Talking After Lights Out: An Ad Hoc Network for Electric Grid Recovery

Jan Janak*, Hema Retty†, Dana Chee‡, Artiom Baloian*, Henning Schulzrinne*
*Department of Computer Science, Columbia University, USA
†FAST Labs, BAE Systems, USA
‡Perspecta Labs, USA
Email: janakj@cs.columbia.edu, dchee@perspectalabs.com, hema.retty@baesystems.com,
ab4659@columbia.edu, hgs@cs.columbia.edu

***Abstract*—When the electrical grid in a region suffers a major outage, e.g., after a catastrophic cyber attack, a "black start" may be required, where the grid is slowly restarted, carefully and incrementally adding generating capacity and demand. To ensure safe and effective black start, the grid control center has to be able to communicate with field personnel and with supervisory control and data acquisition (SCADA) systems. Voice and text communication are particularly critical. As part of the Defense Advanced Research Projects Agency (DARPA) Rapid Attack Detection, Isolation, and Characterization Systems (RADICS) program, we designed, tested and evaluated a self-configuring mesh network prototype called the Phoenix Secure Emergency Network (PhoenixSEN). PhoenixSEN provides a secure drop-in replacement for grid's primary communication networks during black start recovery. The network combines existing and new technologies, can work with a variety of link-layer protocols, emphasizes manageability and auto-configuration, and provides services and applications for coordination of people and devices including voice, text, and SCADA communication. We discuss the architecture of PhoenixSEN and evaluate a prototype on realistic grid infrastructure through a series of DARPA-led exercises.**

***Index Terms*—Ad hoc networks, network architecture, network security.**

I. Introduction

Most electric power outages are locally-contained and recovery can rely on the public or utility-owned communications infrastructure to coordinate restoration and energizing parts of the electrical grid. Large-scale electric power outages, a.k.a blackouts, are rare but do happen [1]–[4]. Recovering from a large-scale outage typically follows a special procedure known as black start. The procedure has been designed to allow the restoration of electricity supplies in a timely manner [5].

In the United States (U.S.), the black start procedure is usually managed by regional transmission organizations (RTOs) that coordinate several electric utilities. For example, PJM, a large RTO, designates specific generators and transmission infrastructure operators as critical for black start [6].

A successful black start requires coordination of electricity supply and demand, typically by incrementally adding both generating capacity and load [7]. Such coordination usually takes place either via phone calls to substation personnel, or via real-time control of supervisory control and data acquisition (SCADA) devices. Both cases require network connectivity. Grid operators often rely on internet service providers (ISPs) for network services [8]. If the ISPs are also impaired by the blackout, network connectivity may be difficult to guarantee. If the blackout is caused by a network-based cyber attack, the attacker may also attempt to actively thwart or delay power restoration, making a bad situation worse [9].

The Defense Advanced Research Projects Agency (DARPA) has recognized the danger network-based cyber attacks represent for the U.S. critical electrical grid infrastructure and launched the Rapid Attack Detection, Isolation, and Characterization Systems (RADICS) program [10]. The goal of the program is to create and evaluate a set of tools to aid the power generation and distribution industry in recovering from a hypothetical large-scale blackout triggered by a network-based cyber attack.

We present the design, prototype implementation, and experimental evaluation of the Phoenix Secure Emergency Network (PhoenixSEN), an ad hoc network for electrical grid recovery. PhoenixSEN is a hybrid, isolated, self-forming network designed to enable the coordination of power restoration. It combines existing and new technologies for rapid deployment into non-cooperative environments, can work with a variety of link-layer protocols, and provides services for coordination of people and (SCADA) devices. PhoenixSEN is designed as a drop-in replacement for grid's primary communication networks severely impaired during a large-scale blackout.

We discuss motivation and related work in Section II. Section III presents a model of the U.S. grid network architecture and discusses DARPA RADICS exercises. We describe the design of PhoenixSEN in Section IV and evaluate a prototype implementation in Section V. We conclude in Section VI.

II. Motivation and Related Work

Modern critical industrial control systems (ICSs), including the electrical grid, require communication to function. Such systems are increasingly targets of cyber attacks [9], [11]. The dangers cyber attacks represent for infrastructure controlling physical processes are well documented [12], [13]. Traditional information technology (IT) system protective measures have

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of DARPA or the U.S. government. Distribution statement A. Distribution approved for public release, distribution unlimited. Not export controlled per ES-FL-020821-0013.

been found inadequate for electrical grid infrastructure [14]. Redundancy and reliability of physical communication infrastructure for rapid black start power delivery restoration was found to be extremely important [2].

The difficulty of experimenting on real systems prompted the research community to create a number of testbeds [15]–[18]. Many of the testbeds researched for this work appear to be using co-simulation, modeling the electrical and networking subsystems separately. Most of the testbeds are software-based, some use real hardware in a limited configuration. Networking subsystems beyond the substation SCADA are rarely emulated.

Network architectures for emergency scenarios have been subject of active research for decades. Many promising architectures based on ad hoc, peer-to-peer, mesh, mobile, delay-tolerant, and opportunistic networking have been proposed [19], [20]. Many of the architectures have been designed to be deployed in isolation without having to interface with legacy systems or link technology provided by third-parties. Thus, they usually lack the necessary flexibility or services to support heterogeneous cyber-physical grid systems.

Observability was found to be a key ingredient of rapid black start recovery. Networking infrastructure achieves observability (topology discovery, bandwidth estimation, intrusion detection) through monitoring. Existing network monitoring and measurement tools [21] can be classified into passive, active, and hybrid [22], [23]. We found that none of the freely available solutions provide sufficient flexibility to discover and monitor substation SCADA networks. Our solution Netmon is an integrated hybrid network monitor inspired by Moloch [24], MRTG [25], and OpenNMS [26]. Netmon integrates with OLSR, DHCP, and intrusion detection subsystems to improve SCADA observability.

Existing grid research mostly focused on the security and restoration of physical grid subsystems. Our work complements such efforts by focusing on the networking subsystems, both within a substation and between substations and utilities. PhoenixSEN is designed to re-establish connectivity, and thus observability and coordination, during black start. The system supports traffic isolation and forensic activities to combat ongoing network-based cyber attacks. To the best of our knowledge, our work is the first attempt to create and evaluate rapidly deployable backup network infrastructure for the electrical grid.

III. Background

A. U.S. Electrical Grid Network Architecture

The U.S. electrical grid is a geographically dispersed system that combines physical infrastructure for producing and delivering electric power with computer-based monitoring, management, and control. The grid has seen incremental upgrades and organic growth, resulting in considerable variability across larger geographical areas and between individual U.S. states. Networking infrastructure plays an increasingly important role and is critical for reliable grid operation.

Several different types of networks are involved in managing the flow of electricity. Fig. 1 shows a simplified model of

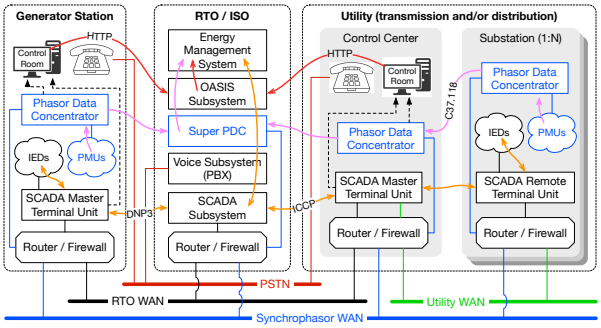

Fig. 1. A simplified model of the U.S. electrical grid networking subsystem. Considerable variety exists across geographic areas and between individual U.S. states (jurisdictions).

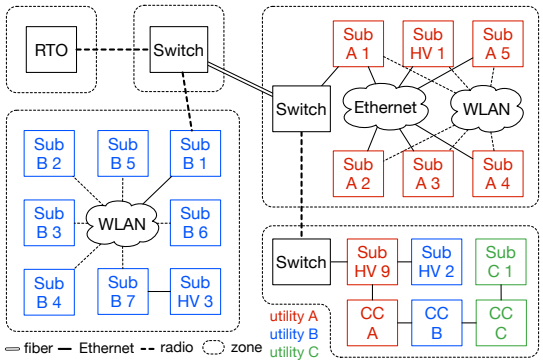

Fig. 2. Physical configuration of the electric grid networking infrastructure during DARPA RADICS exercise 6. Utilities are color-coded.

the grid networking subsystem. The RTO or independent system operator (ISO) operates a redundant wide area network (WAN) to connect to selected utility control centers. The typical RTO/ISO WAN is a redundant Multiprotocol Label Switching (MPLS) network based on external ISPs links. The synchrophasor subsystem will likely use a dedicated WAN with stricter latency and bandwidth guarantees.

Each utility operates a dedicated WAN spanning its (sometimes large) service area that connects the utility's control center with substations. The typical utility WAN is Internet Protocol (IP) based and uses a combination of ISP-owned and utility-owned network infrastructure. A substation with connected devices (e.g., SCADA) usually has a field area network (FAN) that connects substation automation devices as well as any remote devices (metering, data collection) within the substation's service area. Due to the large variety in deployed automation devices, the FAN is perhaps the most heterogeneous network and is typically based on a combination of wired and wireless technologies. SCADA systems use a variety of protocols such as the older Distributed Network Protocol 3 (DNP3), the newer IEC 61850 carried over TCP/IP, and the Inter-Control Center Communications Protocol (ICCN) for data exchange between utilities.

The public internet is commonly used for other communication, e.g., to access the RTO/ISO's Open Access Same-Time Information System (OASIS) portal, or to transfer

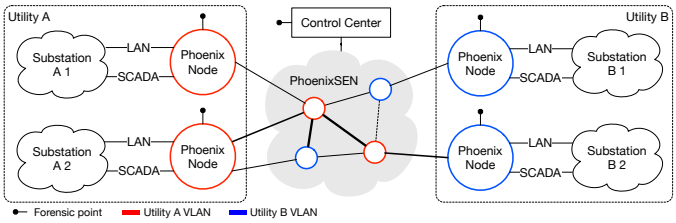

Fig. 3. PhoenixSEN ad hoc network spanning multiple utilities. Each utility is provided with an isolated VLAN spanning its control center and all substations.

metering or billing information between the utility and its customers. Human-to-human voice remains the most important communication modality in emergencies. North American Electric Reliability Corporation (NERC) requires reliable voice communications between the RTO/ISO and utilities. Voice communication is typically provided by the public switched telephone network (PSTN) with cellular or satellite backup.

B. DARPA RADICS Exercises

The DARPA RADICS program organized 7 evaluation exercises from 2016 to 2020 to assess the effectiveness of black start grid restoration technology during a hypothetical large-scale grid cyber attack. The exercises took place in a realistic operational environment with isolated physical grid infrastructure [27]. In this section, we discuss the largest penultimate exercise 6 conducted over 9 days in fall 2019.

Fig. 2 shows the configuration of physical infrastructure during exercise 6. The environment consisted of 21 sites spread across five geographic zones. The simulated RTO was located at Orient Point, NY (Long Island, NY). All remaining sites were located on Plum Island, NY and represented substations and control centers of three independent utilities A, B, and C. The sites were connected with radio, Ethernet, and fiber links. Some links were temporarily provided by the U.S. National Guard, others represented existing utility-owned connections.

Each site or substation was equipped with realistic power grid infrastructure ("substation in a box") that included a remote terminal unit (RTU), real-time automation controller (RTAC), networking and SCADA switches, protective relays, GPS, and various sensors. All infrastructure was designed around commonly deployed U.S. systems to replicate real world conditions. The substations were connected via power lines (not pictured in Fig. 2) to form a multi-utility electrical grid.

The DARPA exercise team repeatedly disrupted the whole system or its parts to simulate blackouts, device malfunctions, or loss of connectivity between sites. Exercise participants were organized into teams: RTO, utility personnel, connectivity, forensics and security. Their goal was to cooperatively restore the system into an operational state as quickly as possible using only the technology and infrastructure available on Plum Island. Communication with the outside world was not permitted.

IV. Phoenix Secure Emergency Network (PhoenixSEN)

PhoenixSEN is a self-configuring ad hoc network designed as a drop-in replacement for the grid's primary networks. The network can be deployed after blackout or under a network-based cyber attack to quickly restore connectivity to control centers (CCs) and substations in a secure manner. PhoenixSEN requires minimal deployment configuration and provides built-in services for voice and SCADA coordination. Uniform hardware and software architecture simplifies deployment to remote substations. The network is compatible with various link technologies including radio, Ethernet, fiber, or powerline and can work with links operated by third-parties.

Fig. 3 shows one possible deployment configuration. A PhoenixSEN node is deployed to each substation, CC, or relay site. The nodes are connected with short-distance and long-distance links, either existing utility-owned links or temporary third-party links provided by, e.g., the National Guard. PhoenixSEN creates a secure isolated virtual network for each utility spanning its CCs and substations. Each PhoenixSEN node provides a separate virtual LAN (VLAN) for voice, SCADA, and forensic traffic to its substation. All essential network services (DHCP, DNS, NTP, VoIP signaling) are provided locally to support communication within the substation even when its PhoenixSEN node is disconnected from the rest of the network.

A dedicated PhoenixSEN CC coordinates the deployment of PhoenixSEN. In the absence of means to communicate with remote substation crew, the CC can use one-way broadcast, e.g., a high-power radio to transmit minimal PhoenixSEN node setup instructions. Once the substation is connected to PhoenixSEN, the CC crew can help configure its infrastructure remotely. PhoenixSEN hardware is designed for crew with technical background, but not necessarily in IT or networking. Detailed instructions are included to allow the substation crew to independently set up the PhoenixSEN node in a minimal operational configuration. Deployed nodes automatically form an ad hoc network based on the Optimized Link State Routing Protocol (OLSR) [28] that eventually connects all utility's sites and infrastructure.

A. PhoenixSEN Node Architecture

The PhoenixSEN node is designed to be deployed from a storage facility to substations by ground transportation or via air lift after blackout. The hardware comes in a weather-resistant enclosure which contains all essential components. Fig. 4 and Fig. 5 show the hardware and software architecture. All nodes have uniform hardware and software configuration to simplify deployment across a large geographic area.

The node is based on the Intel Next Unit of Computing (NUC) computer combined with a VLAN-capable Netgear Ethernet switch. Also included are a pre-configured cordless phone that can call any other PhoenixSEN node, a Wi-Fi access point (AP), an Android smartphone, and a Global Positioning System (GPS) receiver for time synchronization.

The NUC runs Ubuntu Linux 16.04 in a minimal configuration. All custom PhoenixSEN software is pre-installed in the form of Docker containers specifically designed to support software re-building in the field without internet access. This allows fixing critical bugs or vulnerabilities in PhoenixSEN

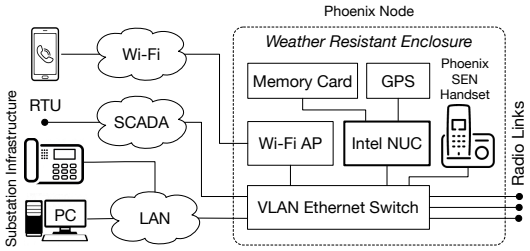

Fig. 4. Hardware architecture of the PhoenixSEN node deployed to utility control centers and substations. The node provides wide area connectivity.

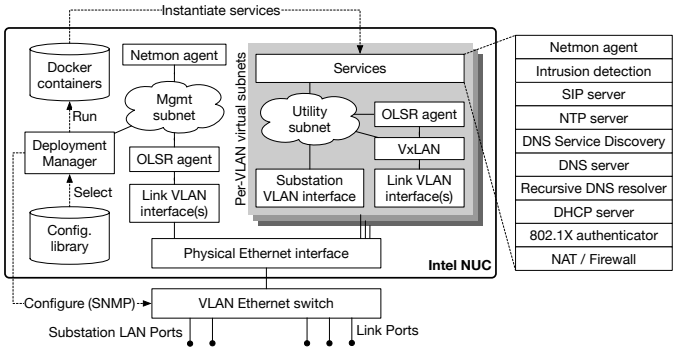

Fig. 5. PhoenixSEN node software architecture. An isolated network environment with all required services is created for each substation LAN.

itself if necessary. The configuration for utility and substation specific services is generated by a configuration synthesis program (not described due to lack of space) for the whole PhoenixSEN prior to deployment. The synthesis uses a model of the networks used by utilities in the deployment area.

Upon deployment, the substation crew provides the node with utility and substation identifiers and the node connects to the corresponding per-utility PhoenixSEN virtual network. The node then creates isolated VLANs for voice, SCADA, and forensic traffic, each implemented with an isolated Linux networking namespace, and exposes the VLANs to the substation. The node also serves as a transparent router for other utilities in the same deployment area.

Each VLAN provides fully isolated network services to the substation including network address translation (NAT), Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), and Network Time Protocol (NTP). Network service virtualization improves robustness and reliability under attack, e.g., if any of the substation SCADA devices have been compromised. Where possible, the services are configured by the synthesis program to match the configuration of the primary ISP networks to minimize the need to reconfigure existing substation infrastructure.

B. Network-Wide Service Discovery

PhoenixSEN is a dynamic ad hoc network without fixed structure. To help applications cope with changing network architecture, PhoenixSEN provides a built-in network-wide service discovery mechanism. Applications based around dynamic

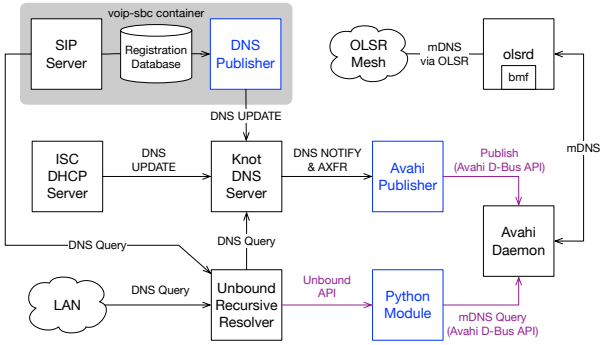

Fig. 6. PhoenixSEN service discovery subsystem architecture. Every node runs the same services. Voice subsystem service discovery is shown in gray.

resource discovery tend to be more robust in case the network is impaired or only partially formed. For maximum interoperability with existing applications, we designed PhoenixSEN service discovery around DNS service discovery (DNS-SD). The primary users are the voice subsystem (Section IV-C) and third-party network and SCADA forensic tools.

At the core of the service discovery subsystem (Fig. 6) is a peer-to-peer DNS service with no single authoritative master server. Instead, every PhoenixSEN node runs a local DNS server (Knot) that stores the DNS service description records published by local services and applications. A multicast DNS (mDNS) [29] publisher (Avahi) disseminates the local records across PhoenixSEN. Clients query the service description records via a recursive DNS resolver (Unbound) also provided by every PhoenixSEN node. The resolver merges records from the local DNS server with the records obtained via mDNS from other PhoenixSEN nodes. This architecture provides an eventual consistency model. Service description conflicts (rare in network like PhoenixSEN) must be resolved by the client.

Efficient network-wide IP multicast is implemented using a Basic Multicast Forwarding (BMF) plugin included with the OLSR agent (olsrd version 1) running on each PhoenixSEN node. The plugin floods IP multicast packets to all nodes in the OLSR network, using OLSR multi-point relays to optimize the flooding. We use the plugin to efficiently disseminate Avahi's mDNS packets across the entire PhoenixSEN.

The network-wide service discovery subsystem provides a simple and intuitive DNS-based interface to applications. To announce service availability, an application on the PhoenixSEN node or in substation LAN can simply publish DNS-SD records via DNS UPDATE to the local DNS server. To discover remote services anywhere in the network, the application can simply query the local resolver for the corresponding DNS-SD records.

To help future practitioners design and build similar systems, we published the network-wide service discovery subsystem as an open source project at https://github.com/janakj/dns2avahi.

C. Voice Communication

PhoenixSEN provides built-in support for Session Initiation Protocol (SIP) [30] based real-time voice and text communication using the following modalities: two-party calls, multi-party

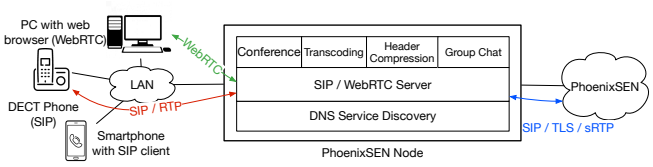

Fig. 7. VoIP subsystem architecture with support for SIP and WebRTC. Network-wide DNS service discovery helps eliminate bottlenecks.

conferencing, and text messaging (Fig. 7). The included Voice over IP (VoIP) clients can be used to place calls and send text messages to any site in any utility across the entire network. We designed the service to require no manual configuration.

Each PhoenixSEN node comes with a VoIP handset and an Android smartphone with a pre-installed SIP client. Both devices are pre-configured for immediate use. The PhoenixSEN node also provides a JavaScript WebRTC application that can turn any device with a compatible web browser into an additional VoIP and chat client. Every PhoenixSEN node runs a full set of VoIP services. The clients register with the VoIP server on the PhoenixSEN node within the same substation. Even if the substation is disconnected from PhoenixSEN, local calls remain possible. The VoIP server enforces transcoding, authentication, and encryption on all calls to PhoenixSEN.

The VoIP subsystem has a peer-to-peer architecture that requires no fixed dialing plan. VoIP clients and servers use network-wide service discovery (Section IV-B) to locate each other. Upon VoIP client registration, the VoIP server publishes a custom DNS service discovery record mapping the client's number to the name of the PhoenixSEN node where it has registered. When a remote VoIP server receives a call for the client, it uses DNS service discovery to map the called number to the PhoenixSEN node where the client is registered and forwards the call there. The mDNS subsystem pro-actively disseminates DNS records across PhoenixSEN. Thus, calls to existing (registered) numbers are resolved from local DNS cache in constant time. Calls too non-existing (unregistered) numbers take a few seconds to fail, until multicast DNS reports that the corresponding record was not found.

D. Network Monitoring

In a geographically dispersed network such as PhoenixSEN, monitoring and situation awareness are important for successful deployment and operation. Netmon is a near real-time network monitoring service designed for PhoenixSEN. Through Netmon's web-based user interface (UI) (Fig. 8), the CC crew can see the formed PhoenixSEN topology and any alerts generated by security events and incidents. Forensic crews can use Netmon to discover the devices at each substation and inspect their state while investigating a potential network-based cyber attack.

Fig. 9 shows the architecture of Netmon. Every PhoenixSEN node runs a Netmon agent process that collects information about the state of the node using OS-level instrumentation and any devices reachable via substation local area networks (LANs) using active network scanning. The collected data includes network interface statistics, the state of the node's

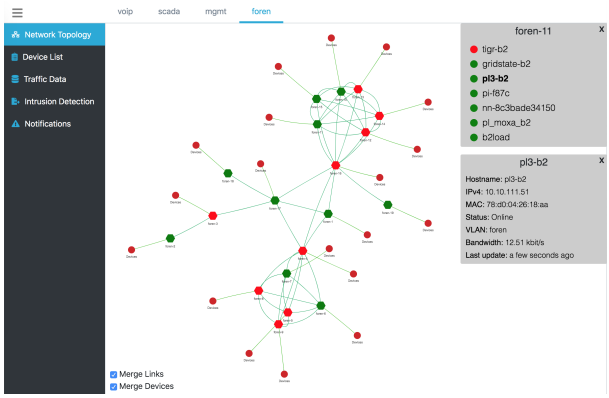

Fig. 8. A PhoenixSEN network topology graph as shown by Netmon during exercise 6. Elements that might require attention are highlighted in red.

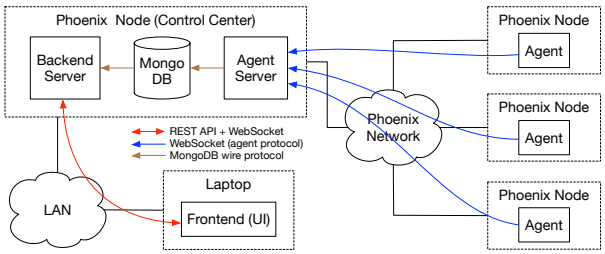

Fig. 9. Architecture of the Netmon monitoring subsystem. PhoenixSEN nodes run agents that stream collected data to a backend server in the control center.

OLSR links, and a list of discovered LAN devices. To discover LAN devices, the agent generates periodic gratuitous ARP requests on all LAN network interfaces. Once a device has been discovered, the agent probes the state of selected User Datagram Protocol (UDP) and Transmission Control Protocol (TCP) ports on the device.

The collected data is streamed in near real-time over a persistent WebSocket [31] connection to a Netmon server in the PhoenixSEN CC. The communication takes place over an isolated management PhoenixSEN virtual network and is thus secure against passive and active attacks. When an agent gets disconnected from the server, it temporarily stores all collected data in a local database. The data will be uploaded to the server later once the agent has reconnected.

The Netmon server correlates and persistently stores the data collected from all agents across PhoenixSEN. The data is indexed to allow time-based addressing and aggregation, e.g., to retrieve the state of the network at a particular time. This feature allows debugging or post-mortem analysis after an exercise when the physical network infrastructure is no longer operational. As new data becomes available, the server pushes updates to any connected UI clients. This allows the UI to show the most recent state of PhoenixSEN.

V. Experimental Evaluation

During DARPA RADICS exercises, PhoenixSEN provided connectivity within utility substations, between utilities, and to the RTO. Fig. 2 shows the physical configuration of the testbed

TABLE I
CONFIGURATION OF DARPA RADICS EXERCISE 6 NETWORK

Grid parameters	Utility A	Utility B	Utility C
Control centers (CC)	1	1	1
Substations	5	7	1
High voltage (HV) substations	2	2	0
RTOs	1 (shared across utilities)		
Maximum RTO-substation distance	5 links		
PhoenixSEN parameters	SCADA	VoIP	Forensic
Connected devices	33	69	265
PhoenixSEN nodes	21 (shared across VLANs)		
Physical links per node	min: 1, max: 3		
OLSR Routes per node	min: 1, average: 5, max: 9		
Inter-utility relay nodes	3 (HV 9, CC A, CC B)		
OLSR network diameter	6 (all VLANs)		

network during exercise 6. We chose the penultimate exercise 6 to evaluate PhoenixSEN because it included the largest number of participants and featured the most sophisticated testbed network configuration.

A PhoenixSEN node was installed at each substation and connected the substation infrastructure to the rest of the environment via external links. PhoenixSEN carried three kinds of traffic in isolated VLANs: VoIP, SCADA, and forensic. Forensic traffic was generated by other exercise participants in order to analyze and mitigate simulated problems in the power grid infrastructure.

To see whether the OLSR overlay correctly spanned all substations, we analyzed the data collected by Netmon at all nodes during the exercise. Fig. 10 shows the OLSR mesh topology as a graph. The graph was identical for all VLANs. The data shows the overlay network correctly spanning all sites. The diameter of the overlay network was 6, which is optimal for this physical configuration. PhoenixSEN correctly handled a variety of links including point-to-multipoint wireless, point-to-point radio and fiber links, Ethernet, and utility fiber connections. As long as all links were operational, PhoenixSEN converged to an optimal topology without network partitioning.

The convergence time in OLSR is a function of the intervals of "hello" and "traffic control" messages periodically transmitted by nodes. In PhoenixSEN we used 2s and 5s intervals, respectively. The exercise network had a network diameter of 6 which gives a restart convergence time of 30s or less when all nodes and links are operational. Thus, for networks of similar size, the recovery time will be dominated by factors other than OLSR convergence time.

As can be seen in Fig. 2, utility B control center and the entire utility C were connected to the rest of the system via utility A control center. This configuration provided an opportunity to test intra-utility traffic relaying where a PhoenixSEN node at one utility relays traffic for other utilities in the area. Based on the Netmon data, this feature worked correctly during the entire exercise. Utility A substations were connected via redundant links (Fig. 2). PhoenixSEN correctly handled that situation as well. When the Ethernet connection was operational, it received all traffic between the substations. When the Ethernet connection got disconnected, OLSR would transparently re-route all traffic to the backup wireless network.

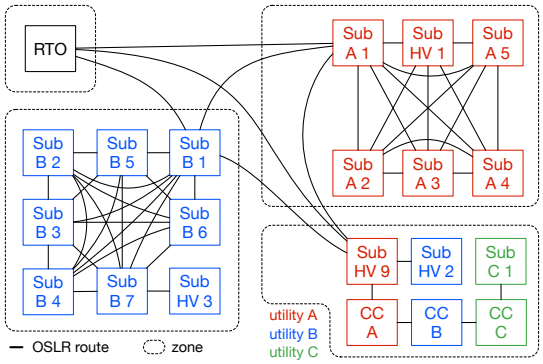

Fig. 10. OLSR overlay network (routes) shown as a graph over physical network configuration. The overlay was identical across all three VLANs.

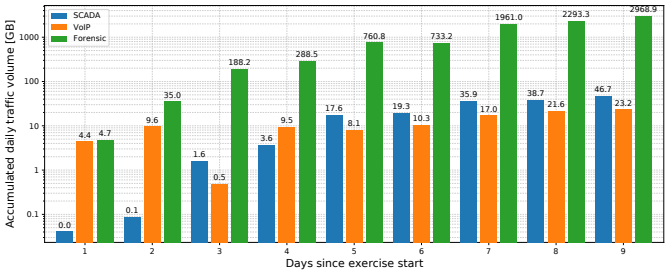

Fig. 11. Daily traffic volumes in the three VLANs supported by PhoenixSEN during DARPA RADICS Exercise 6.

Fig. 11 shows daily traffic volumes for each of the three VLANs, as reported by Netmon. The initial ramp-up period during the first days of the exercise was caused by partially operational testbed. As expected, forensic traffic generated by a large number of exercise participants dominates VoIP and SCADA traffic. Utilities planning backup networking infrastructure may need to take forensic overhead into consideration.

VI. CONCLUSION

We presented the design, implementation, and experimental evaluation of PhoenixSEN, an ad hoc network for real-time coordination of people and SCADA devices. PhoenixSEN has been designed for the needs of the power distribution industry during black start. The network is flexible and can be used as a temporary replacement for third-party ISP networks. PhoenixSEN is designed to speed up electric grid restoration amidst a persistent or ongoing network-based cyber attack.

We tested and evaluated a PhoenixSEN prototype (Fig. 12) in a series of DARPA-led cyber security exercises on isolated realistic electrical grid infrastructure. During the exercise, PhoenixSEN provided connectivity for utility substations, control centers, the RTO, and also supported forensic activities that were part of the recovery process. At the peak of exercise activities, the network consisted of 21 PhoenixSEN nodes and simultaneously connected over 350 devices.

The network performed well and correctly handled links with different kinds of technology including radio point-to-point links, wireless point-to-multipoint connections, and utility fiber.

Fig. 12. Nine PhoenixSEN nodes set up in a lab at the University of Illinois at Urbana-Champaign. The full PhoenixSEN prototype consisted of 21 nodes.

The combined OLSR-VxLAN mesh network correctly handled simulated outages and disconnects and quickly converged to an optimal topology in all cases. Uniform hardware and software architecture proved to be very useful during deployment.

One of our original goals was to design a self-configuring backup network for the electric grid that could be operated by utility personnel. We believe that goal was only partially reached. The current prototype requires deployment planning and one-time configuration. Both steps assume background in computer networking. Simplifying the deployment process further is one of our next research goals.

Acknowledgment

We thank Clifton Lin, Charles Tao, Defu Li, Ranga Reddy, and James Dolan, all members of the BAE Systems team, for fruitful discussions and help in developing and testing the prototype. We are grateful to Frafos GmbH for a free license to use the ABC SBC platform and for help with configuration.

References

[1] The Wall Street Journal, "The Texas grid came close to an even bigger disaster during February freeze," https://www.wsj.com/articles/texas-electrical-grid-bigger-disaster-february-freeze-black-starts-11622124896, 2020, [Online; accessed: May 2020].

[2] G. Andersson, P. Donalek, R. Farmer, N. Hatziargyriou, I. Kamwa, P. Kundur, N. Martins, J. Paserba, P. Pourbeik, J. Sanchez-Gasca *et al.*, "Causes of the 2003 major grid blackouts in North America and Europe, and recommended means to improve system dynamic performance," *IEEE transactions on Power Systems*, vol. 20, no. 4, pp. 1922–1928, 2005.

[3] CNN, "Massive failure leaves Argentina, Paraguay and Uruguay with no power," https://edition.cnn.com/2019/06/16/world/power-outage-argentina-uruguay-paraguay, Aug. 2019, [Online; accessed: Aug. 2019].

[4] BBC, "Major power failure affects homes and transport," https://www.bbc.com/news/uk-49300025, Aug. 2019, [Online; accessed: Aug. 2019].

[5] National Grid ESO, "Black start," https://www.nationalgrideso.com/balancing-services/system-security-services/black-start, [Online; accessed: Sep. 2019.

[6] PJM, "PJM manual 12: Balancing operations," https://www.pjm.com/~/media/documents/manuals/m12.ashx, [Online; accessed: Sep. 2019].

[7] M. Adibi and L. Fink, "Power system restoration planning," *IEEE Transactions on Power Systems*, vol. 9, no. 1, pp. 22–28, 1994.

[8] National Institute of Standards and Technology, "NIST framework and roadmap for smart grid interoperability standards, release 3.0," http://dx.doi.org/10.6028/NIST.SP.1108r3, Sep. 2014, [Online; accessed: Sep 2021].

[9] D. E. Whitehead, K. Owens, D. Gammel, and J. Smith, "Ukraine cyber-induced power outage: Analysis and practical mitigation strategies," in *2017 70th Annual Conference for Protective Relay Engineers (CPRE)*, 2017, pp. 1–8.

[10] Defense Advanced Research Projects Agency (DARPA), "Rapid Attack Detection, Isolation, and Characterization Systems (RADICS) program," https://www.darpa.mil/program/rapid-attack-detection-isolation-and-characterization-systems, [Online; accessed: Aug. 30, 2019].

[11] National Electric Sector Cybersecurity Organization Resource (NESCOR, "Analysis of selected electric sector high risk failure scenarios," https://smartgrid.epri.com/NESCOR.aspx, [Online; accessed Oct. 18, 2019].

[12] S. Karnouskos, "Stuxnet worm impact on industrial cyber-physical system security," in *IECON 2011 - 37th Annual Conference of the IEEE Industrial Electronics Society*, Nov 2011, pp. 4490–4494.

[13] J. Slay and M. Miller, "Lessons learned from the Maroochy water breach," in *International Conference on Critical Infrastructure Protection*. Springer, 2007, pp. 73–82.

[14] Idaho National Laboratory. (2011, Sep.) Vulnerability analysis of energy delivery control systems (INL/EXT-10-18381).

[15] M. McDonald, J. Mulder, B. Richardson, R. Cassidy, A. Chavez, N. Pattengale, G. Pollock, J. Urrea, M. Schwartz, W. Atkins *et al.*, "Modeling and simulation for cyber-physical system security research, development and applications," *Sandia National Laboratories, Tech. Rep. Sandia Report SAND2010-0568*, 2010.

[16] M. M. Roomi, P. P. Biswas, D. Mashima, Y. Fan, and E.-C. Chang, "False data injection cyber range of modernized substation system," in *2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*. IEEE, 2020, pp. 1–7.

[17] A. Ashok, S. Krishnaswamy, and M. Govindarasu, "PowerCyber: A remotely accessible testbed for cyber physical security of the smart grid," in *2016 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*. IEEE, 2016, pp. 1–5.

[18] P. Gunathilaka, D. Mashima, and B. Chen, "Softgrid: A software-based smart grid testbed for evaluating substation cybersecurity solutions," in *Proceedings of the 2nd ACM Workshop on Cyber-Physical Systems Security and Privacy*, 2016, pp. 113–124.

[19] F. Legendre, T. Hossmann, F. Sutton, and B. Plattner, "30 years of wireless ad hoc networking research: what about humanitarian and disaster relief solutions? what are we still missing?" in *Proceedings of the 1st International Conference on Wireless Technologies for Humanitarian Relief (ACWR'11)*, 2011, pp. 217–217.

[20] G. V. Kumar, Y. V. Reddyr, and D. M. Nagendra, "Current research work on routing protocols for MANET: a literature survey," *international Journal on computer Science and Engineering*, vol. 2, no. 03, pp. 706–713, 2010.

[21] L. Cottrell, "Network monitoring tools," https://www.slac.stanford.edu/xorg/nmtf/nmtf-tools.html, 2020, [Online; accessed: May 8, 2020].

[22] B. B. Lowekamp, "Combining active and passive network measurements to build scalable monitoring systems on the grid," *SIGMETRICS Perform. Eval. Rev.*, vol. 30, no. 4, p. 19–26, Mar. 2003. [Online]. Available: https://doi.org/10.1145/773056.773061

[23] W. John, S. Tafvelin, and T. Olovsson, "Passive internet measurement: Overview and guidelines based on experiences," *Computer Communications*, vol. 33, no. 5, pp. 533–550, 2010.

[24] "Moloch full packet capture," https://molo.ch, 2020, [Online; accessed: May 8, 2020].

[25] T. Oetiker, "The Multi Router Traffic Grapher," https://oss.oetiker.ch/mrtg, 2020, [Online; accessed: May 8, 2020].

[26] The OpenNMS Group, "The OpenNMS platform," https://www.opennms.com, 2020, [Online; accessed: May 8, 2020].

[27] DARPA, "Technologies to rapidly restore the electrical grid after cyber attack come online," https://www.darpa.mil/news-events/2021-02-23, 2021, [Online; accessed: Jun 1, 2020].

[28] T. Clausen and P. Jacquet, "Optimized Link State Routing Protocol (OLSR)," RFC 3626, Internet Engineering Task Force, Oct. 2003. [Online]. Available: http://www.ietf.org/rfc/rfc3626.txt

[29] S. Cheshire and M. Krochmal, "Multicast DNS," RFC 6762, Internet Engineering Task Force, Feb. 2013. [Online]. Available: http://www.ietf.org/rfc/rfc6762.txt

[30] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler, "SIP: Session Initiation Protocol," RFC 3261, Internet Engineering Task Force, Jun. 2002. [Online]. Available: http://www.ietf.org/rfc/rfc3261.txt

[31] I. Fette and A. Melnikov, "The WebSocket Protocol," RFC 6455, Internet Engineering Task Force, Dec. 2011. [Online]. Available: http://www.ietf.org/rfc/rfc6455.txt