

Digital Forensics Compute Cluster: A High Speed Distributed Computing Capability for Digital Forensics

Authors : Daniel Gonzales, Zev Winkelman, Trung Tran, Ricardo Sanchez, Dulani Woods, John Hollywood

Abstract : We have developed a distributed computing capability, Digital Forensics Compute Cluster (DFORC2) to speed up the ingestion and processing of digital evidence that is resident on computer hard drives. DFORC2 parallelizes evidence ingestion and file processing steps. It can be run on a standalone computer cluster or in the Amazon Web Services (AWS) cloud. When running in a virtualized computing environment, its cluster resources can be dynamically scaled up or down using Kubernetes. DFORC2 is an open source project that uses Autopsy, Apache Spark and Kafka, and other open source software packages. It extends the proven open source digital forensics capabilities of Autopsy to compute clusters and cloud architectures, so digital forensics tasks can be accomplished efficiently by a scalable array of cluster compute nodes. In this paper, we describe DFORC2 and compare it with a standalone version of Autopsy when both are used to process evidence from hard drives of different sizes.

Keywords : digital forensics, cloud computing, cyber security, spark, Kubernetes, Kafka

Conference Title : ICCTICF 2017 : International Conference on Cyber Threat Intelligence and Cybercrime Forensics

Conference Location : Vancouver, Canada

Conference Dates : August 07-08, 2017