sensor data bus 146 (which in the example of a power utility may include operational and non-operational data); (2) a dedicated event processing bus 147 (which may include event data); (3) an operations service bus 130 (which in the example of a power utility may serve to provide information about the smart grid to the utility back office applications); and (4) an enterprise service bus for the back office IT systems (shown in Figures 1A-C as the enterprise integration environment bus 114 for serving enterprise IT 115). The separate data buses may be achieved in one or more ways. For example, two or more of the data buses, such as the high speed sensor data bus 146 and the event processing bus 147, may be different segments in a single data bus. Specifically, the buses may have a segmented structure or platform. As discussed in

more detail below, hardware and/or software, such as one or more switches, may be used to route data on different segments of the data bus.

[0061] As another example, two or more of the data buses may be on separate buses, such as separate physical buses in terms of the hardware needed to transport data on the separate buses. Specifically, each of the buses may include cabling separate from each other. Further, some or all of the separate buses may be of the same type. For example, one or more of the buses may comprise a local area network (LAN), such as Ethernet[®] over unshielded twisted pair cabling and Wi-Fi. As discussed in more detail below, hardware and/or software, such as a router, may be used to route data on data onto one bus among the different physical buses.

[0062] As still another example, two or more of the buses may be on different segments in a single bus structure and one or more buses may be on separate physical buses. Specifically, the high speed sensor data bus 146 and the event processing bus 147 may be different segments in a single data bus, while the enterprise integration environment bus 114 may be on a physically separate bus.

[0063] Though Figures 1A-C depict four buses, fewer or greater numbers of buses may be used to carry the four listed types of data. For example, a single unsegmented bus may be used to communicate the sensor data and the event processing data (bringing the total number of buses to three), as discussed below. And, the system may operate without the operations service bus 130 and/or the enterprise integration environment bus 114.

[0064] The IT environment may be SOA-compatible. Service Oriented Architecture (SOA) is a computer systems architectural style for creating and using business processes, packaged as services, throughout their lifecycle. SOA also defines and provisions the IT infrastructure to allow different applications to exchange data and participate in business processes. Although, the use of SOA and the enterprise service bus are optional.

[0065] In the example of a power grid, the figures illustrate different elements within the overall architecture, such as the following: (1) INDE CORE 120; (2) INDE SUBSTATION 180; and (3) INDE DEVICE 188. This division of the elements within the overall architecture is for illustration purposes. Other division of the elements may be used. And, the division of elements may be different for different industries. The INDE architecture may be used to support both distributed and centralized approaches to grid intelligence, and to provide mechanisms for dealing with scale in large implementations.

[0066] The INDE Reference Architecture is one example of the technical architecture that may be implemented. For example, it may be an example of a *meta-architecture*, used to provide a starting point for developing any number of specific technical architectures, one for each industry solution (*e.g.*, different solutions for different industries) or one for each application within an industry (*e.g.*, a first solution for a first utility power grid and a second solution for a second utility power grid), as discussed below. Thus, the specific solution for a particular industry, or a particular application within an industry (such as an application to a particular utility) may include one, some, or all of the elements in the INDE Reference Architecture. And, the INDE Reference Architecture may provide a standardized starting point for solution development. Discussed below is the methodology for determining the specific technical architecture for a particular industry or a particular application within an industry (such as a particular power grid).

[0067] The INDE Reference Architecture may be an enterprise wide architecture. Its purpose may be to provide the framework for end to end management of data and analytics, such as end to end management of grid data and analytics and integration of these into utility systems and processes. Since advanced network technology (such as smart grid technology) affects every aspect of utility business processes, one should be mindful of the effects not just at the network level (such as the grid), operations, and customer premise levels, but also at the back office and enterprise levels. Consequently the INDE Reference Architecture can and does reference enterprise level SOA, for example, in order to support the SOA environment for interface purposes. This should not be taken as a requirement that a industry, such as a utility, must convert their existing IT environment to SOA before the advanced network, such as a smart grid, can be built and used. An enterprise service bus is a useful mechanism for facilitating IT integration, but it is not required in order to implement the rest of the solution. The discussion below focuses on different components of the INDE smart grid elements for a utility solution; however, one, some, or all of the components of the INDE may be applied to different industries, such as telecommunication, vehicle travel, and energy exploration.

INDE Component Groups

[0068] As discussed above, the different components in the INDE Reference Architecture may include, for example: (1) INDE CORE 120; (2) INDE SUBSTATION 180; and (3) INDE DEVICE 188. The following sections discuss these three example element groups

of the INDE Reference Architecture and provide descriptions of the components of each group.

INDE CORE

[0069] Figure 2 illustrates the INDE CORE 120, which is the portion of INDE Reference

Architecture that may reside in an operations control center, as shown in Figures 1A-C.

The INDE CORE 120 may contain a unified data architecture for storage of grid data and an integration schema for analytics to operate on that data. This data architecture may use the International Electrotechnical Commission (IEC) Common Information Model (CIM) as its top level schema. The IEC CIM is a standard developed by the electric power industry that has been officially adopted by the IEC, aiming to allow application software to exchange information about the configuration and status of an electrical network.

[0070] In addition, this data architecture may make use of federation middleware 134 to connect other types of utility data (such as, for example, meter data, operational and historical data, log and event files), and connectivity and meta-data files into a single data

architecture that may have a single entry point for access by high level applications, including enterprise applications. Real time systems may also access key data stores via the high speed data bus and several data stores can receive real time data. Different types of data may be transported within one or more buses in the smart grid. As discussed below in the INDE SUBSTATION 180 section, substation data may be collected and stored locally at the substation. Specifically, a database, which may be associated with and proximate to the substation, may store the substation data. Analytics pertaining to the substation level may also be performed at the substation computers and stored at the substation database, and all or part of the data may be transported to the control center.

[0071] The types of data transported may include operation and non-operational data, events, grid connectivity data, and network location data. Operational data may include, but is not limited to, switch state, feeder state, capacitor state, section state, meter state, FCI state, line sensor state, voltage, current, real power, reactive power, etc. Non-operational data may include, but is not limited to, power quality, power reliability, asset health, stress data, etc. The operational and non-operational data may be transported using an operational/non-operational data bus 146. Data collection applications in the electric power transmission and/or electricity distribution of the power grid may be responsible for sending some or all of the data to the operational/non-operational data bus

146. In this way, applications that need this information may be able to get the data by subscribing to the information or by invoking services that may make this data available.

[0072] Events may include messages and/or alarms originating from the various devices and sensors that are part of the smart grid, as discussed below. Events may be directly

5 generated from the devices and sensors on the smart grid network as well as generated by the various analytics applications based on the measurement data from these sensors and devices. Examples of events may include meter outage, meter alarm, transformer outage, etc. Grid components like grid devices (smart power sensors (such as a sensor with an embedded processor that can be programmed for digital processing capability)

10 temperature sensors, etc.), power system components that includes additional embedded processing (RTUs, etc), smart meter networks (meter health, meter readings, etc), and mobile field force devices (outage events, work order completions, etc) may generate event data, operational and non-operational data. The event data generated within the smart grid may be transmitted via an event bus 147.

15 [0073] Grid connectivity data may define the layout of the utility grid. There may be a base layout which defines the physical layout of the grid components (sub stations, segments, feeders, transformers, switches, reclosers, meters, sensors, utility poles, etc) and their inter-connectivity at installation. Based on the events within the grid

(component failures, maintenance activity, etc), the grid connectivity may change on a
20 continual basis. As discussed in more detail below, the structure of how the data is stored as well as the combination of the data enable the historical recreation of the grid layout at various past times. Grid connectivity data may be extracted from the Geographic Information System (GIS) on a periodic basis as modifications to the utility grid are made and this information is updated in the GIS application.

25 [0074] Network location data may include the information about the grid component on the communication network. This information may be used to send messages and information to the particular grid component. Network location data may be either entered manually into the Smart Grid database as new Smart Grid components are installed or is extracted from an Asset Management System if this information is
30 maintained externally.

[0075] As discussed in more detail below, data may be sent from various components in the grid (such as INDE SUBSTATION 180 and/or INDE DEVICE 188). The data may be sent to the INDE CORE 120 wirelessly, wired, or a combination of both. The data

may be received by utility communications networks 160, which may send the data to routing device 190. Routing device 190 may comprise software and/or hardware for managing routing of data onto a segment of a bus (when the bus comprises a segmented bus structure) or onto a separate bus. Routing device may comprise one or more switches
5 or a router. Routing device 190 may comprise a networking device whose software and hardware routes and/or forwards the data to one or more of the buses. For example, the routing device 190 may route operational and non-operational data to the operational/non-operational data bus 146. The router may also route event data to the event bus 147.

[0076] The routing device 190 may determine how to route the data based on one or more
10 methods. For example, the routing device 190 may examine one or more headers in the transmitted data to determine whether to route the data to the segment for the operational/non-operational data bus 146 or to the segment for the event bus 147.

Specifically, one or more headers in the data may indicate whether the data is operation/non-operational data (so that the routing device 190 routes the data to the
15 operational/non-operational data bus 146) or whether the data is event data (so that the routing device 190 routes the event bus 147). Alternatively, the routing device 190 may examine the payload of the data to determine the type of data (e.g., the routing device 190 may examine the format of the data to determine if the data is operational/non-operational data or event data).

[0077] One of the stores, such as the operational data warehouse 137 that stores the operational data, may be implemented as true distributed database. Another of the stores, the historian (identified as historical data 136 in Figures 1 and 2), may be implemented as a distributed database. The other “ends” of these two databases may be located in the INDE SUBSTATION 180 group (discussed below). Further, events may be stored
25 directly into any of several data stores via the complex event processing bus.

Specifically, the events may be stored in event logs 135, which may be a repository for all the events that have published to the event bus 147. The event log may store one, some, or all of the following: event id; event type; event source; event priority; and event generation time. The event bus 147 need not store the events long term, providing the
30 persistence for all the events.

[0078] The storage of the data may be such that the data may be as close to the source as possible or practicable. In one implementation, this may include, for example, the substation data being stored at the INDE SUBSTATION 180. But this data may also be

required at the operations control center level 116 to make different types of decisions that consider the grid at a much granular level. In conjunction with a distributed intelligence approach, a distributed data approach may be adopted to facilitate data availability at all levels of the solution through the use of database links and data services

as applicable. In this way, the solution for the historical data store (which may be accessible at the operations control center level 116) may be similar to that of the operational data store. Data may be stored locally at the substation and database links configured on the repository instance at the control center, provide access to the data at the individual substations. Substation analytics may be performed locally at the substation using the local data store. Historical/collective analytics may be performed at the operations control center level 116 by accessing data at the local substation instances using the database links. Alternatively, data may be stored centrally at the INDE CORE 120. However, given the amount of data that may need to be transmitted from the INDE DEVICES 188, the storage of the data at the INDE DEVICES 188 may be preferred.

Specifically, if there are thousands or tens of thousands of substations (which may occur in a power grid), the amount of data that needs to be transmitted to the INDE CORE 120 may create a communications bottleneck.

[0079] Finally, the INDE CORE 120 may program or control one, some or all of the INDE SUBSTATION 180 or INDE DEVICE 188 in the power grid (discussed below).

For example, the INDE CORE 120 may modify the programming (such as download an updated program) or provide a control command to control any aspect of the INDE SUBSTATION 180 or INDE DEVICE 188 (such as control of the sensors or analytics). Other elements, not shown in Figure 2, may include various integration elements to support this logical architecture.

[0080] Table 1 describes the certain elements of INDE CORE 120 as depicted in Figure 2.

INDE CORE Element	Description
CEP Services 144	Provides high speed, low latency event stream processing, event filtering, and multi-stream event correlation
Centralized Grid Analytics Applications 139	May consist of any number of commercial or custom analytics applications that are used in a non-real time manner, primarily operating from the data stores in CORE
Visualization/Notification Services 140	Support for visualization of data, states and event streams, and automatic notifications based on event triggers

Application Management Services 141	Services (such as Applications Support Services 142 and Distributed Computing Support 143) that support application launch and execution, web services, and support for distributed computing and automated remote program download (e.g., OSGi)
Network Management Services 145	Automated monitoring of communications networks, applications and databases; system health monitoring, failure root cause analysis (non-grid)
Grid Meta-Data Services 126	Services (such as Connectivity Services 127, Name Translation 128, and TEDS Service 129) for storage, retrieval, and update of system meta-data, including grid and communication/sensor net connectivity, point lists, sensor calibrations, protocols, device set points, etc
Grid Data/Analytics Services 123	Services (such as Sensor Data Services 124 and Analytics Management Services 125) to support access to grid data and grid analytics; management of analytics
Meter Data Management System 121	Meter data management system functions (e.g., Lodestar)
AMOS Meter Data Services	See discussion below
Real Time Complex Event Processing Bus 147	Message bus dedicated to handling event message streams – purpose of a dedicated bus is to provide high bandwidth and low latency for highly bursty event message floods. The event message may be in the form of XML message. Other types of messages may be used. Events may be segregated from operational/non-operational data, and may be transmitted on a separate or dedicated bus. Events typically have higher priority as they usually require some immediate action from a utility operational perspective (messages from defective meters, transformers, etc) The event processing bus (and the associated event correlation processing service depicted in Figure 1) may filter floods of events down into an interpretation that may better be acted upon by other devices. In addition, the event processing bus may take multiple event streams, find various patterns occurring across the multiple event streams, and provide an interpretation of multiple event streams. In this way, the event processing bus may not simply examine the event data from a single device, instead looking at multiple device (including multiple classes of devices that may be seemingly unrelated) in order to find correlations. The analysis of the single or multiple event streams may be rule based
Real Time Op/Non-Op Data	Operational data may include data reflecting the

Bus 146	current state of the electrical state of the grid that may be used in grid control (e.g., currents, voltages, real power, reactive power, etc.). Non-operational data may include data reflecting the “health” or condition of a device. Operational data has previously been transmitted directly to a specific device (thereby creating a potential “silo” problem of not making the data available to other devices or other applications). For example, operational data previously was transmitted to the SCADA (Supervisory Control And Data Acquisition) system for grid management (monitor and control grid). However, using the bus structure, the operational data may also be used for load balancing, asset utilization/optimization, system planning, etc., as discussed for example in Figures 10-19. Non-operational data was previously obtained by sending a person in the field to collect the operational data (rather than automatically sending the non-operational data to a central repository). Typically, the operational and non-operational data are generated in the various devices in the grid at predetermined times. This is in contrast to the event data, which typically is generated in bursts, as discussed below. A message bus may be dedicated to handling streams of operational and non-operational data from substations and grid devices. The purpose of a dedicated bus may be to provide constant low latency through put to match the data flows; as discussed elsewhere, a single bus may be used for transmission of both the operation and non-operational data and the event processing data in some circumstances (effectively combining the operation/non-operational data bus with the event processing bus).
Operations Service Bus 130	Message bus that supports integration of typical utility operations applications (EMS (energy management system), DMS (distribution management system), OMS (outage management system), GIS (geographic information system), dispatch) with newer smart grid functions and systems (DRMS (demand response management system), external analytics, CEP, visualization). The various buses, including the Operation/Non-

	operational Data bus 146, the Event data bus 147, and the operations Service Bus 130 may obtain weather feeds, etc. via a security framework 117. The operations service bus 130 may serve as the provider of information about the smart grid to the utility back office applications, as shown in Figure 1. The analytics applications may turn the raw data from the sensors and devices on the grid into actionable information that will be available to utility applications to perform actions to control the grid. Although most of the interactions between the utility back office applications and the INDE CORE 120 is expected to occur thru this bus, utility applications will have access to the other two buses and will consume data from those buses as well (for example, meter readings from the op/non-op data bus 146, outage events from the event bus 147)
CIM Data Warehouse 132	Top level data store for the organization of grid data; uses the IEC CIM data schema; provides the primary contact point for access to grid data from the operational systems and the enterprise systems. Federation Middleware allow communication to the various databases.
Connectivity Warehouse 131	The connectivity warehouse 131 may contain the electrical connectivity information of the components of the grid. This information may be derived from the Geographic Information System (GIS) of the utility which holds the as built geographical location of the components that make up the grid. The data in the connectivity warehouse 131 may describe the hierarchical information about all the components of the grid (substation, feeder, section, segment, branch, t-section, circuit breaker, recloser, switch, etc – basically all the assets). The connectivity warehouse 131 may have the asset and connectivity information as built. Thus, the connectivity warehouse 131 may comprise the asset database that includes all the devices and sensors attached to the components of the grid.
Meter Data Warehouse 133	The meter data warehouse 133 may provide rapid access to meter usage data for analytics. This repository may hold all the meter reading information from the meters at customer premises. The data collected from the meters may be stored in meter data warehouse 133 and provided to other utility applications for billing (or other back-office operations) as well as other analysis.
Event Logs 135	Collection of log files incidental to the operation of various utility systems. The event logs 135 may be used for post mortem analysis of events and for data

	mining
Historical Data 136	Telemetry data archive in the form of a standard data historian. Historical data 136 may hold the time series non-operational data as well as the historical operational data. Analytics pertaining to items like power quality, reliability, asset health, etc may be performed using data in historical data 136. Additionally, as discussed below, historical data 136 may be used to derive the topology of the grid at any point in time by using the historical operational data in this repository in conjunction with the as-built grid topology stored in the connectivity data mart. Further, the data may be stored as a flat record, as discussed below.
Operational Data 137	Operational Data 137 may comprise a real time grid operational database. Operational Data 137 may be built in true distributed form with elements in the substations (with links in Operational Data 137) as well as the Operations center. Specifically, the Operational Data 137 may hold data measurements obtained from the sensors and devices attached to the grid components. Historical data measurements are not held in this data store, instead being held in historical data 136. The data base tables in the Operational Data 137 may be updated with the latest measurements obtained from these sensors and devices.
DFR/SER Files 138	Digital fault recorder and serial event recorder files; used for event analysis and data mining; files generally are created in the substations by utility systems and equipment

Table 1: INDE CORE Elements

[0081] As discussed in Table 1, the real time data bus 146 (which communicates the operation and non-operational data) and the real time complex event processing bus 147 (which communicates the event processing data) into a single bus 346. An example of this is illustrated in the block diagram 300 in Figures 3A-C.

[0082] As shown in Figures 1A-C, the buses are separate for performance purposes. For CEP processing, low latency may be important for certain applications which are subject to very large message bursts. Most of the grid data flows, on the other hand, are more or less constant, with the exception of digital fault recorder files, but these can usually be retrieved on a controlled basis, whereas event bursts are asynchronous and random.

[0083] Figure 1 further shows additional elements in the operations control center 116 separate from the INDE CORE 120. Specifically, Figure 1 further shows Meter

Data Collection Head End(s) 153, a system that is responsible for communicating with meters (such as collecting data from them and providing the collected data to the utility).

Demand Response Management System 154 is a system that communicates with equipment at one or more customer premises that may be controlled by the utility.

5 Outage Management System 155 is a system that assists a utility in managing outages by tracking location of outages, by managing what is being dispatched, and by how they are being fixed. Energy Management System 156 is a transmission system level control system that controls the devices in the substations (for example) on the transmission

10 grid. Distribution Management System 157 is a distribution system level control system that controls the devices in the substations and feeder devices (for example) for

distribution grids. IP Network Services 158 is a collection of services operating on one or more servers that support IP-type communications (such as DHCP and FTP).

Dispatch Mobile Data System 159 is a system that transmits/receives messages to mobile data terminals in the field. Circuit & Load Flow Analysis, Planning, Lightning Analysis

15 and Grid Simulation Tools 152 are a collection of tools used by a utility in the design, analysis and planning for grids. IVR (integrated voice response) and Call Management 151 are systems to handle customer calls (automated or by attendants). Incoming telephone calls regarding outages may be automatically or manually entered and

20 forwarded to the Outage Management System 155. Work Management System 150 is a system that monitors and manages work orders. Geographic Information System 149 is a database that contains information about where assets are located geographically and how the assets are connected together. If the environment has a Services Oriented Architecture (SOA), Operations SOA Support 148 is a collection of services to support the SOA environment.

25 **[0084]** One or more of the systems in the Operations Control Center 116 that are outside of the INDE Core 120 are legacy product systems that a utility may have.

Examples of these legacy product systems include the Operations SOA Support 148, Geographic Information System 149, Work Management System 150, Call Management 151, Circuit & Load Flow Analysis, Planning, Lightning Analysis and Grid Simulation

30 Tools 152, Meter Data Collection Head End(s) 153, Demand Response Management System 154, Outage Management System 155, Energy Management System 156,

Distribution Management System 157, IP Network Services 158, and Dispatch Mobile Data System 159. However, these legacy product systems may not be able to process or

handle data that is received from a smart grid. The INDE Core 120 may be able to receive the data from the smart grid, process the data from the smart grid, and transfer the processed data to the one or more legacy product systems in a fashion that the legacy product systems may use (such as particular formatting particular to the legacy product system). In this way, the INDE Core 120 may be viewed as a middleware.

[0085] The operations control center 116, including the INDE CORE 120, may communicate with the Enterprise IT 115. Generally speaking, the functionality in the Enterprise IT 115 comprises back-office operations. Specifically, the Enterprise IT 115 may use the enterprise integration environment bus 114 to send data to various systems within the Enterprise IT 115, including Business Data Warehouse 104, Business Intelligence Applications 105, Enterprise Resource Planning 106, various Financial Systems 107, Customer Information System 108, Human Resource System 109, Asset Management System 110, Enterprise SOA Support 111, Network Management System 112, and Enterprise Messaging Services 113. The Enterprise IT 115 may further include a portal 103 to communicate with the Internet 101 via a firewall 102.

INDE SUBSTATION

[0086] Figure 4 illustrates an example of the high level architecture for the INDE SUBSTATION 180 group. This group may comprise elements that are actually hosted in the substation 170 at a substation control house on one or more servers co-located with the substation electronics and systems.

[0087] Table 2 below lists and describes certain INDE SUBSTATION 180 group elements. Data security services 171 may be a part of the substation environment; alternatively, they may be integrated into the INDE SUBSTATION 180 group.

INDE SUBSTATION ELEMENTS	Description
Non-Operational Data Store 181	Performance and health data; this is a distributed data historian component
Operational Data Store 182	Real time grid state data; this is part of a true distributed database
Interface/Communications Stack 187	Support for communications, including TCP/IP, SNMP, DHCP, SFTP, IGMP, ICMP, DNP3, IEC 61850, etc.
Distributed/remote computing support 186	Support for remote program distribution, inter-process communication, etc. (DCE, JINI, OSGi for example)
Signal/Waveform Processing 185	Support for real time digital signal processing components; data normalization; engineering units conversions

Detection/Classification Processing 184	Support for real time event stream processing, detectors and event/waveform classifiers (ESP, ANN, SVM, etc.)
Substation Analytics 183	Support for programmable real time analytics applications ; DNP3 scan master; The substation analytics may allow for analysis of the real-time operational and non-operational data in order to determine if an “event” has occurred. The “event” determination may be rule-based with the rules determining whether one of a plurality of possible events occurring based on the data. The substation analytics may also allow for automatic modification of the operation of the substation based on a determined event. In this way, the grid (including various portions of the grid) may be “self-healing.” This “self-healing” aspect avoids the requirement that the data be transmitted to a central authority, the data be analyzed at the central authority, and a command be sent from the central authority to the grid before the problem in the grid be corrected. In addition to the determination of the “event,” the substation analytics may also generate a work-order for transmission to a central authority. The work-order may be used, for example, for scheduling a repair of a device, such as a substation.
Substation LAN 172	Local networking inside the substation to various portions of the substation, such as microprocessor relays 173, substation instrumentation 174, event file recorders 175, and station RTUs 176.
Security services 171	The substation may communicate externally with various utility communications networks via the security services layer.

Table 2 INDE SUBSTATION Elements

[0088] As discussed above, different elements within the smart grid may include additional functionality including additional processing/analytical capability and database resources. The use of this additional functionality within various elements in the smart grid enables distributed architectures with centralized management and administration of applications and network performance. For functional, performance, and scalability reasons, a smart grid involving thousands to tens of thousands of INDE SUBSTATIONS 180 and tens of thousands to millions of grid devices may include distributed processing, data management, and process communications.

[0089] The INDE SUBSTATION 180 may include one or more processors and one or more memory devices (such as substation non-operational data 181 and substation

operations data 182). Non-operational data 181 and substation operations data 182 may be associated with and proximate to the substation, such as located in or on the INDE SUBSTATION 180. The INDE SUBSTATION 180 may further include components of the smart grid that are responsible for the observability of the smart grid at a substation level. The INDE SUBSTATION 180 components may provide three primary functions: operational data acquisition and storage in the distributed operational data store; acquisition of non-operational data and storage in the historian; and local analytics processing on a real time (such as a sub-second) basis. Processing may include digital signal processing of voltage and current waveforms, detection and classification processing, including event stream processing; and communications of processing results to local systems and devices as well as to systems at the operations control center 116. Communication between the INDE SUBSTATION 180 and other devices in the grid may be wired, wireless, or a combination of wired and wireless. For example, the transmission of data from the INDE SUBSTATION 180 to the operations control center 116 may be wired. The INDE SUBSTATION 180 may transmit data, such as operation/non-operational data or event data, to the operations control center 116. Routing device 190 may route the transmitted data to one of the operational/non-operational data bus 146 or the event bus 147.

[0090] Demand response optimization for distribution loss management may also be performed here. This architecture is in accordance with the distributed application architecture principle previously discussed.

[0091] For example, connectivity data may be duplicated at the substation 170 and at the operations control center 116, thereby allowing a substation 170 to operate independently even if the data communication network to the operations control center 116 is not functional. With this information (connectivity) stored locally, substation analytics may be performed locally even if the communication link to the operations control center is inoperative.

[0092] Similarly, operational data may be duplicated at the operations control center 116 and at the substations 170. Data from the sensors and devices associated with a particular substation may be collected and the latest measurement may be stored in this data store at the substation. The data structures of the operational data store may be the same and hence database links may be used to provide seamless access to data that resides on the substations thru the instance of the operational data store at the control center. This

provides a number of advantages including alleviating data replication and enabling substation data analytics, which is more time sensitive, to occur locally and without reliance on communication availability beyond the substation. Data analytics at the operations control center level 116 may be less time sensitive (as the operations control center 116 may typically examine historical data to discern patterns that are more predictive, rather than reactive) and may be able to work around network issues if any.

[0093] Finally, historical data may be stored locally at the substation and a copy of the data may be stored at the control center. Or, database links may be configured on the repository instance at the operations control center 116, providing the operations control center access to the data at the individual substations. Substation analytics may be performed locally at the substation 170 using the local data store. Specifically, using the additional intelligence and storage capability at the substation enables the substation to analyze itself and to correct itself without input from a central authority. Alternatively, historical/collective analytics may also be performed at the operations control center level 116 by accessing data at the local substation instances using the database links.

[0094] INDE DEVICE

[0095] The INDE DEVICE 188 group may comprise any variety of devices within the smart grid, including various sensors within the smart grid, such as various distribution grid devices 189 (*e.g.*, line sensors on the power lines), meters 163 at the customer premises, etc. The INDE DEVICE 188 group may comprise a device added to the grid with particular functionality (such as a smart Remote Terminal Unit (RTU) that includes dedicated programming), or may comprise an existing device within the grid with added functionality (such as an existing open architecture pole top RTU that is already in place in the grid that may be programmed to create a smart line sensor or smart grid device).

The INDE DEVICE 188 may further include one or more processors and one or more memory devices.

[0096] Existing grid devices may not be open from the software standpoint, and may not be capable of supporting much in the way of modern networking or software services. The existing grid devices may have been designed to acquire and store data for occasional offload to some other device such as a laptop computer, or to transfer batch files via PSTN line to a remote host on demand. These devices may not be designed for operation in a real time digital network environment. In these cases, the grid device data may be obtained at the substation level 170, or at the operations control center level 116,

depending on how the existing communications network has been designed. In the case of meters networks, it will normally be the case that data is obtained from the meter data collection engine, since meter networks are usually closed and the meters may not be addressed directly. As these networks evolve, meters and other grid devices may be individually addressable, so that data may be transported directly to where it is needed, which may not necessarily be the operations control center 116, but may be anywhere on the grid.

[0097] Devices such as faulted circuit indicators may be married with wireless network interface cards, for connection over modest speed (such as 100 kbps) wireless networks.

These devices may report status by exception and carry out fixed pre-programmed functions. The intelligence of many grid devices may be increased by using local smart RTUs. Instead of having poletop RTUs that are designed as fixed function, closed architecture devices, RTUs may be used as open architecture devices that can be programmed by third parties and that may serve as an INDE DEVICE 188 in the INDE

Reference Architecture. Also, meters at customers' premises may be used as sensors. For example, meters may measure consumption (such as how much energy is consumed for purposes of billing) and may measure voltage (for use in volt/VAr optimization).

[0098] Figures 5A-B illustrate an example architecture for INDE DEVICE 188 group. Table 3 describes the certain INDE DEVICE 188 elements. The smart grid device may include an embedded processor, so the processing elements are less like SOA services and more like real time program library routines, since the DEVICE group is implemented on a dedicated real time DSP or microprocessor.

INDE DEVICE ELEMENTS	Description
Ring buffers 502	Local circular buffer storage for digital waveforms sampled from analog transducers (voltage and current waveforms for example) which may be used hold the data for waveforms at different time periods so that if an event is detected, the waveform data leading up to the event may also be stored
Device status buffers 504	Buffer storage for external device state and state transition data
Three phase frequency tracker 506	Computes a running estimate of the power frequency from all three phases; used for frequency correction to other data as well as in grid stability and power quality measures (especially as relates to DG)
Fourier transform block 508	Conversion of time domain waveforms to frequency domain to enable frequency domain

	analytics
Time domain signal analytics 510	Processing of the signals in the time domain; extraction of transient and envelop behavior measures
Frequency domain signal analytics 512	Processing of the signals in the frequency domain; extraction of RMS and power parameters
Secondary signal analytics 514	Calculation and compensation of phasors; calculation of selected error/fault measures
Tertiary signal analytics 516	Calculation of synchrophasors based on GPS timing and a system reference angle
Event analysis and triggers 518	Processing of all analytics for event detection and triggering of file capture. Different types of INDE DEVICES may include different event analytical capability. For example, a line sensor may examine ITIC events, examining spikes in the waveform. If a spike occurs (or a series of spikes occur), the line sensor, with the event analytical capability, may determine that an “event” has occurred and also may provide a recommendation as to the cause of the event. The event analytical capability may be rule-based, with different rules being used for different INDE DEVICES and different applications.
File storage - capture/formatting/transmission 520	Capture of data from the ring buffers based on event triggers
Waveform streaming service 522	Support for streaming of waveforms to a remote display client
Communications stack	Support for network communications and remote program load
GPS Timing 524	Provides high resolution timing to coordinate applications and synchronize data collection across a wide geographic area. The generated data may include a GPS data frame time stamp 526.
Status analytics 528	Capture of data for status messages

Table 3 INDE DEVICE Elements

[0099] Figure 1A further depicts customer premises 179, which may include one or more Smart Meters 163, an in-home display 165, one or more sensors 166, and one or more controls 167. In practice, sensors 166 may register data at one or more devices at the customer premises 179. For example, a sensor 166 may register data at various major appliances within the customer premises 179, such as the furnace, hot water heater, air conditioner, etc. The data from the one or more sensors 166 may be sent to the Smart Meter 163, which may package the data for transmission to the operations control center 116 via utility communication network 160. The in-home display 165 may provide the customer at the customer premises with an output device to view, in

real-time, data collected from Smart Meter 163 and the one or more sensors 166. In addition, an input device (such as a keyboard) may be associated with in-home display 165 so that the customer may communicate with the operations control center 116. In one embodiment, the in-home display 165 may comprise a computer resident at the customer premises.

[00100] The customer premises 165 may further include controls 167 that may control one or more devices at the customer premises 179. Various appliances at the customer premises 179 may be controlled, such as the heater, air conditioner, etc., depending on commands from the operations control center 116.

[00101] As depicted in Figure 1A, the customer premises 169 may communicate in a variety of ways, such as via the Internet 168, the public-switched telephone network (PSTN) 169, or via a dedicated line (such as via collector 164). Via any of the listed communication channels, the data from one or more customer premises 179 may be sent. As shown in Figure 1, one or more customer premises 179 may comprise a Smart Meter Network 178 (comprising a plurality of smart meters 163), sending data to a collector 164 for transmission to the operations control center 116 via the utility management network 160. Further, various sources of distributed energy generation/storage 162 (such as solar panels, etc.) may send data to a monitor control 161 for communication with the operations control center 116 via the utility management network 160.

[00102] As discussed above, the devices in the power grid outside of the operations control center 116 may include processing and/or storage capability. The devices may include the INDE SUBSTATION 180 and the INDE DEVICE 188. In addition to the individual devices in the power grid including additional intelligence, the individual devices may communicate with other devices in the power grid, in order to exchange information (include sensor data and/or analytical data (such as event data)) in order to analyze the state of the power grid (such as determining faults) and in order to change the state of the power grid (such as correcting for the faults). Specifically, the individual devices may use the following: (1) intelligence (such as processing capability); (2) storage (such as the distributed storage discussed above); and (3) communication (such as the use of the one or more buses discussed above). In this way, the individual devices in the power grid may communicate and cooperate with one another without oversight from the operations control center 116.

[00103] For example, the INDE architecture disclosed above may include a device that senses at least one parameter on the feeder circuit. The device may further include a processor that monitors the sensed parameter on the feeder circuit and that analyzes the sensed parameter to determine the state of the feeder circuit. For example, the analysis of the sense parameter may comprise a comparison of the sensed parameter with a predetermined threshold and/or may comprise a trend analysis. One such sensed parameter may include sensing the waveforms and one such analysis may comprise determining whether the sensed waveforms indicate a fault on the feeder circuit. The device may further communicate with one or more substations. For example, a particular substation may supply power to a particular feeder circuit. The device may sense the state of the particular feeder circuit, and determine whether there is a fault on the particular feeder circuit. The device may communicate with the substation. The substation may analyze the fault determined by the device and may take corrective action depending on the fault (such as reducing the power supplied to the feeder circuit). In the example of the device sending data indicating a fault (based on analysis of waveforms), the substation may alter the power supplied to the feeder circuit without input from the operations control center 116. Or, the substation may combine the data indicating the fault with information from other sensors to further refine the analysis of the fault. The substation may further communicate with the operations control center 116, such as the outage intelligence application (such as discussed Figures 13A-B) and/or the fault intelligence application (such as discussed in Figures 14A-C). Thus, the operations control center 116 may determine the fault and may determine the extent of the outage (such as the number of homes affected by the fault). In this way, the device sensing the state of the feeder circuit may cooperatively work with the substation in order to correct a potential fault with or without requiring the operations control center 116 to intervene.

[00104] As another example, a line sensor, which includes additional intelligence using processing and/or memory capability, may produce grid state data in a portion of the grid (such as a feeder circuit). The grid state data may be shared with the demand response management system 155 at the operations control center 116. The demand response management system 155 may control one or more devices at customer sites on the feeder circuit in response to the grid state data from the line sensor. In particular, the demand response management system 155 may command the energy management system 156 and/or the distribution management system 157 to reduce load on the feeder

circuit by turning off appliances at the customer sites that receive power from the feeder circuit in response to line sensor indicating an outage on the feeder circuit. In this way, the line sensor in combination with the demand response management system 155 may shift automatically load from a faulty feeder circuit and then isolate the fault.

- 5 **[00105]** As still another example, one or more relays in the power grid may have a microprocessor associated with it. These relays may communicate with other devices and/or databases resident in the power grid in order to determine a fault and/or control the power grid.

INDS Concept and Architecture

10 Outsourced Smart Grid Data/Analytics Services Model

- [00106]** One application for the smart grid architecture allows the utility to subscribe to grid data management and analytics services while maintaining traditional control systems and related operational systems in-house. In this model, the utility may install and own grid sensors and devices (as described above), and may either own and
- 15 operate the grid data transport communication system, or may outsource it. The grid data may flow from the utility to a remote Intelligent Network Data Services (INDS) hosting site, where the data may be managed, stored, and analyzed. The utility may then subscribe to data and analytics services under an appropriate services financial model. The utility may avoid the initial capital expenditure investment and the ongoing costs of
- 20 management, support, and upgrade of the smart grid data/analytics infrastructure, in exchange for fees. The INDE Reference Architecture, described above, lends itself to the outsourcing arrangement described herein.

INDS Architecture for Smart Grid Services

- [00107]** In order to implement the INDS services model, the INDE Reference
- 25 Architecture may be partitioned into a group of elements that may be hosted remotely, and those that may remain at the utility. Figures 6A-C illustrate how the utility architecture may look once the INDE CORE 120 has been made remote. A server may be included as part of the INDE CORE 120 that may act as the interface to the remote systems. To the utility systems, this may appear as a virtual INDE CORE 602.
- 30 **[00108]** As the overall block diagram 600 in Figures 6A-C show, the INDE SUBSTATION 180 and INDE DEVICE 188 groups are unchanged from that depicted in Figures 1A-C. The multiple bus structure may also still be employed at the utility as well.

[00109] The INDE CORE 120 may be remotely hosted, as the block diagram 700 in Figure 7 illustrates. At the hosting site, INDE COREs 120 may be installed as needed to support utility INDS subscribers (shown as North American INDS Hosting Center 702). Each CORE 120 may be a modular system, so that adding a new subscriber is a routine operation. A party separate from the electric utility may manage and support the software for one, some, or all of the INDE COREs 120, as well as the applications that are downloaded from the INDS hosting site to each utility's INDE SUBSTATION 180 and INDE DEVICES 188.

[00110] In order to facilitate communications, high bandwidth low latency communications services, such as via network 704 (e.g., a MPLS or other WAN), may be use that can reach the subscriber utility operations centers, as well as the INDS hosting sites. As shown in Figure 7, various areas may be served, such as California, Florida, and Ohio. This modularity of the operations not only allows for efficient management of various different grids. It also allows for better inter-grid management. There are instances where a failure in one grid may affect operations in a neighboring grid. For example, a failure in the Ohio grid may have a cascade effect on operations in a neighboring grid, such as the mid-Atlantic grid. Using the modular structure as illustrated in Figure 7 allows for management of the individual grids and management of inter-grid operations. Specifically, an overall INDS system (which includes a processor and a memory) may manage the interaction between the various INDE COREs 120. This may reduce the possibility of a catastrophic failure that cascades from one grid to another. For example, a failure in the Ohio grid may cascade to a neighboring grid, such as the mid-Atlantic grid. The INDE CORE 120 dedicated to managing the Ohio grid may attempt to correct for the failure in the Ohio grid. And, the overall INDS system may attempt to reduce the possibility of a cascade failure occurring in neighboring grids.

Specific examples of functionality in INDE CORE

[00111] As shown in Figures 1, 6, and 7, various functionalities (represented by blocks) are included in the INDE CORE 120, two of which depicted are meter data management services (MDMS) 121 and metering analytics and services 122. Because of the modularity of the architecture, various functionality, such as MDMS 121 and metering analytics and services 122, may be incorporated.

Observability Processes

[00112] As discussed above, one functionality of the application services may include observability processes. The observability processes may allow the utility to “observe” the grid. These processes may be responsible for interpreting the raw data received from all the sensors and devices on the grid and turning them into actionable information. Figure 8 includes a listing of some examples of the observability processes.

[00113] Figures 9A-B illustrate a flow diagram 900 of the Grid State Measurement & Operations Processes. As shown, the Data Scanner may request meter data, as shown at block 902. The request may be sent to one or more grid devices, substation computers, and line sensor RTUs. In response to the request, the devices may collect operations data, as shown at blocks 904, 908, 912, and may send data (such as one, some or all of the operational data, such as Voltage, Current, Real Power, and Reactive Power data), as shown at blocks 906, 910, 914. The data scanner may collect the operational data, as shown at block 926, and may send the data to the operational data store, as shown at block 928. The operational data store may store the operational data, as shown at block 938. The operational data store may further send a snapshot of the data to the historian, as shown at block 940, and the historian may store the snapshot of the data, as shown at block 942.

[00114] The meter state application may send a request for meter data to the Meter DCE, as shown in block 924, which in turn sends a request to one or more meters to collect meter data, as shown at block 920. In response to the request, the one or more meters collects meter data, as shown at block 916, and sends the voltage data to the Meter DCE, as shown at block 918. The Meter DCE may collect the voltage data, as shown at block 922, and send the data to the requestor of the data, as shown at block 928. The meter state application may receive the meter data, as shown at block 930, and determine whether it is for a single value process or a voltage profile grid state, as shown at block 932. If it is for the single value process, the meter data is send to the requesting process, as shown at block 936. If the meter data is for storage to determine the grid state at a future time, the meter data is stored in the operational data store, as shown at block 938. The operational data store further sends a snapshot of the data to the historian, as shown at block 940, and the historian stores the snapshot of the data, as shown at block 942.

[00115] Figures 9A-B further illustrate actions relating to demand response (DR). Demand response refers to dynamic demand mechanisms to manage customer consumption of electricity in response to supply conditions, for example, having

electricity customers reduce their consumption at critical times or in response to market prices. This may involve actually curtailing power used or by starting on site generation which may or may not be connected in parallel with the grid. This may be different from energy efficiency, which means using less power to perform the same tasks, on a

5 continuous basis or whenever that task is performed. In demand response, customers, using one or more control systems, may shed loads in response to a request by a utility or market price conditions. Services (lights, machines, air conditioning) may be reduced according to a preplanned load prioritization scheme during the critical timeframes. An alternative to load shedding is on-site generation of electricity to supplement the power
10 grid. Under conditions of tight electricity supply, demand response may significantly reduce the peak price and, in general, electricity price volatility.

[00116] Demand response may generally be used to refer to mechanisms used to encourage consumers to reduce demand, thereby reducing the peak demand for electricity. Since electrical systems are generally sized to correspond to peak demand
15 (plus margin for error and unforeseen events), lowering peak demand may reduce overall plant and capital cost requirements. Depending on the configuration of generation capacity, however, demand response may also be used to increase demand (load) at times of high production and low demand. Some systems may thereby encourage energy storage to arbitrage between periods of low and high demand (or low and high prices). As
20 the proportion of intermittent power sources such as wind power in a system grows, demand response may become increasingly important to effective management of the electric grid.

[00117] The DR state application may request the DR available capacity, as shown at block 954. The DR management system may then request available capacity from one
25 or more DR home devices, as shown at block 948. The one or more home devices may collect available DR capacity in response to the request, as shown at block 944, and send the DR capacity and response data to the DR management system, as shown at block 946. The DR management system may collect the DR capacity and response data, as shown at block 950, and send the DR capacity and response data to the DR state application, as
30 shown at block 952. The DR state application may receive the DR capacity and response data, as shown at block 956, and send the capacity and response data to the operational data store, as shown at block 958. The operational data store may store the DR capacity and response data, as shown at block 938. The operational data store may further send a

snapshot of the data to the historian, as shown at block 940, and the historian may store the snapshot of the data, as shown at block 942.

[00118] The substation computer may request application data from the substation application, as shown at block 974. In response, the substation application may request application from the substation device, as shown at block 964. The substation device may collect the application data, as shown at block 960, and send the application data to the substation device (which may include one, some or all of Voltage, Current, Real Power, and Reactive Power data), as shown at block 962. The substation application may collect the application data, as shown at block 966, and send the application data to the requestor (which may be the substation computer), as shown at block 968. The substation computer may receive the application data, as shown at block 970, and send the application data to the operational data store, as shown at block 972.

[00119] The grid state measurement and operational data process may comprise deriving the grid state and grid topology at a given point in time, as well as providing this information to other system and data stores. The sub-processes may include: (1) measuring and capturing grid state information (this relates to the operational data pertaining to the grid that was discussed previously); (2) sending grid state information to other analytics applications (this enables other applications, such as analytical applications, access to the grid state data); (3) persisting grid state snapshot to connectivity / operational data store (this allows for updating the grid state information to the connectivity/operational data store in the appropriate format as well as forwarding this information to the historian for persistence so that a point in time grid topology may be derived at a later point in time); (4) deriving grid topology at a point in time based on default connectivity and current grid state (this provides the grid topology at a given point in time by applying the point in time snapshot of the grid state in the historian to the base connectivity in the connectivity data store, as discussed in more detail below); and (5) providing grid topology information to applications upon request.

[00120] With regard to sub-process (4), the grid topology may be derived for a predetermined time, such as in real-time, 30 seconds ago, 1 month ago, etc. In order to recreate the grid topology, multiple databases may be used, and a program to access the data in the multiple databases to recreate the grid topology. One database may comprise a relational database that stores the base connectivity data (the “connectivity database”). The connectivity database may hold the grid topology information as built in order to

determine the baseline connectivity model. Asset and topology information may be updated into this database on a periodic basis, depending on upgrades to the power grid, such as the addition or modification of circuits in the power grid (*e.g.*, additional feeder circuits that are added to the power grid). The connectivity database may be considered “static” in that it does not change. The connectivity database may change if there are changes to the structure of the power grid. For example, if there is a modification to the feeder circuits, such as an addition of a feeder circuit, the connectivity database may change.

[00121] One example of the structure 1800 of the connectivity database may be derived from the hierarchical model depicted in Figures 18A-D. The structure 1800 is divided into four sections, with Figure 18A being the upper-left section, Figure 18B being the upper-right section, Figure 18C being the bottom-left section, and Figure 18D being the bottom-right section. Specifically, Figures 18A-D are an example of an entity relationship diagram, which is an abstract method to represent the baseline connectivity database. The hierarchical model in Figures 18A-D may hold the meta-data that describes the power grid and may describe the various components of a grid and the relationship between the components.

[00122] A second database may be used to store the “dynamic” data. The second database may comprise a non-relational database. One example of a non-relational database may comprise a historian database, which stores the time series non-operational data as well as the historical operational data. The historian database may stores a series of “flat” records such as: (1) time stamp; (2) device ID; (3) a data value; and (4) a device status. Furthermore, the stored data may be compressed. Because of this, the operation/non-operational data in the power grid may be stored easily, and may be manageable even though a considerable amount of data may be available. For example, data on the order of 5 Terabytes may be online at any given time for use in order to recreate the grid topology. Because the data is stored in the simple flat record (such as no organizational approach), it allows efficiency in storing data. As discussed in more detail below, the data may be accessed by a specific tag, such as the time stamp.

[00123] Various analytics for the grid may wish to receive, as input, the grid topology at a particular point in time. For example, analytics relating to power quality, reliability, asset health, etc. may use the grid topology as input. In order to determine the grid topology, the baseline connectivity model, as defined by the data in the connectivity

database, may be accessed. For example, if the topology of a particular feeder circuit is desired, the baseline connectivity model may define the various switches in the particular feeder circuit in the power grid. After which, the historian database may be accessed (based on the particular time) in order to determine the values of the switches in the particular feeder circuit. Then, a program may combine the data from the baseline connectivity model and the historian database in order to generate a representation of the particular feeder circuit at the particular time.

[00124] A more complicated example to determine the grid topology may include multiple feeder circuits (*e.g.*, feeder circuit A and feeder circuit B) that have an inter-tie switch and sectionalizing switches. Depending on the switch states of certain switches (such as the inter-tie switch and/or the sectionalizing switches), sections of the feeder circuits may belong to feeder circuit A or feeder circuit B. The program that determines the grid topology may access the data from both the baseline connectivity model and the historian database in order to determine the connectivity at a particular time (*e.g.*, which circuits belong to feeder circuit A or feeder circuit B).

[00125] Figure 10 illustrates a flow diagram 1000 of the Non-Operational Data processes. The non-operational extract application may request non-operational data, as shown at block 1002. In response, the data scanner may gather non-operational data, as shown at block 1004, where by various devices in the power grid, such as grid devices, substation computers, and line sensor RTUs, may collect non-operational data, as shown at blocks 1006, 1008, 1110. As discussed above, non-operational data may include temperature, power quality, etc. The various devices in the power grid, such as grid devices, substation computers, and line sensor RTUs, may send the non-operational data to the data scanner, as shown at blocks 1012, 1014, 1116. The data scanner may collect the non-operational data, as shown at block 1018, and send the non-operational data to the non-operational extract application, as shown at block 1020. The non-operational extract application may collect the non-operational data, as shown at block 1022, and send the collected non-operational data to the historian, as shown at block 1024. The historian may receive the non-operational data, as shown at block 1026, store the non-operational data, as shown at block 1028, and send the non-operational data to one or more analytics applications, as shown at block 1030.

[00126] Figure 11 illustrates a flow diagram 1100 of the Event Management processes. Data may be generated from various devices based on various events in the

power grid and sent via the event bus 147. For example, the meter data collection engine may send power outage/restoration notification information to the event bus, as shown at block 1102. The line sensors RTUs generate a fault message, and may send the fault message to the event bus, as shown at block 1104. The substation may analytics may generate a fault and/or outage message, and may send the fault and/or outage message to the event bus, as shown at block 1106. The historian may send signal behavior to the event bus, as shown at block 1108. And, various processes may send data via the event bus 147. For example, the fault intelligence process, discussed in more detail in Figures 14A-C, may send a fault analysis event via the event bus, as shown at block 1110. The outage intelligence process, discussed in more detail in Figures 13A-B, may send an outage event via the event bus, as shown at block 1112. The event bus may collect the various events, as shown at block 1114. And, the Complex Event Processing (CEP) services may process the events sent via the event bus, as shown at block 1120. The CEP services may process queries against multiple concurrent high speed real time event message streams. After processing by the CEP services, the event data may be sent via the event bus, as shown at block 1118. And the historian may receive via the event bus one or more event logs for storage, as shown at block 1116. Also, the event data may be received by one or more applications, such as the outage management system (OMS), outage intelligence, fault analytics, etc., as shown at block 1122. In this way, the event bus may send the event data to an application, thereby avoiding the “silo” problem of not making the data available to other devices or other applications.

[00127] Figures 12A-C illustrate a flow diagram 1200 of the Demand Response (DR) Signaling processes. DR may be requested by the distribution operation application, as shown at block 1244. In response, the grid state/connectivity may collect DR availability data, as shown at block 1202, and may send the data, as shown at block 1204. the distribution operation application may distribute the DR availability optimization, as show at block 1246, via the event bus (block 1254), to one or more DR Management Systems. The DR Management System may send DR information and signals to one or more customer premises, as shown at block 1272. The one or more customer premises may receive the DR signals, as shown at block 1266, and send the DR response, as shown at block 1268. The DR Management may receive the DR response, as shown at block 1274, and send DR responses to one, some or all of the operations data bus 146, the billing database, and the marketing database, as shown at block 1276. The

billing database and the marketing database may receive the responses, as shown at blocks 1284, 1288. The operations data bus 146 may also receive the responses, as shown at block 1226, and send the DR responses and available capacity to the DR data collection, as shown at block 1228. The DR data collection may process the DR responses and available capacity, as shown at block 1291, and send the data to the operations data bus, as shown at block 1294. The operations data bus may receive the DR availability and response, as shown at block 1230, and send it to the grid state/connectivity. The grid state/connectivity may receive the data, as shown at block 1208. The received data may be used to determine the grid state data, which may be sent (block 1206) via the operations data bus (block 1220). The distribution operation application may receive the grid state data (as an event message for DR optimization), as shown at block 1248. Using the grid state data and the DR availability and response, the distribution operation application may run distribution optimization to generate distribution data, as shown at block 1250. The distribution data may be retrieved by the operations data bus, as shown at block 1222, and may be sent to the connectivity extract application, as shown at block 1240. The operational data bus may send data (block 1224) to the distribution operation application, which in turn may send one or more DR signals to one or more DR Management Systems (block 1252). The event bus may collect signals for each of the one or more DR Management Systems (block 1260) and send the DR signals to each of the DR Management Systems (block 1262). The DR Management System may then process the DR signals as discussed above.

[00128] The communication operation historian may send data to the event bus, as shown at block 1214. The communication operation historian may also send generation portfolio data, as shown at block 1212. Or, an asset management device, such as a Ventyx[®], may request virtual power plant (VPP) information, as shown at block 1232. The operations data bus may collect the VPP data, as shown at block 1216, and send the data to the asset management device, as shown at block 1218. The asset management device may collect the VPP data, as shown at block 1234,, run system optimization, as shown at block 1236, and send VPP signals to the event bus, as shown at block 1238.

The event bus may receive the VPP signals, as shown at block 1256, and send the VPP signals to the distribution operation application, as shown at block 1258. The distribution operation application may then receive and process the event messages, as discussed above.

[00129] The connection extract application may extract new customer data, as shown at block 1278, to be sent to the Marketing Database, as shown at block 1290. The new customer data may be sent to the grid state/connectivity, as shown at block 1280, so that the grid state connectivity may receive new DR connectivity data, as shown at block 1210.

[00130] The operator may send one or more override signals when applicable, as shown at block 1242. The override signals may be sent to the distribution operation application. The override signal may be sent to the energy management system, as shown at block 1264, the billing database, as shown at block 1282, and/or the marketing database, as shown at block 1286.

[00131] Figures 13A-B illustrate a flow diagram 1300 of the Outage Intelligence processes. Various devices and applications may send power outage notification, as shown at blocks 1302, 1306, 1310, 1314, 1318. The outage events may be collected by the event bus, as shown at block 1324, which may send the outage events to the complex event processing (CEP), as shown at block 1326. Further, various devices and applications may send power restoration status, as shown at block 1304, 1308, 1312, 1316, 1320. The CEP may receive outage and restoration status messages (block 1330), process the events (block 1332), and send event data (block 1334). The outage intelligence application may receive the event data (block 1335) and request grid state and connectivity data (block 1338). The operational data bus may receive the request for grid state and connectivity data (block 1344) and forward it to one or both of the operational data store and the historian. In response, the operational data store and the historian may send the grid state and connectivity data (blocks 1352, 1354) via the operational data bus (block 1346) to the outage intelligence application (block 1340). It is determined whether the grid state and connectivity data indicate whether this was a momentary, as shown at block 1342. If so, the momentaries are sent via the operational data bus (block 1348) to the momentaries database for storage (block 1350). If not, an outage case is created (block 1328) and the outage case data is stored and processed by the outage management system (block 1322).

[00132] The outage intelligence processes may: detect outages; classify & log momentaries; determine outage extent; determine outage root cause(s); track outage restoration; raise outage events; and update system performance indices.

[00133] Figures 14A-C illustrate a flow diagram 1400 of the Fault Intelligence processes. The complex event processing may request data from one or more devices, as shown at block 1416. For example, the grid state and connectivity in response to the request may send grid state and connectivity data to the complex event processing, as shown at block 1404. Similarly, the historian in response to the request may send real time switch state to the complex event processing, as shown at block 1410. And, the complex event processing may receive the grid state, connectivity data, and the switch state, as shown at block 1418. The substation analytics may request fault data, as shown at block 1428. Fault data may be sent by a variety of devices, such as line sensor RTUs, and substation computers, as shown at blocks 1422, 1424. The various fault data, grid state, connectivity data, and switch state may be sent to the substation analytics for event detection and characterization, as shown at block 1430. The event bus may also receive event messages (block 1434) and send the event messages to the substation analytics (block 1436). The substation analytics may determine the type of event, as shown at block 1432. For protection and control modification events, the substation computers may receive a fault event message, as shown at block 1426. For all other types of events, the event may be received by the event bus (block 1438) and sent to the complex event processing (block 1440). The complex event processing may receive the event data (block 1420) for further processing. Similarly, the grid state and connectivity may send grid state data to the complex event processing, as shown at block 1406. And, the Common Information Model (CIM) warehouse may send meta data to the complex event processing, as shown at block 1414.

[00134] The complex event processing may send a fault event message, as shown at block 1420. The event bus may receive the message (block 1442) and send the event message to the fault intelligence application (block 1444). The fault intelligence application may receive the event data (block 1432) and request grid state, connectivity data, and switch state, as shown at block 1456. In response to the request, the grid state and connectivity send grid state and connectivity data (block 1408), and the historian send switch state data (block 1412). The fault intelligence receives the data (block 1458), analyzes the data, and sends event data (block 1460). The event data may be received by the event bus (block 1446) and sent to the fault log file (block 1448). The fault log file may log the event data (block 1402). The event data may also be received by the operational data bus (block 1462) and send to one or more applications (block 1464). For

example, the outage intelligence application may receive the event data (block 1466), discussed above with respect to Figures 13A-B. The work management system may also receive the event data in the form of a work order, as shown at block 1468. And, other requesting applications may receive the event data, as shown at block 1470.

5 **[00135]** The fault intelligent processes may be responsible for interpreting the grid data to derive information about current and potential faults within the grid. Specifically, faults may be detected using the fault intelligent processes. A fault is typically a short circuit caused when utility equipment fails or alternate path for current flow is created, for example, a downed power line. This processes may be used to detect typical faults
10 (typically handled by the conventional fault detection and protection equipment – relays, fuses, etc) as well as high impedance faults within the grid that are not easily detectable using fault currents.

[00136] The fault intelligence process may also classify and categorize faults. This allows for faults to be classified and categorized. Currently, no standard exists for a
15 systematic organization and classification for faults. A de-facto standard may be established for the same and implemented. The fault intelligence process may further characterize faults.

[00137] The fault intelligence may also determine fault location. Fault location in the distribution system may be a difficult task due to its high complexity and difficulty
20 caused by unique characteristics of the distribution system such as unbalanced loading, three-, two-, and single- phase laterals, lack of sensors/measurements, different types of faults, different causes of short circuits, varying loading conditions, long feeders with multiple laterals and network configurations that are not documented. This process enables the use a number of techniques to isolate the location of the fault with as much
25 accuracy as the technology allows.

[00138] The fault intelligence may further raise fault events. Specifically, this process may create and publish fault events to the events bus once a fault has been detected, classified, categorized, characterized and isolated. This process may also be responsible for collecting, filtering, collating and de-duplicating faults so that an
30 individual fault event is raised rather than a deluge based on the raw events that are typical during a failure. Finally, the fault intelligence may log fault events to the event log database.

[00139] Figures 15A-B illustrate a flow diagram 1500 of the Meta-data Management processes. Meta-data management processes may include: point list management; and communication connectivity & protocol management; and element naming & translation; sensor calibration factor management; and real time grid topology data management. The base connectivity extract application may request base connectivity data, as shown at block 1502. The Geographic Information Systems (GIS) may receive the request (block 1510) and send data to the base connectivity extract application (block 1512). The base connectivity extract application may receive the data (block 1504), extract, transform and load data (block 1506) and send base connectivity data to the connectivity data mart (block 1508). The connectivity data mart may thereafter receive the data, as shown at block 1514.

[00140] The connectivity data mart may comprise a custom data store that contains the electrical connectivity information of the components of the grid. As shown in Figures 15A-B, this information may be derived typically from the Geographic Information System (GIS) of the utility, which holds the as built geographical location of the components that make up the grid. The data in this data store describes the hierarchical information about all the components of the grid (substation, feeder, section, segment, branch, t-section, circuit breaker, recloser, switch, etc – basically all the assets). This data store may have the asset and connectivity information as built.

[00141] The meta data extract application may request meta data for grid assets, as shown at block 1516. The meta data database may receive the request (block 1524) and send meta data (block 1526). The meta data extract application may receive the meta data (block 1518), extract, transform and load meta data (block 1520), and send the meta data to the CIM data warehouse (block 1522).

[00142] The CIM (Common Information Model) data warehouse may then store the data, as shown at block 1528. CIM may prescribe utility standard formats for representing utility data. The INDE smart grid may facilitate the availability of information from the smart grid in a utility standard format. And, the CIM data warehouse may facilitate the conversion of INDE specific data to one or more formats, such as a prescribed utility standard format.

[00143] The asset extract application may request information on new assets, as shown at block 1530. The asset registry may receive the request (block 1538) and send information on the new assets (block 1540). The asset extract application may receive the

information on the new assets (block 1532), extract transform and load data (block 1534), and send information on the new assets to the CIM data warehouse (block 1536).

[00144] The DR connectivity extract application may request DR connectivity data, as shown at block 1542. The operational data bus may send the DR connectivity data request to the marketing database, as shown at block 1548. The marketing database may receive the request (block 1554), extract transform, load DR connectivity data (block 1556), and send the DR connectivity data (block 1558). The operational data bus may send the DR connectivity data to the DR connectivity extract application (block 1550). The DR connectivity extract application may receive the DR connectivity data (block 1544), and send the DR connectivity data (block 1546) via the operational data bus (block 1552) to the grid state and connectivity DM, which stores the DR connectivity data (block 1560).

[00145] Figure 16 illustrates a flow diagram 1600 of the Notification Agent processes. A notification subscriber may log into a webpage, as shown at block 1602.

The notification subscriber may create/modify/delete scenario watch list parameters, as shown at block 1604. The web page may store the created/modified/deleted scenario watch list, as shown at block 1608, and the CIM data warehouse may create a list of data tags, as shown at block 1612. A name translate service may translate the data tags for the historian (block 1614) and send the data tags (block 1616). The web page may send the data tag list (block 1610) via the operational data bus, which receives the data tag list (block 1622) and sends it to the notification agent (block 1624). The notification agent retrieves the list (block 1626), validates and merges the lists (block 1628), and checks the historian for notification scenarios (block 1630). If exceptions matching the scenarios are found (block 1632), a notification is sent (block 1634). The event bus receives the notification (block 1618) and sends it to the notification subscriber (block 1620). The notification subscriber may receive the notification via a preferred medium, such as text, e-mail, telephone call, etc., as shown at block 1606.

[00146] Figure 17 illustrates a flow diagram 1700 of the Collecting Meter Data (AMI) processes. The current collector may request residential meter data, as shown at block 1706. One or more residential meters may collect residential meter data in response to the request (block 1702) and send the residential meter data (block 1704). The current collector may receive the residential meter data (block 1708) and send it to the operational data bus (block 1710). The meter data collection engine may request

commercial and industrial meter data, as shown at block 1722. One or more commercial and industrial meters may collect commercial and industrial meter data in response to the request (block 1728) and send the commercial and industrial meter data (block 1730).

The meter data collection engine may receive the commercial and industrial meter data (block 1724) and send it to the operational data bus (block 1726).

[00147] The operational data bus may receive residential, commercial, and industrial meter data (block 1712) and send the data (block 1714). The data may be received by the meter data repository database (block 1716) or may be received by the billing processor (block 1718), which may in turn be sent to one or more systems, such as a CRM (customer relationship management) system (block 1720).

[00148] The observability processes may further include remote asset monitoring processes. Monitoring the assets within a power grid may prove difficult. There may be different portions of the power grid, some of which are very expensive. For example, substations may include power transformers (costing upwards of \$1 million), and circuit breakers. Oftentimes, utilities would do little, if anything, in the way of analyzing the assets and maximizing the use of the assets. Instead, the focus of the utility was typically to ensure that the power to the consumer was maintained. Specifically, the utility was focused on scheduled inspections (which would typically occur at pre-determined intervals) or “event-driven” maintenance (which would occur if a fault occurred in a portion of the grid).

[00149] Instead of the typical scheduled inspections or “event-driven” maintenance, the remote asset monitoring processes may focus on condition-based maintenance. Specifically, if one portion (or all) of the power grid may be assessed (such as on a periodic or continual basis), the health of the power grid may be improved.

[00150] As discussed above, data may be generated at various portions of the power grid and transmitted to (or accessible by) a central authority. The data may then be used by the central authority in order to determine the health of the grid. Apart from analyzing the health of the grid, a central authority may perform utilization monitoring. Typically, equipment in the power grid is operated using considerable safety margins. One of the reasons for this is that utility companies are conservative by nature and seek to maintain power to the consumer within a wide margin of error. Another reason for this is that the utility companies monitoring the grid may not be aware of the extent a piece of equipment in the power grid is being utilized. For example, if a power company is

transmitting power through a particular feeder circuit, the power company may not have a means by which to know if the transmitted power is near the limit of the feeder circuit (for example, the feeder circuit may become excessively heated). Because of this, the utility companies may be underutilizing one or more portions of the power grid.

5 [00151] Utilities also typically spend a considerable amount of money to add capacity to the power grid since the load on the power grid has been growing (*i.e.*, the amount of power consumed has been increasing). Because of the Utilities' ignorance, Utilities will upgrade the power grid unnecessarily. For example, feeder circuits that are not operating near capacity may nonetheless be upgraded by reconductoring (*i.e.*, bigger
10 wires are laid in the feeder circuits), or additional feeder circuits may be laid. This cost alone is considerable.

[00152] The remote asset monitoring processes may monitor various aspects of the power grid, such as: (1) analyzing current asset health of one or more portions of the grid; (2) analyzing future asset health of one or more portions of the grid; and (3) analyzing
15 utilization of one or more portions of the grid. First, one or more sensors may measure and transmit to remote asset monitoring processes in order to determine the current health of the particular portion of the grid. For example, a sensor on a power transform may provide an indicator of its health by measuring the dissolved gases on the transformer. The remote asset monitoring processes may then use analytic tools to determine if the
20 particular portion of the grid (such as the power transform is healthy or not healthy). If the particular portion of the grid is not healthy, the particular portion of the grid may be fixed.

[00153] Moreover, the remote asset monitoring processes may analyze data generated from portions of the grid in order to predict the future asset health of the
25 portions of the grid. There are things that cause stress on electrical components. The stress factors may not necessarily be constant and may be intermittent. The sensors may provide an indicator of the stress on a particular portion of the power grid. The remote asset monitoring processes may log the stress measurements, as indicated by the sensor data, and may analyze the stress measurement to predict the future health of the portion of
30 the power grid. For example, the remote asset monitoring processes may use trend analysis in order to predict when the particular portion of the grid may fail, and may schedule maintenance in advance of (or concurrently with) the time when the particular portion of the grid may fail. In this way, the remote asset monitoring processes may

predict the life of a particular portion of the grid, and thus determine if the life of that portion of the grid is too short (*i.e.*, is that portion of the grid being used up too quickly).

[00154] Further, the remote asset monitoring processes may analyze the utilization of a portion of the power grid in order to manage the power grid better. For example, the remote asset monitoring processes may analyze a feeder circuit to determine what its operating capacity is. In this feeder circuit example, the remote asset monitoring processes may determine that the feeder circuit is currently being operated at 70%. The remote asset monitoring processes may further recommend that the particular feeder circuit may be operated at a higher percentage (such as 90%), while still maintaining acceptable safety margins. The remote asset monitoring processes may thus enable an effective increase in capacity simply through analyzing the utilization.

Methodology for Determining Specific Technical Architecture

[00155] There are various methodologies for determining the specific technical architecture that may use one, some, or all of the elements of the INDE Reference Architecture. The methodology may include a plurality of steps. First, a baseline step may be performed in generating documentation of the as-is state of the utility, and a readiness assessment for transition to a Smart Grid. Second, a requirements definition step may be performed in generating the definition of the to-be state and the detailed requirements to get to this state.

[00156] Third, a solution development step may be performed in generating the definition of the solution architecture components that will enable the Smart Grid including the measurement, monitoring and control. For the INDE architecture, this may include the measuring devices, the communication network to pass data from the devices to the INDE CORE 120 applications, the INDE CORE 120 applications to persist and react to the data, analytical applications to interpret the data, the data architecture to model the measured and interpreted data, the integration architecture to exchange data and information between INDE and utility systems, the technology infrastructure to run the various applications and databases and the standards that may be followed to enable an industry standard portable and efficient solution.

[00157] Fourth, a value modeling may be performed in generating the definition of key performance indicators and success factors for the Smart Grid and the implementation of the ability to measure and rate the system performance against the

desired performance factors. The disclosure above relates to the Architecture development aspect of step 3.

[00158] Figures 19A-B illustrate an example of a blueprint progress flow graphic.

Specifically, Figures 19A-B illustrate a process flow of the steps that may be undertaken to define the smart grid requirements and the steps that may be executed to implement the smart grid. The smart grid development process may begin with a smart grid vision development, which may outline the overall goals of the project, that may lead to the smart grid roadmapping process. The roadmapping process may lead to blueprinting and to value modeling.

[00159] Blueprinting may provide a methodical approach to the definition of the smart grid in the context of the entire utility enterprise. Blueprinting may include an overall roadmap, which may lead to a baseline and systems evaluation (BASE) and to a requirements definition and analytics selection (RDAS). The RDAS process may create the detailed definition of the utility's specific smart grid.

[00160] The BASE process may establish the starting point for the utility, in terms of systems, networks, devices, and applications to support smart grid capabilities. The first part of the process is to develop a systems inventory of the grid, which may include: grid structure (such as generation, transmission lines, transmission substations, sub transmission lines, distribution substations, distribution feeders, voltage classes); grid devices (such as switches, reclosers, capacitors, regulators, voltage drop compensators, feeder inter-ties); substation automation (such as IEDs, substation LANs, instrumentation, station RTUs/computers); distribution automation (such as capacitor and switch control; fault isolation and load rollover controls; LTC coordination systems; DMS; Demand Response Management System); and grid sensors (such as sensor types, amounts, uses, and counts on distribution grids, on transmission lines and in substations); etc. Once the inventory is complete, an evaluation of the utility against a high level smart grid readiness model may be created. An as-is dataflow model and a systems diagram may also be created.

[00161] The architecture configuration (ARC) process may develop a preliminary smart grid technical architecture for the utility by combining the information from the BASE process, requirements and constraints from the RDAS process and the INDE Reference Architecture to produce a technical architecture that meets the specific needs of the utility and that takes advantage of the appropriate legacy systems and that conforms to

the constraints that exist at the utility. Use of the INDE Reference Architecture may avoid the need to invent a custom architecture and ensures that accumulated experience and best practices are applied to the development of the solution. It may also ensure that the solution can make maximum use of reusable smart grid assets.

- 5 **[00162]** The sensor network architecture configuration (SNARC) process may provide a framework for making the series of decisions that define the architecture of a distributed sensor network for smart grid support. The framework may be structured as a series of decision trees, each oriented to a specific aspect of sensor network architecture. Once the decisions have been made, a sensor network architecture diagram may be
- 10 created.
- [00163]** The sensor allocation via T-section recursion (SATSECTR) process may provide a framework for determining how many sensors should be placed on the distribution grid to obtain a given level of observability, subject to cost constraints. This process may also determine the sensor types and locations.
- 15 **[00164]** The solution element evaluation and components template (SELECT) process may provide a framework for evaluation of solution component types and provides a design template for each component class. The template may contain a reference model for specifications for each of the solution elements. These templates may then be used to request vendor quotations and to support vendor/product evaluations.
- 20 **[00165]** The upgrade planning for applications and networks (UPLAN) process may provide for development of a plan to upgrade of existing utility systems, applications, and networks to be ready for integration into a smart grid solution. The risk assessment and management planning (RAMP) process may provide an assessment of risk associated with specific elements of the smart grid solution created in the ARC process. The
- 25 UPLAN process may assesses the level or risk for identified risk elements and provides an action plan to reduce the risk before the utility commits to a build-out. The change analysis and management planning (CHAMP) process may analyze the process and organizational changes that may be needed for the utility to realize the value of the smart grid investment and may provide a high level management plan to carry out these changes
- 30 in a manner synchronized with the smart grid deployment. The CHAMP process may result in a blueprint being generated.
- [00166]** The roadmap in the value modeling process may lead to specifying value metrics, which may lead to estimation of cost and benefits. The estimation may lead to

the building of one or more cases, such as a rate case and business case, which in turn may lead to a case closure. The output of blueprinting and the value modeling may be sent to the utility for approval, which may result in utility system upgrades and smart grid deployments and risk reduction activities. After which, the grid may be designed, built and tested, and then operated.

Alternative INDE High Level Architecture Description

[00167] In one example, the overall INDE architecture may be applied to an industry including both mobile and stationary sensors. The INDE architecture may be implemented to receive sensor data and respond accordingly through both distributed and centralized intelligence. Figures 21 through 28 illustrate examples of the INDE architecture implemented in various vehicle travel industries.

Overall Architecture

[00168] Turning to the drawings, wherein like reference numerals refer to like elements, Figures 21A-C illustrate one example of the overall architecture for INDE. This architecture may serve as a reference model that provides for end to end collection, transport, storage, and management of network data related to a one or more particular industries. It may also provide analytics and analytics management, as well as integration of the foregoing into processes and systems. Hence, it may be viewed as an enterprise-wide architecture. Certain elements, such as operational management and aspects of the network itself, are discussed in more detail below.

[00169] The architecture depicted in Figures 21A-C may include up to four data and integration buses: (1) a high speed sensor data bus 2146 (which may include operational and non-operational data); (2) a dedicated event processing bus 2147 (which may include event data); (3) an operations service bus 2130 (which may serve to provide information about the network back office applications); and (4) an enterprise service bus for the back office IT systems (shown in Figures 1A-C as the enterprise integration environment bus 2114 for serving enterprise IT 2115). The separate data buses may be achieved in one or more ways. For example, two or more of the data buses, such as the high speed sensor data bus 2146 and the event processing bus 2147, may be different segments in a single data bus. Specifically, the buses may have a segmented structure or platform. As discussed in more detail below, hardware and/or software, such as one or more switches, may be used to route data on different segments of the data bus.

[00170] As another example, two or more of the data buses may be on separate buses, such as separate physical buses in terms of the hardware needed to transport data on the separate buses. Specifically, each of the buses may include cabling separate from each other. Further, some or all of the separate buses may be of the same type. For
5 example, one or more of the buses may comprise a local area network (LAN), such as Ethernet[®] over unshielded twisted pair cabling and Wi-Fi. As discussed in more detail below, hardware and/or software, such as a router, may be used to route data onto one bus among the different physical buses.

[00171] As still another example, two or more of the buses may be on different
10 segments in a single bus structure and one or more buses may be on separate physical buses. Specifically, the high speed sensor data bus 2146 and the event processing bus 2147 may be different segments in a single data bus, while the enterprise integration environment bus 2114 may be on a physically separate bus.

[00172] Though Figures 21A-C depict four buses, fewer or greater numbers of
15 buses may be used to carry the four listed types of data. For example, a single unsegmented bus may be used to communicate the sensor data and the event processing data (bringing the total number of buses to three), as discussed below. And, the system may operate without the operations service bus 2130 and/or the enterprise integration environment bus 2114.

[00173] The IT environment may be SOA-compatible. Service Oriented
20 Architecture (SOA) is a computer systems architectural style for creating and using business processes, packaged as services, throughout their lifecycle. SOA also defines and provisions the IT infrastructure to allow different applications to exchange data and participate in business processes. Although, the use of SOA and the enterprise service
25 bus are optional.

[00174] In an example of a generic industry, the figures illustrate different elements within the overall architecture, such as the following: (1) INDE CORE 2120; and (2) INDE DEVICE 2188. This division of the elements within the overall architecture is for illustration purposes. Other division of the elements may be used. And, the division of
30 elements may be different for different industries. The INDE architecture may be used to support both distributed and centralized approaches to intelligence, and to provide mechanisms for dealing with scale in large implementations.

[00175] The INDE Reference Architecture is one example of the technical architecture that may be implemented. For example, it may be an example of a *meta-architecture*, used to provide a starting point for developing any number of specific technical architectures, one for each industry solution (*e.g.*, different solutions for different industries) or one for each application within an industry (*e.g.*, a first solution for a first vehicle travel network and a second solution for a second vehicle travel network), as discussed below. Thus, the specific solution for a particular industry or a particular application within an industry (such as an application to a particular utility) may include one, some, or all of the elements in the INDE Reference Architecture. And, the INDE Reference Architecture may provide a standardized starting point for solution development. Discussed below is the methodology for determining the specific technical architecture for a particular industry or a particular application within an industry.

[00176] The INDE Reference Architecture may be an enterprise wide architecture. Its purpose may be to provide the framework for end to end management of data and analytics, and integration of these into systems and processes. Since advanced network technology affects every aspect of business processes, one should be mindful of the effects not just at the network level, operations, and customer premise levels, but also at the back office and enterprise levels. Consequently the INDE Reference Architecture can and does reference enterprise level SOA, for example, in order to support the SOA environment for interface purposes. This should not be taken as a requirement that an industry, such as, must convert their existing IT environment to SOA before the advanced network can be built and used. An enterprise service bus is a useful mechanism for facilitating IT integration, but it is not required in order to implement the rest of the solution. The discussion below focuses on different components of the INDE elements for vehicle travel; however, one, some, or all of the components of the INDE elements may be applied to different industries, such as telecommunication and energy exploration.

INDE Component Groups

[00177] As discussed above, the different components in the INDE Reference Architecture may include, for example: (1) INDE CORE 2120; and (2) INDE DEVICE 2188. The following sections discuss these example element groups of the INDE Reference Architecture and provide descriptions of the components of each group.

INDE CORE

[00178] Figure 22 illustrates the INDE CORE 2120, which is the portion of INDE Reference Architecture that may reside in an operations control center, as shown in Figures 21A-C. The INDE CORE 2120 may contain a unified data architecture for storage of data and an integration schema for analytics to operate on that data.

5 [00179] In addition, this data architecture may make use of federation middleware 2134 to connect other types of data (such as, for example, sensor data, operational and historical data, log and event files), and connectivity and meta-data files into a single data architecture that may have a single entry point for access by high level applications, including enterprise applications. Real time systems may also access key data stores via
10 the high speed data bus and several data stores can receive real time data. Different types of data may be transported within one or more buses in the INDE architecture.

[00180] The types of data transported may include operation and non-operational data, events, network connectivity data, and network location data. The operational and non-operational data may be transported using an operational/non-operational data bus
15 2146. Data collection applications may be responsible for sending some or all of the data to the operational/non-operational data bus 2146. In this way, applications that need this information may be able to get the data by subscribing to the information or by invoking services that may make this data available.

[00181] Events may include messages and/or alarms originating from the various
20 devices and sensors that are part of an industry network, as discussed below. Events may be directly generated from the devices and sensors on the as well as generated by the various analytics applications based on the measurement data from these sensors and devices.

[00182] As discussed in more detail below, data may be sent from various
25 components in the (such as INDE DEVICE 2188). The data may be sent to the INDE CORE 2120 wirelessly, wired, or a combination of both. The data may be received by utility communications networks 2160, which may send the data to routing device 2189. Routing device 2189 may comprise software and/or hardware for managing routing of data onto a segment of a bus (when the bus comprises a segmented bus structure) or onto
30 a separate bus. Routing device may comprise one or more switches or a router. Routing device 2189 may comprise a networking device whose software and hardware routes and/or forwards the data to one or more of the buses. For example, the routing device

2189 may route operational and non-operational data to the operational/non-operational data bus 2146. The routing device 2189 may also route event data to the event bus 2147.

[00183] The routing device 2189 may determine how to route the data based on one or more methods. For example, the routing device 2189 may examine one or more headers in the transmitted data to determine whether to route the data to the segment for the operational/non-operational data bus 2146 or to the segment for the event bus 2147. Specifically, one or more headers in the data may indicate whether the data is operation/non-operational data (so that the routing device 2189 routes the data to the operational/non-operational data bus 2146) or whether the data is event data (so that the routing device 2189 routes the event bus 2147). Alternatively, the routing device 2189 may examine the payload of the data to determine the type of data (e.g., the routing device 2189 may examine the format of the data to determine if the data is operational/non-operational data or event data).

[00184] One of the stores, such as the operational data warehouse 2137 that stores the operational data, may be implemented as true distributed database. Another of the stores, the historian (identified as historical data 2136 in Figures 21 and 22), may be implemented as a distributed database. Further, events may be stored directly into any of several data stores via the complex event processing bus. Specifically, the events may be stored in event logs 2135, which may be a repository for all the events that have published to the event bus 2147. The event log may store one, some, or all of the following: event id; event type; event source; event priority; and event generation time. The event bus 2147 need not store the events long term, providing the persistence for all the events.

[00185] The storage of the data may be such that the data may be as close to the source as possible or practicable. In one implementation, this may include, for example, the substation data being stored at the INDE DEVICE 2188. But this data may also be required at the operations control center level 2116 to make different types of decisions that consider the network at a much granular level. In conjunction with a distributed intelligence approach, a distributed data approach may be been adopted to facilitate data availability at all levels of the solution through the use of database links and data services as applicable. In this way, the solution for the historical data store (which may be accessible at the operations control center level 2116) may be similar to that of the operational data store. Historical/collective analytics may be performed at the operations

control center level 2116 by accessing data at INDE DEVICE level. Alternatively, data may be stored centrally at the INDE CORE 2120. However, given the amount of data that may need to be transmitted amongst the INDE DEVICES 2188, the storage of the data at the INDE DEVICES 2188 may be preferred. Specifically, if there are thousands or tens of thousands of sensors, the amount of data that needs to be transmitted to the INDE CORE 2120 may create a communications bottleneck.

[00186] Finally, the INDE CORE 2120 may program or control one, some or all of the INDE DEVICE 2188 in the network. For example, the INDE CORE 2120 may modify the programming (such as download an updated program) or provide a control command to control any aspect of the INDE DEVICE 2188 (such as control of the sensors or analytics). Other elements, not shown in Figure 2, may include various integration elements to support this logical architecture.

[00187] Table 2 describes the certain elements of INDE CORE 2120 as depicted in Figure 21.

INDE CORE Element	Description
CEP Services 2144	Provides high speed, low latency event stream processing, event filtering, and multi-stream event correlation
Centralized Analytics Applications 2139	May consist of any number of commercial or custom analytics applications that are used in a non-real time manner, primarily operating from the data stores in INDE CORE.
Visualization/Notification Services 2140	Support for visualization of data, states and event streams, and automatic notifications based on event triggers
Application Management Services 2141	Services (such as Applications Support Services 2142 and Distributed Computing Support 2143) that support application launch and execution, web services, and support for distributed computing and automated remote program download (e.g., OSGi)
Network Management Services 2145	Automated monitoring of communications networks, applications and databases; system health monitoring, failure root cause analysis.
Meta-Data Services 2126	Services (such as Connectivity Services 2127, Name Translation 2128, and TEDS Service 2129) for storage, retrieval, and update of system meta-data, including communication/sensor net connectivity, point lists, sensor calibrations, protocols, device set points, etc
Analytics Services 2123	Services (such as Sensor Data Services 2124 and Analytics Management Services 2125) to support access to sensor data and sensor analytics; management of analytics.

Sensor Data Management System 2121	Sensor data management system functions.
Real Time Complex Event Processing Bus 2147	Message bus dedicated to handling event message streams – purpose of a dedicated bus is to provide high bandwidth and low latency for highly bursty event message floods. The event message may be in the form of XML message. Other types of messages may be used. Events may be segregated from operational/non-operational data, and may be transmitted on a separate or dedicated bus. Events typically have higher priority as they usually require some immediate action from an operational perspective (messages from defective vehicle equipment). The event processing bus (and the associated event correlation processing service depicted in Figure 21) may filter floods of events down into an interpretation that may better be acted upon by other devices. In addition, the event processing bus may take multiple event streams, find various patterns occurring across the multiple event streams, and provide an interpretation of multiple event streams. In this way, the event processing bus may not simply examine the event data from a single device, instead looking at multiple device (including multiple classes of devices that may be seemingly unrelated) in order to find correlations. The analysis of the single or multiple event streams may be rule based
Real Time Op/Non-Op Data Bus 2146	Operational data may include data reflecting the current state of the particular industry network. Non-operational data may include data reflecting the “health” or condition of a device. Operational data has previously been transmitted directly to a specific device (thereby creating a potential “silo” problem of not making the data available to other devices or other applications). However, using the bus structure, the operational data may also be used for asset utilization/optimization, system planning, etc.,. Non-operational data was previously obtained by sending a person in the field to collect the operational data (rather than automatically sending the non-operational data to a central repository). Typically, the operational and non-operational data are generated in the various devices in the grid at predetermined times. This is in contrast to the event data, which typically is generated in bursts, as

	discussed below. A message bus may be dedicated to handling streams of operational and non-operational data from substations and grid devices. The purpose of a dedicated bus may be to provide constant low latency through put to match the data flows; as discussed elsewhere, a single bus may be used for transmission of both the operation and non-operational data and the event processing data in some circumstances (effectively combining the operation/non-operational data bus with the event processing bus).
Operations Service Bus 2130	Message bus that supports integration of typical industry operations applications The operations service bus 2130 may serve as the provider of information about the smart grid to the utility back office applications, as shown in Figure 21. The analytics applications may turn the raw data from the sensors and devices on the grid into actionable information that will be available to utility applications to perform actions to control the grid. Although most of the interactions between the utility back office applications and the INDE CORE 2120 is expected to occur thru this bus, utility applications will have access to the other two buses and will consume data from those buses as well (for example, meter readings from the op/non-op data bus 2146, outage events from the event bus 2147)
Sensor Data Warehouse 2133	The sensor data warehouse 2133 may provide rapid access to sensor usage data for analytics. This repository may hold all the sensor reading information from the sensors. The data collected from the sensors may be stored in sensor data warehouse 2133 and provided to other applications,) as well as other analysis.
Event Logs 2135	Collection of log files incidental to the operation of various industry systems. The event logs 2135 may be used for post mortem analysis of events and for data mining
Historical Data 2136	Telemetry data archive in the form of a standard data historian. Historical data 2136 may hold the time series non-operational data as well as the historical operational data. Analytics pertaining to items like reliability, asset health, etc may be performed using data in historical data 2136.
Operational Data 2137	Operational Data 2137 may comprise a real time operational database. Operational Data 2137 may be built in true distributed form. Specifically, the

	Operational Data 2137 may hold data measurements obtained from the sensors and devices. Historical data measurements are not held in this data store, instead being held in historical data 2136. The data base tables in the Operational Data 2137 may be updated with the latest measurements obtained from these sensors and devices.
--	--

Table 2: INDE CORE Elements

[00188] As discussed in Table 2, the real time data bus 2146 (which communicates the operation and non-operational data) and the real time complex event processing bus 2147 (which communicates the event processing data) into a single bus 2346. An example of this is illustrated in the block diagram 2300 in Figures 23A-C.

[00189] As shown in Figures 21A-C, the buses are separate for performance purposes. For CEP processing, low latency may be important for certain applications which are subject to very large message bursts. Most of the grid data flows, on the other hand, are more or less constant, with the exception of digital fault recorder files, but these can usually be retrieved on a controlled basis, whereas event bursts are asynchronous and random.

[00190] Figure 21 further shows additional elements in the operations control center 2116 separate from the INDE CORE 2120. Specifically, Figure 21 further shows Sensor Data Collection Head End(s) 2153, a system that is responsible for communicating with meters (such as collecting data from them and providing the collected data to the utility). IP Network Services 2158 is a collection of services operating on one or more servers that support IP-type communications (such as DHCP and FTP). Dispatch Mobile Data System 2159 is a system that transmits/receives messages to mobile data terminals in the field. Work Management System 2150 is a system that monitors and manages work orders. Geographic Information System 2149 is a database that contains information about where assets are located geographically and how the assets are connected together. If the environment has a Services Oriented Architecture (SOA), Operations SOA Support 2148 is a collection of services to support the SOA environment.

[00191] One or more of the systems in the operations control center 2116 that are outside of the INDE Core 2120 are legacy product systems that a utility may have. Examples of these legacy product systems include the Operations SOA Support 2148, Sensor Data Collection Head End(s) 2153, IP Network Services 2158, and Dispatch Mobile Data System 2159. However, these legacy product systems may not be able to process or handle data that is received from a smart grid. The INDE Core 2120 may be

able to receive the data from the smart grid, process the data from the smart grid, and transfer the processed data to the one or more legacy product systems in a fashion that the legacy product systems may use (such as particular formatting particular to the legacy product system). In this way, the INDE Core 2120 may be viewed as a middleware.

5 [00192] The operations control center 2116, including the INDE CORE 2120, may communicate with the Enterprise IT 2115. Generally speaking, the functionality in the Enterprise IT 2115 comprises back-office operations. Specifically, the Enterprise IT 2115 may use the enterprise integration environment bus 2114 to send data to various systems within the Enterprise IT 2115, including Business Data Warehouse 2104,
10 Business Intelligence Applications 2105, Enterprise Resource Planning 2106, various Financial Systems 2107, Customer Information System 2108, Human Resource System 2109, Asset Management System 2110, Enterprise SOA Support 2111, Network Management System 2112, and Enterprise Messaging Services 2113. The Enterprise IT 2115 may further include a portal 2103 to communicate with the Internet 2101 via a
15 firewall 2102.

INDE DEVICE

[00193] The INDE DEVICE 2188 group may comprise any variety of devices used to provide data associated with a particular device. In one example, the device group 2188 may include stationary sensor units 2190 and mobile sensor units 2192. Each
20 stationary sensor unit 2190 and mobile sensor unit 2192 may include one or more sensors, processors, memory devices, communication modules, and/or power modules allowing receipt of any data from the sensors, as well as subsequent processing and/or transmission of raw or processed sensor data. Raw or processed sensor data from the stationary sensor units 2190 and mobile sensor units 2192 may be processed by one or more gateways
25 2194. In one example, each gateway 2194 may be one or more devices capable of encoding and transmitting data to an operations control center 2116. Raw or processed sensor data from the stationary sensor units 2190 and the mobile sensor units 2192 may also be provided to a data collector 2196. The data collector 2196 may include one or more processors, memory devices, communication modules, and power modules. The
30 data collector 2196 may be a memory device and processor configured to collect, store, and transmit data. The data collector 2196 may communicate with the stationary sensor units 2190 and the mobile sensor units 2192 to collect data and transmit collected data to one or more gateways 2194.

[00194] In one example, the stationary sensor units 2190 may detect conditions associated with one or more of the mobile sensor units 2192 or other stationary sensor units 2190. The mobile sensor units 2192 may detect conditions associated with the stationary sensor units 2192 or may detect other conditions associated with other mobile sensor units 2192. During operation, event data may be generated by the stationary sensor units 2190 and the mobile sensor units 2192. The event data may be indicative of abnormal or undesired conditions of a vehicle travel network. Such event data may be transmitted from the stationary sensor units 2190 and the mobile sensor units 2192 through the gateways 2194 to the central authority. In one example, event data may be received by a routing device 2189. The event data may be provided to the event bus 2147 by the routing device 2189. The received event data may be processed by the operations control center 2116 to allow an appropriate response to be generated.

[00195] Figures 24A-24C is a block diagram of the INDE architecture to operate with a rail travel network. The INDE system of Figures 24A-24B may receive event data from stationary sensor units 2190 and mobile sensor units 2192 positioned on rail cars 2400 as shown in FIG. 25. In one example, the stationary sensor units 2190 and mobile sensor units 2192 may be those disclosed in United States Patent Publication No. 2009/0173840, which is incorporated by reference herein.

[00196] Referring to FIG. 25, in one example, a freight train 2500 may include rail cars 2400 of various types such as box cars, cabooses, coal cars, engine cars, and any other car configured to be conveyed via rail. The engine car 2502 may be powered by a diesel engine, battery, flywheel, fuel cell or any combination thereof. Each rail car 2400 may include one or more mobile sensor units 2192. The mobile sensor units 2192 may communicate with one other allowing communication amongst mobile sensor units 2192 of the same rail car 2400 or different rail cars 2400 attached to the same string of rail cars 2400 or other rail cars 2400 (not shown) detached from the string, such as those located in a train yard. Each mobile sensor unit 2192 may have a unique ID and each particular rail car 2400 may have a unique ID maintained by each mobile sensor unit 2192 associated with the particular rail car 2400. The ID's may be provided via RFID, for example.

[00197] In one example, the stationary sensor units 2190A may be configured to act as a "hot box" detector configured to monitor heating associated with a rail car wheels, axles, etc. The term "hot box," as is known in the art, may refer to a rail car experiencing overheating at one or more axle bearings and/or other wheel-based component on a piece

of railway rolling stock. Stationary sensor units 2190A may be placed along railroad tracks 2501. Each stationary sensor unit 2190A may be fitted with one or more infrared (IR) sensors to determine heating patterns of the bearings/axles/wheels of the rail cars 2400 as the rail cars pass through a sensing zone of a particular stationary sensor unit 2190A. Abnormal heating may indicate various issues such as rail car load imbalance, rail car structural issues, track issues, etc. If an overheated bearing is detected, a type of alarm can be triggered to alert the engineer to stop the train and correct the potentially dangerous situation which, if allowed to continue, could result in a train derailment. An example of a hot box detector is disclosed in U.S. Patent No. 4,659,043, which is hereby incorporated by reference. The stationary sensor units 2190A may be configured to process the IR sensor data to generate event data based on the alarm, such as event messages, to be received by the event bus 2147 for subsequent processing.

[00198] The stationary sensor units 2192B may also serve as a defect detector. A defect detector may be a device used on railroads to detect axle and signal problems in passing trains. The defect detectors can be integrated into the rail tracks and may include sensors to detect one or more kinds of problems that could occur. Defect detectors enable railroads to eliminate the caboose at the rear of the train, as well as various station agents stationed along the route to detect unsafe conditions. The defect detector may be integrated or associated with a wired or wireless transmitter. As trains pass the defect detectors, the defect detector may output the railroad name, milepost or location, track number (if applicable), number of axles on the train that passed and the indication “no defects” to indicate that no problems were detected on the train. Further, the location’s ambient temperature and train speed may be output. When a problem is detected, the detector may output an alarm indication, followed by a description of the problem and the axle position within the train where the problem occurred.

[00199] The stationary sensor units 2190C may also be configured to act as a “silver boxes,” as is known in the art, configured to receive raw or processed data received by one or more stationary sensor units 2190A and 2190B. The stationary sensor units 2192C may receive data from a respective group of stationary sensor units 2190A based on various common factors, such as geographic location, for example. In this regard, the stationary sensory units 2190C may be act as a data collector 2196 as shown in Figures 21A-21C.

[00200] During operation, a train 2500 having a string of rail cars 2400 may travel along rail tracks 2502. As the train 2500 travels, the stationary sensor units 2190A may detect information regarding each rail car 2400, such as bearing temperature. Each stationary sensor unit 2190A may also communicate with each mobile sensor unit 2192.

5 Communication may allow each stationary sensor unit 2190A to perform a health check of each rail car 2400 and associated mobile sensor units 2192. Any indication of abnormal or undesired conditions associated with a particular rail car 2400 may be relayed to the stationary sensor units 2190C. Conditions detected may relate to rail car structure, rail car environment (e.g., temperature), rail car content (e.g., weight,
10 distribution, quantity, etc.), rail car motion, rail car position, or any other parameter of interest regarding a rail car 2400. The conditions detected may also relate to security, such as when a rail car door is opened, which may indicate attempted theft or vandalism. The event data 2508 may be used to alert a particular organization that may own a particular rail car. Thus, the operations control center 2116 may oversee an entire railway
15 network, but companies owning individual rail cars 2400 may be alerted when an event data is transmitted regarding a particular rail car(s) 2400 owned by a particular company. Alert messages may be provided via an interface, subscription service, email, text message and/or any other communication manners capable of providing such alerts.

[00201] In one example, one of the rail cars 2400, such as the engine 2502, may
20 have a mobile sensor unit 2192 serving as a master mobile sensor unit 2504 to receive data from each mobile stationary unit 2192 associated with the rail cars 2400 of the current rail car string. When rail cars 2400 are connected to form a particular train, each mobile sensor unit 2192 may register with the master mobile sensor unit 2504. The master mobile sensor unit 2504 may receive periodic or continuous streams of raw or
25 processed data from the mobile sensor units 2192. This allows the engineer to determine the health of each rail car 2400 during use.

[00202] In one example, each mobile sensor unit 2190 may include a global positioning system (GPS) module, allowing each individual mobile sensor unit 2192 to determine a respective geographic location. Each mobile sensor unit 2190 may receive
30 GPS signals 2506 to determine geographic location. This information may be relayed to the stationary sensor units 2192A when a particular rail car is within proximity to allow transmission of such information. Each mobile sensor unit 2192 may, automatically or upon request, relay the geographic position to the master mobile sensor unit 2504, which

may be subsequently relayed through an onboard gateway 2194 to the operations control center 2116. In one example, the gateway 2194 may be that described in United States Patent Publication No. 2009/0173840. Each mobile sensor unit 2192A may also wirelessly transmit a GPS signal, allowing each rail car 2400 to be individually tracked.

5 Such an arrangement may allow an entire train to be tracked when only a single rail car has clear access to GPS satellites, such as when a train is moving through a tunnel.

[00203] Each of the stationary sensor units 2190 and the mobile sensor units 2192 may at least a part of the sensor data and generate an “event”. The event may comprise a no-incident event or an incident event. The no-incident event may indicate that there is
10 no incident (such as no fault) to report. The incident event may indicate that an incident has occurred or is occurring, such as a fault that has occurred or is occurring in the section of the rail travel network.

[00204] The stationary sensor units 2190 and the mobile sensor units 2192 may include one or both of: (1) intelligence to determine whether there is an incident; and (2)
15 ability to take one or more actions based on the determination whether there is an incident. In particular, the memory in one or more of the stationary sensor units 2190 and the mobile sensor units 2192 may include one or more rules to determine different types of incidents based on the data generated by one or more sensors. Or, the memory in the stationary sensor units 2190 and the mobile sensor units 2192 may include one or more
20 look-up tables to determine different types of incidents based on the data generated by one or more sensors. Further, the stationary sensor units 2190 and the mobile sensor units 2192 may include the ability to take one or more actions based on the determination whether there is an incident.

[00205] Apart from working alone, the electronic elements in the rail travel network
25 may work together as part of the distributed intelligence of the rail travel network. For example, stationary sensor units 2190 and the mobile sensor units 2192 may share data or share processing power to jointly determine whether there is an incident and take one or more actions based on the determination whether there is an incident.

[00206] The actions include, but are not limited to: (1) sending the incident
30 determination on the event bus 2147; (2) sending the incident determination along with a recommended action on the event bus 2147; and (3) taking an action to modify the state of one or more sections of the rail travel network or one or more vehicles traveling on the rail travel network. For example, the stationary sensor units 2190 may control one or

more switches in a section of the rail travel network (such as redirecting traffic onto a separate rail line, open a lane for travel in different directions, etc.). Or, the stationary sensor units 2190 may modify the parameters of one or more sensors in a section of the rail travel network (such as command the sensors to be more sensitive in its readings, command the sensors to generate more frequent readings, etc.). As still another example, the stationary sensor units 2190 and the mobile sensor units 2192 may control one or more vehicles traveling on the rail travel network. For example, a locomotive may include remote command-control capability, whereby the locomotive may be capable of receiving a wirelessly transmitted command to control one or more aspects of the locomotive. The one or more aspects that the command controls may include, but are not limited to, speed of the locomotive, generating a whistle (or other type of noise), and generating a light (or other type of visual indication). The receiver of the locomotive may receive the command and the processor of the locomotive may control one or more aspects of the locomotive based on the command (such as modifying operation of the engine).

[00207] The rail travel network may include distributed intelligence. As discussed above, different stationary sensor units 2190 and the mobile sensor units 2192 within the rail travel network may include additional functionality including additional processing/analytical capability and database resources. The use of this additional functionality within various stationary sensor units 2190 and the mobile sensor units 2192 in the rail travel network enables distributed architectures with centralized management and administration of applications and network performance. For functional, performance, and scalability reasons, a rail travel network involving thousands of stationary sensor units 2190 and the mobile sensor units 2192 may include distributed processing, data management, and process communications.

[00208] Non-operational data and operational data may be associated with and proximate to the stationary sensor units 2190 and the mobile sensor units 2192. The stationary sensor units 2190 and the mobile sensor units 2192 may further include components of the rail travel network that are responsible for the observability of the rail travel network at various sections. The stationary sensor units 2190 and the mobile sensor units 2192 may provide three primary functions: operational data acquisition and storage in the distributed operational data store; acquisition of non-operational data and storage in the historian; and local analytics processing on a real time (such as a sub-

second) basis. Processing may include digital signal processing, detection and classification processing, including event stream processing; and communications of processing results to local systems and devices as well as to systems at the operations control center 2116. Communication between the stationary sensor units 2190 and the mobile sensor units 2192 and other devices in the rail travel network may be wired, wireless, or a combination of wired and wireless. The electronic element may transmit data, such as operation/non-operational data or event data, to the operations control center 2116. A routing device may route the transmitted data to one of the operational/non-operational data bus or the event bus.

[00209] One or more types of data may be duplicated at the electronic element and at the operations control center 2116, thereby allowing an electronic element to operate independently even if the data communication network to the operations control center 2116 is not functional. With this information (connectivity) stored locally, analytics may be performed locally even if the communication link to the operations control center 2116 is inoperative.

[00210] Similarly, operational data may be duplicated at the operations control center 2116 and at the electronic elements. Data from the sensors and devices associated with a particular electronic element may be collected and the latest measurement may be stored in this data store at the electronic element. The data structures of the operational data store may be the same and hence database links may be used to provide seamless access to data that resides on the electronic element thru the instance of the operational data store at the operations control center 2116. This provides a number of advantages including alleviating data replication and enabling data analytics, which is more time sensitive, to occur locally and without reliance on communication availability beyond the electronic element. Data analytics at the operations control center 2116 level may be less time sensitive (as the operations control center 2116 may typically examine historical data to discern patterns that are more predictive, rather than reactive) and may be able to work around network issues if any.

[00211] Finally, historical data may be stored locally at the electronic element and a copy of the data may be stored at the operations control center 2116. Or, database links may be configured on the repository instance at the operations control center 2116, providing the operations control center access to the data at the individual electronic elements. Electronic element analytics may be performed locally at the electronic

element using the local data store. Specifically, using the additional intelligence and storage capability at the electronic element enables the electronic element to analyze itself and to correct itself without input from a central authority.

[00212] Alternatively, historical/collective analytics may also be performed at the operations control center 2116 level by accessing data at the local electronic element instances using the database links.

[00213] Further, various analytics may be used to analyze the data and/or the events. One type of analytics may comprise spatial visualization. Spatial visualization ability or visual-spatial ability is the ability to manipulate 2-dimensional and 3-dimensional figures. Spatial visualization may be performed using one or more electronic elements, or may be performed by the central authority. Further, spatial visualization may be used with a variety of industry networks, including utility networks and vehicle travel networks.

[00214] In one example, during operation, event data 2508 may be produced by each mobile sensor unit 2192. The event data 2508 may be transmitted to the master mobile sensor unit 2504. The master mobile sensor unit 2504 may transmit the event data wirelessly from the gateway to the operations control center 2116 for processing via a gateway 2194. In alternative examples, each rail car 2400 may include a respective gateway 2194 allowing data to be transmitted directly from the mobile sensor unit of the rail car 2400. This allows rail cars 2400 to communicate when not linked with an engine 2502, such as those being stored in a train yard, to communicate event data 2508 to be received by the operations control center 2116. In other alternative examples, each train yard may have one or more stationary sensor units 2190 and gateway 2194 to communicate with stored rail cars 2400 and to transmit any event data 2508. In a similar fashion, the stationary sensor units 2190A and 2190B may transmit sensor data to the stationary sensor and event data through a similar gateway(s) 2194 to relay such information to the operations control center 2116.

[00215] As described above, the network of Figures 24A-4C may also allow distributed analysis such that event data 2508 is processed at the stationary sensor units 2190A and 2190B and at the master mobile sensor module 2504. Such processing may allow for any issues to be analyzed and provide a solution or course of action. Such issue solution may be conveyed may be used to automatically control the train 2500 any in any capacity as the train 2500 is configured to allow or may alert human operators to control

the train 2500 accordingly. The solution may also be conveyed to the operations control center 2116 allowing the operations control center 2116 to perform actions remotely in order to confirm the issue solution and implement the solution accordingly.

[00216] FIG. 26A-26C is a block diagram of an implementation of the INDE

architecture related to an electric train network, such as a commuter train. An electric train network may include one or more electric trains that may be powered overhead electric lines or third rail. In one example, an electric train 2600 may include one or more cars 2602. Each car may be individually powered by an external source (e.g., third rail or overhead lines) or internal sources (e.g., battery or fuel cell). Each car may include one or more mobile sensor units 2192. Each mobile sensor unit 2192 may detect various conditions associated with various predetermined parameters of the train 2600.

[00217] In the example shown in Figures 26A-26C, the electric train 2600 may be powered by a third rail 2604. The third rail 2604 may be connected to one or more stationary sensor modules 2192 that may monitor the power flowing through the third rail 2604. The stationary sensor modules 2190 may determine the health of the rail system and transmit any events related to abnormal, undesired, or status check in the form of event messages. The event messages may be transmitted by a gateway 2194 to be received by the operations control center 2116.

[00218] Each electric rail car 2602 may include one or more mobile sensor units 2192. One of the mobile sensor units 2192 may serve as a master mobile sensor unit, such as the master mobile sensor unit 2504. The mobile sensor units 2192 may accumulate information regarding the respective rail car in a fashion similar to that discussed with regard to Figures 25A-5B. The master mobile sensor unit for the electric train 2600 may transmit event messages generated by the other mobile sensor units 2192 to the central authority via the gateway 2120.

[00219] Figures 27A-27C is a block diagram of an implementation of the INDE architecture applied to a road-based cargo transport network, such as that used in the trucking industry. In one example, one or more mobile sensor units 2192 may be include in cargo containers 2700 such as those hauled by diesel-engine trucks 2704. Each mobile sensor unit 2192 may be similar to that of the mobile sensor units 2192 discussed with regard to Figures 25A-26C. Each mobile sensor unit 2192 may detect various conditions of the cargo containers 2700 and relay them via an onboard gateway 2194 to the operations control center 2116. Stationary sensor units 2190 may be distributed at

customer facilities allowing cargo to be checked using communication between the stationary sensor units 2190 and the mobile sensor units 2192. The mobile sensor units 2192 may be used for cargo tracking, cargo container environment, theft/vandalism detection and any other appropriate uses as described with regard to Figures 24A-25.

5 [00220] Figures 28A-28C is a block diagram of an implementation of the overall INDE architecture applied to a network of automobiles that may be petroleum-fueled, electric, hybrid-fueled, bio-fueled or fueled by any other suitable manner. In one example, a vehicle 2800 may include one or more mobile sensor units 2192 allowing various conditions of the vehicle 2800 to be monitored. Each vehicle may include a gateway 2194 or communicate through an external gateway 2194 to communicate event data directly to INDE core 120 or other mobile sensor units 2192. Similar to other examples discussed, the mobile sensor units 2192 may include distributed intelligence that may perform analytics involved with the vehicle or may do so through interaction with INDE core 2120. Stationary sensor units 2190 may be used to communicate with the mobile sensor units 2192 allowing evaluation of a vehicle 800 having a mobile sensor unit 2192 in proximity to communicate with stationary sensor units 190. The stationary sensor units 2190 may include or share a gateway 194 allowing event data to be transmitted to the INDE CORE 2120 or may transmit it directly to the INDE CORE 120. The stationary sensor units 2190 may be implemented by car rental companies, car sales lots, or individual owners to receive event data associated with the condition and/or location of a vehicle 2800.

[00221] Figure 30 shows an example of the INDE systems 2000 that may be remotely hosted, as the block diagram illustrates. At a hosting site 3000, network cores 2002 may be installed as needed to support INDS subscribers for a particular industry. In one example, as subscriber 3002 may require use of various industries, such as rail, trucking, and airline. An INDE system 2000 may be modular allowing more industry types to be added, or in alternative examples a new subscriber. A party separate from the electric utility may manage and support the software for one, some, or all of the INDE systems 2000, as well as the applications that are downloaded from the INDS hosting site to be used for system endpoints 2006 and system infrastructure 2008. In order to facilitate communications, high bandwidth low latency communications services, such as via network 3004 (e.g., a MPLS or other WAN), may be use that can reach the subscriber utility operations centers, as well as the INDS hosting sites.

-67-

[00222] While this invention has been shown and described in connection with the preferred embodiments, it is apparent that certain changes and modifications in addition to those mentioned above may be made from the basic features of this invention. In addition, there are many different types of computer software and hardware that may be utilized in practicing the invention, and the invention is not limited to the examples described above. The invention was described with reference to acts and symbolic representations of operations that are performed by one or more electronic devices. As such, it will be understood that such acts and operations include the manipulation by the processing unit of the electronic device of electrical signals representing data in a structured form. This manipulation transforms the data or maintains it at locations in the memory system of the electronic device, which reconfigures or otherwise alters the operation of the electronic device in a manner well understood by those skilled in the art. The data structures where data is maintained are physical locations of the memory that have particular properties defined by the format of the data. While the invention is described in the foregoing context, it is not meant to be limiting, as those of skill in the art will appreciate that the acts and operations described may also be implemented in hardware. Accordingly, it is the intention of the Applicants to protect all variations and modification within the valid scope of the present invention. It is intended that the invention be defined by the following claims, including all equivalents.

[00223] Throughout this specification and claims which follow, unless the context requires otherwise, the word "comprise", and variations such as "comprises" and "comprising", will be understood to imply the inclusion of a stated integer or step, or group of integers or steps, but not the exclusion of any other integer or step or group of integers or steps.

-67a-

[00224] The reference to any prior art in this specification is not, and should not be taken as, an acknowledgement or any suggestion that the prior art forms part of the common general knowledge in Australia.

THE CLAIMS DEFINING THE INVENTION ARE AS FOLLOWS:

1. An integration framework for facilitating communication with a central authority that manages an industry network, the integration framework
5 including:
 - a first bus for communicating operational data to the central authority, the operational data including a real-time measurement for at least a part of the industry network; and
 - a second bus for communicating event data to the central authority,
10 the second bus being separate from the first bus, the event data being distinct from and derived from the real-time measurement and including at least one analytical determination based on the at least one real time measurement,
 - wherein the operational data is communicated via the first bus and
15 not communicated via the second bus, and
 - wherein the event data is communicated via the second bus and not communicated via the first bus.
2. An integration framework according to claim 1, wherein the industry
20 network includes a vehicle travel network.
3. An integration framework according to either claim 1 or claim 2, wherein the vehicle travel network includes a rail travel network.
- 25 4. An integration framework according to any one of the preceding claims, wherein the first bus includes a first segment and the second bus includes a second segment.
- 30 5. An integration framework according to claim 4, further including at least one switch for analyzing at least a part of data received and for routing the data to one of the first segment and the second segment.

6. An integration framework according to either claim 4 or claim 5, wherein the first bus includes a first dedicated segment dedicated exclusively to communicating the operational data; and
- 5 wherein the second bus includes a second dedicated segment dedicated exclusively to communicating the event data.
7. An integration framework according to any one of the preceding claims, wherein
- 10 the first bus is physically separate from the second bus.
8. An integration framework according to claim 7, further including a router for analyzing at least a part of data received and for routing the data to one of the first bus and the second bus.
- 15
9. An integration framework according to either claim 7 or claim 8, wherein the router analyzes at least one header in the data to determine whether to route the data to the first bus or to the second bus.
- 20 10. An integration framework for facilitating communication with a central authority that manages an industry network, the integration framework including:
- a system infrastructure including:
- a plurality of stationary sensors configured to detect at least
- 25 one aspect regarding infrastructure of the industry network and generate infrastructure data indicative of the at least one aspect; and
- an infrastructure analysis module configured to receive the infrastructure data and to receive endpoint data from one or more mobile sensors attached to mobile units of the industry network, the
- 30 infrastructure analysis module configured to generate event data based on the infrastructure and endpoint data;

an operational bus in communication with the plurality of stationary sensors and the infrastructure analysis module, the operational bus configured to receive the infrastructure and endpoint data and to communicate the infrastructure and endpoint data to the central authority, the infrastructure and endpoint data including at least one real-time measurement for at least a part of the industry network; and

an event bus in communication with the infrastructure analysis module, the event bus configured to receive the event data and to communicate the event data to the central authority, the event bus being separate from the operational bus, the event data being distinct from and derived from the real-time measurement and including at least one analytical determination as to operation of the industry network based on the at least one real time measurement; and

a network core configured to receive the event data exclusively through the event bus and to generate a command based on the event data, wherein events within the event data include undesired or abnormal conditions occurring within the industry system;

wherein the command is sent to an endpoint that includes at least one of the mobile sensors that generated the endpoint data, for execution of the command to control a mobile unit of the industry network in response to the undesired or abnormal conditions.

25

11. An integration framework according to claim 10, wherein the network core is further configured to generate the command based on the at least one of the endpoint data and the infrastructure data.

30

12. An integration framework according to either claim 10 or claim 11, wherein an event includes an alert that is also sent to an organization that owns the one or more mobile units.

13. An integration framework according to claim 12, wherein the alert includes a security message indicative of attempted theft or vandalism.
- 5 14. An integration framework according to any one of claims 10 to 13, wherein the stationary sensors include memory and rules stored in the memory, wherein the stationary sensors are configured to execute the rules to determine an event based on the infrastructure and endpoint data.
- 10 15. An integration framework according to claim 14, wherein the event bus is in further communication with the stationary sensors to receive event data from the stationary sensors.
- 15 16. An integration framework according to any one of claims 10 to 15, wherein the industry system is a railway system and the mobile units include individual railcars.
- 20 17. An integration framework according to claim 16, wherein the stationary sensors are further configured to control one or more switches in a section of the railway system.
- 25 18. An integration framework according to either claim 16 or claim 17, wherein the stationary sensors are configured to modify one or more parameters of one or more other sensors in a section of the railway system.
19. An integration framework according to any one of claims 16 to 18, wherein at least one of the stationary or mobile sensors is configured to control one or more railcars traveling within the railway system.
- 30 20. An integration framework according to any one claims 16 to 19, wherein at least some of the stationary sensors include defect detectors that are configured to detect a potential unsafe condition of a railcar.

21. An integration framework according to any one of claims 10 to 20, wherein the industry system is a trucking system and the mobile sensors are attached to cargo containers to track the cargo containers.
- 5 22. An integration framework according to claim 21, wherein the stationary sensors are distributed at customer facilities and are configured to receive endpoint data from the mobile sensors of the cargo containers.
- 10 23. An integration framework according to any one of claims 10 to 22, wherein the generated events include a no-incident event or an incident event.
- 15 24. An integration framework according to any one of claims 10 to 23, wherein one of the mobile sensors includes a master mobile sensor configured to receive endpoint data, including registration, from each other mobile sensor.
- 20 25. An integration framework for facilitating communication with a central authority that manages an industry network, the integration framework including:
- a plurality of sensors and industry network components configured to generate operational data and event data within the industry network;
 - an operational bus in communication with the plurality of sensors
 - 25 and industry network components, the operational bus configured to receive the operational data and to communicate the operational data to the central authority, the operational data including a real-time measurement for at least one sensor or industry network component; and
 - an event bus in communication with the plurality of sensors and
 - 30 industry network components, the event bus configured to receive the event data and to communicate the event data to the central authority, the event bus being separate from the operational bus, the event data being

distinct from and derived from the real-time measurement and including at least one analytical determination based on the at least one real time measurement,

5 wherein the operational data is communicated via the operational bus and not communicated via the event bus, and

wherein the event data is communicated via the event bus and not communicated via the operational bus.

26. An integration framework according to claim 25, further including:
10 a server coupled with the operational bus and the event bus and configured to receive and store the operational data, the server further configured to:

analyze the operational data with respect to at least one rule;
generate at least one event based on the analysis; and
15 send the at least one event to at least one of the sensors and network industry components to trigger self-healing within at least a section of the industry network.

27. An integration framework according to claim 26, wherein the server
20 resides in a central control center, wherein the at least one event triggers a modification to operation of the central control center.

28. An integration framework according to claim 26 or claim 27, wherein the server is further configured to generate a work-order for transmission to
25 the central authority.

29. A computer-implemented method for communicating data to a central authority that manages an industry network, the method including:
communicating, at least partly wirelessly, operational data to the
30 central authority on a first bus, the operational data including at least one real-time measurement of at least a part of the industry network; and

communicating, at least partly wirelessly, event data to the central authority on a second bus, the second bus being separate from the first bus, the event data being distinct from and derived from the real-time measurement and including at least one analytical determination as to
5 operation of the industry network based on the at least one real time measurement,

wherein the operational data is communicated via the first bus and not communicated via the second bus, and

wherein the event data is communicated via the second bus and not
10 communicated via the first bus.

30. A method according to claim 29, wherein:

said first bus is an operational bus;

said second bus is an event bus;

15 said operational data includes at least one real-time measurement received from at least one sensor or industry network component of a plurality of sensors and industry network components; and

said event data is received from at least some of the plurality of sensors and industry network components.

20

31. A method according to claim 30, further including:

analyzing, by a router, at least a part of data received; and


routing, by the router, the data to one of the operational bus and the event bus based on the analysis of the at least a part of the data received.


25


32. A method according to either claim 30 or claim 31, wherein the operational bus further communicates non-operational data, the non-operational data including at least one of performance, asset health, and stress data.

30

33. A method according to any one of claims 30 to 32, further including:
filtering, by the event bus, a plurality of event streams to interpret the
plurality of event streams according to at least one event pattern; and
sending, by the event bus, the interpretation of the plurality of event
5 streams to at least one of the plurality of sensor and industry network
components.
34. A method according to any one of claims 30 to 33, further including:
receiving and storing the operational data by a server coupled with
10 the operational bus and the event bus;
analyzing, by the server, the operational data with respect to at least
one rule;
generating, by the server, at least one event based on the analysis;
and
15 sending, by the server, the at least one event to at least one of the
sensors and industry network components to trigger self-healing within at
least a section of the industry network.
35. A method according to claim 34, wherein the server resides in a
20 central control center of the industry network, further including:
triggering, by the server, a modification to operation of the central
control center based on the at least one event.
36. A method according to either claim 34 or claim 35, further including:
25 generating a work order by the server for transmission to the central
authority.
37. An integration framework for facilitating communication according to
claim 1 or claim 10 or claim 25, or a computer-implemented method
30 according to claim 29, substantially as hereinbefore described with
reference to the accompanying Figures.













FIG. 2










FIG. 4








FIG. 5A












FIG. 7






FIG. 8





* V = VOLTAGE, I = CURRENT, P = REAL POWER, Q = REACTIVE POWER

FIG. 9B




FIG. 10





FIG. 11

**FIG. 12A**




FIG. 12B




FIG. 12C




FIG. 13A




FIG. 13B

FAULT INTELLIGENCE




FIG. 14B





FIG. 14C

**FIG. 15A**




FIG. 15B




FIG. 16





FIG. 17






FIG. 18B




FIG. 18C




FIG. 18D




FIG. 19A

BLUEPRINT PROCESS LIST

BASE	BASELINE AND SYSTEMS EVALUATION
RDAS	REQUIREMENTS DEFINITION AND ANALYTICS SELECTION
ARC	ARCHITECTURE CONFIGURATION
SELECT	SOLUTION ELEMENT EVALUATION AND COMPONENTS TEMPLATE
UPLAN	UPGRADE PLANNING FOR APPLICATIONS AND NETWORKS
RAMP	RISK ASSESSMENT AND MANAGEMENT PLANNING
CHAMP	CHANGE ANALYSIS AND MANAGEMENT PLANNING
SNARC	SENSOR NETWORK ARCHITECTURE CONFIGURATION
SATSECTR	SENSOR ALLOCATION VIA T-SECTION RECURSION




FIG. 19B





FIG. 20






FIG. 21B





FIG. 21C






FIG. 23A




FIG. 23B





FIG. 23C






FIG. 24B




FIG. 24C




FIG. 25




FIG. 26A




FIG. 26B





FIG. 26C






FIG. 27B





FIG. 27C






FIG. 28B




FIG. 28C




FIG. 29




FIG. 30