Identity-based Network Security for Commercial Blockchain Services

Casimer DeCusatis, Marcus Zimmermann
Marist College
Poughkeepsie, NY USA
casimer.decusatis@marist.edu

Anthony Sager
BlackRidge Technology
Reno, NV USA
tsager@blackridge.us

***Abstract*—While blockchain services hold great promise to improve many different industries, there are significant cybersecurity concerns which must be addressed. In this paper, we present experimental test bed results for a novel method of user identity management for cloud-based blockchain applications. Using a BlackRidge Technology endpoint on a Windows host, we insert cryptographic identity tokens on the first packet to request a new session. A corresponding gateway appliance in the cloud enforces security policy, blocking unauthorized access at or below the transport layer. Results of penetration testing a sample Hyperledger 1.0 application are discussed. We also demonstrate network segmentation and traffic separation, which allows multiple organizations to share blockchain infrastructure and facilitates compliance auditing.**

Keywords—Blockchain, cybersecurity, identity, authentication, hyperledger, cloud

I. INTRODUCTION

The use of blockchain technologies to provide immutable, distributed transaction ledgers accessible from a large data network has received significant attention recently as a technology with the potential to disrupt and transform virtually every aspect of global business. To cite just a few examples, within the past two years major financial firms including Goldman Sachs, BNY Mellon, UBS, and the New York City Depository Trust and Clearing Corporation have announced plans to move trillions of dollars to blockchain in 2018 [1]. Blockchain applications have also been implemented in the global shipping industry to track import/export requirements [2], for tracking manufacturing raw materials and components in the airline industry [3], for tracking royalties in digital entertainment systems [4], for guaranteeing the safety of food chains from farm to table [2], and for disrupting electric utility companies using solar panels [5]. Virtually every market vertical has begun to explore or implement blockchains as cornerstones of their next generation technology roadmaps. Many of these new services will be hosted in cloud computing environments; the Head of Financial Services Business Development for Amazon Web Services has commented "distributed ledger technology is at the forefront of any discussion related to innovation" [6].

This widespread interest in blockchain distributed ledgers is accompanied by growing concerns regarding cybersecurity. Both the number and severity of cybersecurity incidents have grown dramatically in the past few years [7, 8]. Many of these attacks can be attributed, at least in part, to a lack of rigorous authentication and identity management, protection against insider threats, and lack of rigorous compliance auditing. For example, in May of 2016 an attack on the global financial messaging network known as SWIFT resulted in over $80 million in losses to banks in Bangladesh, the Philippines, Sri Lanka, Vietnam, and elsewhere [8]. A lack of authentication contributes to distributed denial of service attacks (DDoS), such as when the Mirai botnet unleashed a record-setting 1.2 Terabit/second attack against the service provider Dyn, disrupting their DNS service and impacting millions of users on Twitter, Amazon, Spotify, Netflix, Tumblr, and Reddit [7, 8]. When the Mirai source code was released shortly thereafter, variants reportedly disrupted banks and telecom carriers in Russia, Germany, and the United Kingdom. These examples are part of an unfortunate rising trend in cyberattacks, which motivate an urgent need for improved defensive capabilities. A fundamental issue with current blockchain service proposals is a lack of authentication, identity management, and nonrepudiation, along with the associated resistance to DDoS attacks.

In this paper, we present results from our experimental cybersecurity cloud testbed which implements a novel approach to user identity management in blockchain services. We implement identity-based end-to-end security which extends from the blockchain client to the server-side application and fabric. We also demonstrate identity-based network segmentation and traffic separation, which enables multiple users to securely share the same blockchain infrastructure, reduces the risk of DDoS attacks, and enables automated regulatory compliance audits. Our solution is based on a combination of BlackRidge First Packet Authentication™ and BlackRidge Transport Access Control (TAC) technologies, implemented using software endpoints and gateway appliances from BlackRidge Technology. Experimental results will be provided for a sample resource

978-1-5386-4649-6/18/$31.00

trading application using IBM's version of the open source Hyperledger framework. Our approach can easily be generalized to protect many different types of commercial applications.

The rest of this paper is organized as follows. After an introduction to the problem in section I, we describe the Hyper ledger framework and related application issues in section II. In section III we describe the cloud computing network security test bed, including our original software, test results, and analysis. Finally, section IV summarizes our conclusions and potential objectives for future work.

II. Blockchain Applications on Hyperledger

Although it has some mathematical precedents, modern distributed ledger technology underlying blockchains was developed around 2008 [9]. Timestamped transaction records (which can take the form of name/value pairs) are collected to form a data structure called a block. Each block contains a hash of the prior block and a nonce, forming a digital fingerprint linking the blocks to form a block chain. In a public blockchain, the transaction record or ledger is shared among nodes in a distributed peer-to-peer network, and forms an immutable transaction record. Each node performs computations required to add new transactions or blocks, a process called mining or proof of work. Blocks may also contain metadata including timestamps and transaction roots from a Merkel tree [6-60]. Transaction chains may represent self-enforcing or self-executing service agreements, called smart contracts [60]. A consensus process enables all nodes to agree on the ledger content. Being a distributed database, blockchain is subject to Brewer's Theorem (or the CAP theorem), which states that it is impossible to simultaneously achieve consistency, availability, and partition tolerance [57-59]. Blockchain provides eventual consistency, i.e. the transaction ledgers from all participants will be consistent after some interval (usually between milliseconds and seconds).

Hyperledger is an open source collaborative effort, hosted by The Linux Foundation since 2015 [10]. It was created as an umbrella project to advance blockchain technology and open source tools by addressing key features for an ad hoc industry standard distributed ledger. The first release to be widely adopted (Hyperledger version 0.6) featured a shared ledger and digital asset repository, business logic to automate so-called "smart contracts", and limited cryptographic features which made the ledger tamper resistant. More recently, Hyperledger 1.0 has added significant features and functionality enabling enterprise-grade applications. These features include permission control and confidentiality, un-linkable identity privacy for blockchain participants, a modular and easily auditable consensus protocol, and improved scalability. There have been many enhancements provided by industry partners such as Microsoft's Coco platform [11], Cisco's blockchain internet of things protocol initiative [12], and the Enterprise Ethereum Alliance [13], among others. In particular, IBM created an open source beta program [14] called the High Security Business Network (HSBN), in which network resources such as certificate authorities, source code, and peer clients were provided through an IBM Secure Service Container (SSC). We began this research participating in the IBM HSBN beta program, which provided a method to build distributed blockchains using IBM BlueMix applications on the IBM LinuxONE cloud, which supports z Systems mainframe enterprise platforms. This allowed us to immediately write and test chaincode applications (i.e. application rules) without the need to design and configure a private blockchain network. This beta program was subsequently developed into the IBM Blockchain Platform [15].

Our work extends the blockchain platform in several important ways. First, we introduce a new method for identity-based network security, which extends end-to-end from the client to the blockchain fabric. This is realized by authenticating the first packet of a network connection request using cryptographic identity tokens, which are inserted into the packet header by a BlackRidge software endpoint or hardware gateway at the client, and later authenticated by a BlackRidge gateway at the blockchain fabric server. All unauthorized traffic (including port scans) is dropped at the transport level, so the traffic source does not receive any acknowledgment or feedback which might be used for reconnaissance or enumeration as the first step in a cyberattack. The identity used in the authentication policy is based on the Hyperledger identity management system. In this manner, we isolate and protect blockchain services from unauthorized access; this helps prevent cyberattacks, enables cloud-based blockchain services, and forms the basis for a zero trust blockchain network.

Second, we introduce identity-based network segmentation and traffic separation, which reduces the risk of cyberattacks and enables regulatory compliance audits. Using the First Packet Authentication described previously, we separate internal traffic between clients and administrative functions used in blockchain. The BlackRidge authentication appliance (which may be physical or virtual) includes a separate policy engine and up to eight dynamically adjustable trust levels, which can be used to enhance access control for channels within the ledger. Trust levels for different authorized users can be changed based on business requirements or in response to potential insider threats. Audit trails for all authorized and unauthorized connection attempts to the blockchain are maintained and can be easily audited using software to parse the log contents.

III. Experimental Test Bed Results

The experimental test bed used to investigate Hyperledger security is shown in figure 1. Marbles is a sample blockchain application provided by IBM [14]. This application utilizes three distinct isolated environments, namely client side Javascript, server side Javascript running in a Node.js environment (client side Javascript opens a web socket to server side Node.js), and chaincode written in GoLang which runs on peers in the blockchain network. The application allows users to trade and delete "marbles" or JSON objects

This work was supported by the National Science Foundation under CC*DNI Integration (Area 4): Application Aware Software-Defined Networks for Secure Cloud Services (SecureCloud). Award #1541384.

which act as a stand-in for cryptocurrency transactions; each marble has attributes including a unique identity, color, size, and owner. The JSON objects are converted to strings which are stored in the chaincode (i.e. the blockchain ledger) as key/value pairs, where the key is the marble identity and the value is a JSON string containing the other marble attributes. Interaction with the chaincode is facilitated by the gRPC protocol on a network peer, and by the Hyperledger Fabric Client SDK (see figure 1).

A web sockets interface between the service layer and browser provides block updates and user query handling. This implementation uses an HTTP/2 Websocket interface over TCP/IP; this supports low latency, bidirectional, asynchronous connectivity and potentially makes more efficient use of server resources. On the cloud host, we use a secure API between the service layer and blockchain for query handling and service invocation (such as gRPC, which offers a wide range of verbs including bidirectional, asynchronous local function calls). Peers are supported by chain code and authenticated by certificate authorities (CA). Within the blockchain service, there is an ordering service (i.e. the ledger database) and a number of other CAs (one handles login attempts, while the others authenticate transactions).

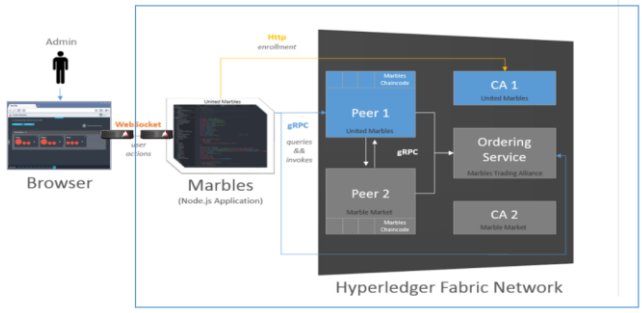

Figure 1 – Hyperledger blockchain authentication test bed. Clients and administrators using a web browser access the cloud-based server hosting Marbles and the Hyperledger Fabric Network.

The recommended architecture implements Node.js on a separate server from the chain code and its related functions. However, we can implement both the Node.js and chain code within a cloud environment, allowing clients to access this system from anywhere in the world using only a web browser. This greatly simplifies the end user implementation, provided that the cloud service provider can be verifiably trusted. The service provider's role as a root of trust is similar to the implicit trust which exists in current financial transaction systems [cite]. As discussed previously, our solution supports automating regulatory compliance auditing.

The Hyperledger network fabric and server-side Node.js server (hosting the Marbles sample application) reside in an enterprise-grade cloud. The blockchain users and administrators access this cloud from a web browser; the BlackRidge endpoint also includes a software implementation on Windows of the identity token insertion services for the authentication system. A BlackRidge authentication gateway (implemented as a hardware appliance) is used at the other end of the connection between the browser and Node.js server, to validate identity tokens. In order for users to access Marbles (or any other secured blockchain application), the first packet of their connection request is authenticated using BlackRidge network security gateways. The BlackRidge software endpoint creates and inserts a time sensitive identity token (valid for four seconds) into the initial TCP/IP connection request. After traversing an untrusted network (such as the Internet) the identity token is authenticated by a BlackRidge gateway (in this case, a physical gateway appliance) installed in the enterprise cloud. Unauthorized packets are discarded, and BlackRidge TAC prevents any further feedback which an attacker might use for reconnaissance. The use of a similar architecture to defend against DDoS attacks by having the authentication gateway blacklist unauthorized traffic has been demonstrated previously [16-18].

To test the authentication solution, we performed a Zenmap port scan of the Node.js and blockchain fabric, with and without identity token enforcement enabled; results are shown in figure 2. Without authentication, our scan identified six open ports which could serve as potential attack vectors for the blockchain service. With authentication enabled, no open ports are identified, and the scanner receives no feedback from the blockchain fabric at or below the transport layer. A record of the scan attempt is logged for future reference. Note that by inhibiting port scanning and blocking unauthorized access attempts before they reach the Node.js server, denial of service attacks can be mitigated. We have repeated this test multiple times with different clients, to confirm that there is no possibility of unauthorized access because of a compromised authentication key; only users with the appropriate key on a specific computer can access the blockchain service.

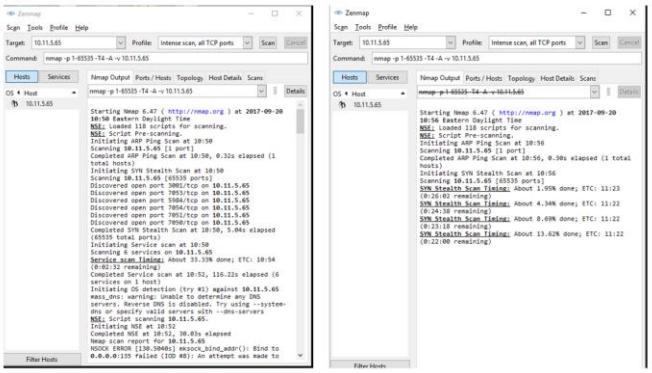

Figure 2 – Zenmap port scan of Marbles and Hyperledger Fabric Network; (left) without authentication and (right) with authentication.

In our next test, we attempted to log into the blockchain fabric with and without identity token enforcement enabled; results are shown in figure 3. Without authentication, we are able to reach the Marbles login screen as the first step in launching an attack on the service, such as attempting to brute

force the password. With authentication enabled, an authorized user will still be presented with the login screen for Marbles, while an unauthorized user receives the message that this site cannot be reached, and no further information is available.

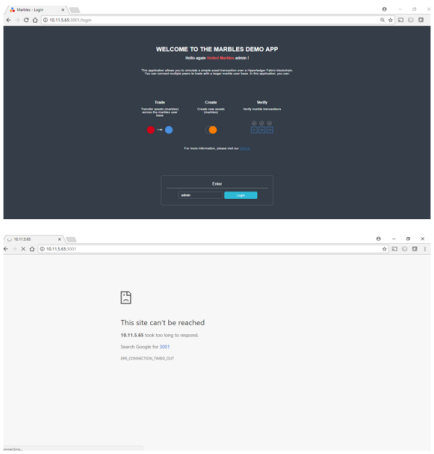

Figure 3 – Results of access attempts to the Marbles login screen: (top) authorized user access and (bottom) unauthorized user access.

As before all authorized and unauthorized access attempts are logged for future reference and auditing purposes. We have also demonstrated the ability to segment the blockchain network using authentication groups tied to the Hyperledger policy engine; this means that multiple companies could share the blockchain infrastructure, while keeping their respective traffic separated.

ACKNOWLEDGMENT

We gratefully acknowledge the support of Marist College and the New York State Cloud Computing and Analytic Center (CCAC), as well as support from the National Science Foundation under CC*DNI Integration (Area 4): Application Aware Software-Defined Networks for Secure Cloud Services (SecureCloud) Award #1541384.

REFERENCES

[1] A. Nordrum, "Wall Street firms to move trillions to blockchain in 2018", IEEE Spectrum, October 2017, https://spectrum.ieee.org/telecom/internet/wall-street-firms-to-move-trillions-to-blockchains-in-2018 (last accessed November 20, 2017)

[2] K. Lewis, "Blockchain: four use cases transforming business", IBM Internet of Things blog, May 2017 https://www.ibm.com/blogs/internet-of-things/iot-blockchain-use-cases/ (last accessed November 20, 2017)

[3] K. Lotay and C. DeCusatis, "Using blockchain technology to digitize supply chain systems", Proc. National Conference on Undergraduate Research, Atlanta, GA, Nov. 3-5, 2017

[4] M. Peck, "Blockchains: how they work", IEEE Spectrum, October 2017, https://spectrum.ieee.org/computing/networks/blockchains-how-they-work-and-why-theyll-change-the-world (last accessed November 20, 2017)

[5] M. Peck and D. Wagman, "Blockchains allow rooftop solar energy trading", IEEE Spectrum, October 2017 https://spectrum.ieee.org/computing/networks/blockchains-will-allow-rooftop-solar-energy-trading-for-fun-and-profit (last accessed November 20, 2017)

[6] A.Flores & K. Gannon, "BlockChain on AWS: Disrupting the Norm", paper GPSD301, AWS Re:Invent 2016 (November 29, 2016) https://www.slideshare.net/AmazonWebServices/aws-reinvent-2016-blockchain-on-aws-disrupting-the-norm-gpst301 (last accessed September 20, 2017)

[7] Cisco Institution, "Cisco 2017 annual cybersecurity report," Cisco, Tech. Rep., 2017.

[8] Mikko Hypponen. and Tomi Tuominen., "F-Secure 2017 State of Cybersecurity report," F-Secure, Tech. Rep., 2017.

[9] S. Nakamoto, "Bitcoin: a peer to peer electronic cash system", Oct. 31, 2008 http://nakamotoinstitute.org/bitcoin/. http://bitcoin.org/bitcoin.pdf. https://github.com/saivann/bitcoinwhitepaper. (last accessed Sept. 20, 2017)

[10] R. Miller, "IBM unveils HyperLedger project", March 19, 2017, https://techcrunch.com/2017/03/19/ibm-unveils-blockchain-as-a-service-based-on-open-source-hyperledger-fabric-technology/ (last accessed September 20, 2017)

[11] M. Russinovich, "The Coco framework for enterprise blockchain networks", August 10, 2017 https://azure.microsoft.com/en-us/blog/announcing-microsoft-s-coco-framework-for-enterprise-blockchain-networks/ (last accessed November 20, 2017)

[12] T. Khurana, "Cisco Internet of Things blockchain protocol initiative", August 1, 2017, https://blogs.cisco.com/tag/blockchain-iot-protocol-initiative (last accessed November 20, 2017)

[13] Enterprise Etherium Alliance press release, July 18, 2017 https://entethalliance.org/enterprise-ethereum-alliance-becomes-worlds-largest-open-source-blockchain-initiative/ (last accessed November 20, 2017)

[14] J. Lang, "Three uses for blockchain in banking", October 23, 2017, https://www.ibm.com/blogs/blockchain/2017/10/three-uses-for-blockchain-in-banking/ (last accessed November 20, 2017)

[15] IBM Press Release, BlockChain at Interconnect 2017 conference, March 21, 2017 https://www.ibm.com/blogs/blockchain/2017/03/guide-everything-blockchain-ibm-interconnect-2017/ (last accessed September 26, 2017)

[16] C. DeCusatis, P. Liengtiraphan, A. Sager, "Advanced Intrusion Prevention for Geographically Dispersed Higher Education Cloud Networks", Proc. IEEE/ACM International Conference on Remote Engineering & Virtual Instrumentation (REV 2017), Columbia University, New York, NY (March 15-17, 2017)

[17] D. Eidle, S. Ni, C. DeCusatis, and T. Sager, "Autonomic security for zero trust networks", Proc. IEEE 8th annual ubiquitous computing, electronics, and mobile communications conference, Columbia University, October 19-21, 2017

[18] G. Leaden, M. Zimmermann, C. DeCusatis, and A. Labouseur, "An API honeypot for DDoS and XSS analysis", Proc. IEEE MIT Undergraduate Research Conference, Cambridge, MA, Nov. 3-5, 2018