



**Pacific Northwest**  
NATIONAL LABORATORY

*Proudly Operated by **Battelle** Since 1965*

# Concept of Operations: Essence

**April 2014**

WJ Hutton



U.S. DEPARTMENT OF  
**ENERGY**

Prepared for the U.S. Department of Energy  
under Contract DE-AC05-76RL01830

## DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor Battelle Memorial Institute, nor any of their employees, makes **any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights.** Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or Battelle Memorial Institute. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

PACIFIC NORTHWEST NATIONAL LABORATORY  
*operated by*  
BATTELLE  
*for the*  
UNITED STATES DEPARTMENT OF ENERGY  
*under Contract DE-AC05-76RL01830*

Printed in the United States of America

Available to DOE and DOE contractors from the  
Office of Scientific and Technical Information,  
P.O. Box 62, Oak Ridge, TN 37831-0062;  
ph: (865) 576-8401  
fax: (865) 576-5728  
email: [reports@adonis.osti.gov](mailto:reports@adonis.osti.gov)

Available to the public from the National Technical Information Service  
5301 Shawnee Rd., Alexandria, VA 22312  
ph: (800) 553-NTIS (6847)  
email: [orders@ntis.gov](mailto:orders@ntis.gov) <<http://www.ntis.gov/about/form.aspx>>  
Online ordering: <http://www.ntis.gov>



This document was printed on recycled paper.

(8/2010)

# **Concept of Operations: Essence**

WJ Hutton

April 2014

Prepared for  
the U.S. Department of Energy  
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory  
Richland, Washington 99352



## Summary

This concept of operations is designed to give the reader a brief overview of the National Rural Electric Cooperative Association's Essence project and a description of the Essence device design. The data collected by the device, how the data are used, and how the data are protected are also discussed in this document.



## **Acknowledgments**

The author would like to thank Carnegie Mellon University, Cigital, and the National Rural Electric Cooperative Association for their support and partnership on the Essence project. The author would also like to acknowledge the specific collaboration of Philip A. Craig, Jr., and Craig Miller.



## Acronyms and Abbreviations

CMU	Carnegie Mellon University
CPU	central processing unit
IP	internet protocol
NRECA	National Rural Electric Cooperative Association
PNNL	Pacific Northwest National Laboratory



# Contents

Summary .....	iii
Acknowledgments.....	v
Acronyms and Abbreviations.....	vii
1.0 Introduction.....	1
2.0 Purpose of the Pilot.....	2
3.0 Device Design.....	3
4.0 Device Placement .....	6
5.0 Device Security.....	7
6.0 What Data are Collected.....	8
7.0 How the Data are Used.....	9
8.0 How the Data are Protected .....	10
9.0 Conclusion .....	11

# Figures

1	Stealth 1402 LPC .....	3
2	Stealth 1402 LPC .....	3
3	Bro Log Sample .....	4
4	Network Visualization .....	5
5	Network Tap Diagram .....	6

## 1.0 Introduction

The Cooperative Research Network, the research branch of the National Rural Electric Cooperative Association (NRECA), was awarded a \$4 million project by the U.S. Department of Energy to research next-generation cybersecurity devices for the power grid under the thought leadership of NRECA chief scientist Craig Miller, Ph.D. The project name is Essence, and NRECA partnered with Carnegie Mellon University (CMU) and Pacific Northwest National Laboratory (PNNL). Large corporations such as Honeywell, CISCO, Intel, and IBM are also interested in collaborating with the Essence project team.

Previous approaches to cybersecurity were prescriptive and required knowledge of a specific threat to the network in order to create white lists, black lists, and detection signatures. The goal of the Essence project is to implement a new, non-prescriptive approach to cybersecurity for power system networks. Essence proposes to use inexpensive, commodity devices to monitor power system networks. Machine learning algorithms from CMU will be used to determine “normal” network behavior, and will assist in making decisions and taking actions. This non-prescriptive approach could compliment prescriptive security controls by informing traditional information technology security devices. The Essence project will also include near real-time packet analysis of the machine-to-machine communications that are typical of power system networks. Messages will be analyzed for syntax as well as semantics to inform decisions and ultimately take automated actions, if desired.

## **2.0 Purpose of the Pilot**

The purpose of the initial pilot is two-fold. The main goal is to collect real-world power system network traffic for research and development purposes. MultiSpeak traffic is the primary interest. Distributed network protocol and legacy serial data are the secondary interest. Serial communications may be made routable for capture and analysis using SerialTap technology developed by PNNL. Soliciting feedback on this document and process is also a goal.

### 3.0 Device Design

The device design is predicated on several criteria: multiple network adapters; an affordable, small form factor that can be weatherproofed if necessary; enough central processing unit (CPU) resources to conduct edge analytics; and sufficient internal storage capacity.

First, multiple network adapters are required if a network tap is used to observe full-duplex network traffic. Typically, one adapter monitors inbound traffic and a second adapter monitors outbound traffic. A third network adaptor can be useful for communicating with a management network. Typical management activities include data exfiltration and device management. However, for the sake of simplicity and security, a management network is not currently used.

An affordable device is necessary if the devices are to be deployed en mass. A small form factor is desired to limit the impact to server environments. Lastly, a device that can be weatherized will allow deployment location flexibility. Initially, sufficient CPU capacity is necessary to keep up with aggregated network packet capture in real time. In the future, additional CPU capacity is desired to provide resources to conduct edge analytics and support automated decisions and actions. The 1402 LPC from Stealth (Figures 1 and 2) meets all of these requirements.



**Figure 1.** Stealth 1402 LPC



**Figure 2.** Stealth 1402 LPC

All of the software installed on the device, including the operating system, is open source. The operating system is Ubuntu 12.04 LTS. Tcpcap is used for full or selective packet capture. The Bro intrusion detection system is used to classify and categorize observed network traffic (Figure 3). Afterglow and associated Perl scripts are used to first transform Bro logs into comma-separated values formatted text files and then transform the comma-separated values files into GraphViz description language. Lastly, Neato is used to render a graph of the observed network traffic, including source internet protocol (IP) address (ovals), destination IP addresses (rectangles), and destination port used (circles) (as shown in Figure 4). The color of these objects can be assigned using regular expressions; for example, to denote different subnets, white listed IP addresses, allowed or disallowed port numbers, etc.

```
>== Total === 2014-03-17-12-05-39 - 2014-03-17-12-59-45
- Connections 2.1k - Payload 5.5m -
Ports      | Sources          | Destinations      | Services          | Protocols | States          |
80         | 192.168.0.179#1 | 192.168.10.22#2   | -                | 6          | 50             |
389        | fe80::20a:9dff:fe10:acff#3 | 130.20.248.22#4 | dns              | 17.4%     | 23.3%         |
443        | ::#5            | 192.168.100.43#6 | http             | 1          | SH             |
53         | 0.0.0.0#7       | 192.168.1.164#8  | dhcp             | 1.9%     | 4.5%          |
123        | fe80::3e07:54ff:fe4b:badd#9 | 130.20.67.21#10 |                  | 0.3%     | OTH            |
137        |                 | 192.168.100.33#11 |                  |           |               |
445        |                 | 192.168.2.122#12 |                  |           |               |
829        |                 | 192.168.10.21#13 |                  |           |               |
139        |                 | 192.168.2.255#14 |                  |           |               |
3268       |                 | 192.168.150.124#15 |                  |           |               |


#1=<??> #2=<??> #3=<??>
#4=<??> #5=<??> #6=<??>
#7=<??> #8=<??> #9=<??>
#10=<??> #11=<??> #12=<??>
#13=<??> #14=<??> #15=<??>

>== Top 10 local networks by number of connections
1  0  10.0.0.0/8      Private IP space
2  0  192.168.0.0/16 Private IP space

>== 2069 connections did not have any local address. Here are the first 10:
192.168.0.179 <-> 10.0.1.131
192.168.0.179 <-> 10.0.1.142
192.168.0.179 <-> 10.0.1.192
fe80::20a:9dff:fe10:acff <-> ff02::2
192.168.0.179 <-> 10.0.1.198
192.168.0.179 <-> 10.0.1.217
192.168.0.179 <-> 10.0.1.138
192.168.0.179 <-> 10.0.1.107
192.168.0.179 <-> 10.0.1.194

>== Incoming === N/A - N/A
- Connections 0 - Payload 0 -
Ports      | Sources          | Destinations      | Services          | Protocols | States          |
```

Figure 3. Bro Log Sample




**Figure 4.** Network Visualization

Hourly Bro connection summary logs are aggregated to provide a picture of every source IP, destination IP, and destination port observed. Additional information such as services, protocols, and states is also available in the log but that information is not currently used. Additional Bro logs record information about network communications, domain name service, files, hypertext transfer protocol traffic, scripts, secure socket layer, syslog, and other abnormal traffic. At this time, these Bro logs are not used by the Essence project. Additional information about Bro logs is available at <http://www.bro.org/documentation/index.html>.

## 4.0 Device Placement

Essence device placement depends on the site's network configuration and the network traffic the project is trying to observe. A discussion with each site will determine the best device placement location within the site's network. Currently, the primary goal is to collect MultiSpeak traffic. The availability of existing network taps or span ports simplifies the deployment of the Essence device. It can simply be plugged in. If no copies of the desired network traffic are available via tap or span port, a network outage may be required to install a tap. PNNL has NetOptics network taps available for temporary installation in order to facilitate research data collection. A diagram of how the tap is installed in a typical network is shown in Figure 5.



**Figure 5.** Network Tap Diagram

## 5.0 Device Security

Reasonable steps were taken to harden the Essence device. The operating system is Ubuntu 12.04 LTS with all recent patches applied. Unneeded packages were removed and unnecessary services were disabled. Likewise, unnecessary user accounts were removed, and the remaining accounts have strong passwords (a 20-character mix of upper and lower case letters, numbers, and printable symbols).

Unused network interfaces were disabled and used network interfaces are addressable only by media access control address. The device does not listen for or accept external connections. Device access is only available via physical presence using the console.

## 6.0 What Data are Collected

What data are collected is largely a function of where the device is placed in the site's network. An accounting mechanism will be used to track all copies of the site data.

Initial research requirements require full-packet capture of desired protocols, specifically MultiSpeak traffic. Additional protocols, such as distributed network protocol and serial communications, may be needed also.

The project is specifically interested in industrial control system/supervisory control and data acquisition network traffic. The nature of this traffic is typically machine-to-machine. It is fairly predictable and often does not require confidentiality. This type of traffic should never contain any personally identifiable information. If it is determined that the Essence device may observe personally identifiable information based on where it is deployed within the site's network, an impact assessment should be conducted.

## **7.0 How the Data are Used**

Any network traffic collected by our device will only be shared with Essence project partners (currently Cooperative Research Network, Cigital, CMU, NRECA, and PNNL). This data will be used by CMU to develop machine-learning algorithms to categorize network traffic as normal or abnormal and provide a confidence value of that categorization. Network data may also be played back to testing, validation, and verification of algorithms developed by CMU.

## 8.0 How the Data are Protected

While the Essence device is within the participating site's security accreditation boundary, the site is responsible for the physical protection of the device and their data stored within the device. While the device is within the accreditation boundary of a research partner, it will fall under that site's cybersecurity plan.

Data obfuscation may be possible if desired by the site. Individual IP addresses would be transformed to specified non-routable IP addresses, maintaining the address classes, subnets, and scope. This process has not been implemented before, but researchers expect it to be trivial. Sensitivities of data sharing will need to be determined before this feature can be developed.

However, the Essence device has an encrypted data store created with TrueCrypt. This encrypted data store can be used to protect site data if the device is removed from a security accreditation boundary or while in transit between accreditation boundaries. TrueCrypt uses advanced encryption standard (2001).

If site data is transmitted between research partners, advanced encryption standard will be used to protect the site data in motion. Site data will not be shared without the prior consent of the participating site. This consent can be revoked at any time. If the site chooses to no longer participate in the Essence project, all copies of the data will be destroyed.

## **9.0 Conclusion**

The Essence project requires operational data from real-world power system networks to for research and development purposes, including the creation of machine-learning algorithms to support next generation, non-prescriptive cybersecurity. If your site uses the MultiSpeak protocol, or you would be interested in field-testing one of our Essence devices in the future, we would love to hear from you.

Please contact Maurice Martin, CRN Program Manager at (703) 907-5694 for more information.



## Distribution

**No. of  
Copies**

1 Maurice Martin  
Cooperative Research Network  
4401 Wilson Boulevard  
Arlington, VA 22203

1 Duane Crum  
Benton PUD  
2721 West 10<sup>th</sup> Avenue  
Kennewick, WA 99336

**No. of  
Copies**

1 Marc Seay  
Rappahannock Electric Cooperative  
247 Industrial Court  
Fredericksburg, VA 22408







**Pacific Northwest**  
NATIONAL LABORATORY

*Proudly Operated by **Battelle** Since 1965*



U.S. DEPARTMENT OF  
**ENERGY**

---

902 Battelle Boulevard  
P.O. Box 999  
Richland, WA 99352  
1-888-375-PNNL (7665)  
[www.pnnl.gov](http://www.pnnl.gov)