WHITE PAPER – JUNE 2017

REDHAT OPENSHIFT CONTAINER PLATFORM PRODUCT APPLICABILITY GUIDE FOR PCI DSS 3.2

APPLICABILITY TO ASSIST CUSTOMERS IN PCI DSS 3.2 DEPLOYMENTS

JASON MACALLISTER
CHRIS KRUEGER | CISSP, QSA

COALFIRE OPINION SERIES
RELEASE VERSION 1.0

North America | Europe
877.224.8077 | info@coalfire.com | coalfire.com

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Red Hat, Inc. (Red Hat) delivers a comprehensive portfolio of products and services built from open source software components using an affordable, predictable subscription model. Red Hat engaged Coalfire Systems, Inc. (Coalfire), a respected Payment Card Industry Qualified Security Assessor (QSA) company, to conduct an independent technical assessment of Red Hat® OpenShift Container Platform (OpenShift) on Red Hat Enterprise Linux (RHEL) and/or Red Hat Atomic Host (Atomic Host). The purpose of this product applicability guide is to identify alignment of OpenShift and the underlying supported operating systems (OS) to the Payment Card Industry Data Security Standard (PCI DSS) v3.2 technical requirements in order to assist payment entities wishing to use the solution in a manner that supports compliance with PCI DSS v3.2.

OpenShift is a container platform that natively integrates open source Linux container technologies and Kubernetes, combining them in an enterprise solution running on RHEL or Atomic Host. OpenShift provides self-service build and deployment automation for containers in addition to operational container features including scaling, monitoring, and management capabilities. Atomic Host is a scaled-down “lightweight” version of the Red Hat Enterprise Linux OS. It has been optimized to run Linux containers and only contains the tools and services necessary to support containerized workloads.

This product applicability guide may be useful to any payment entity that is considering using OpenShift on either RHEL or Atomic Host as part of their cardholder data environment (CDE). The opinion paper discusses the PCI DSS v3.2 requirements that are applicable to OpenShift, RHEL, and Atomic Host. It further identifies and addresses PCI DSS v3.2 requirements that may be addressable by the built-in capabilities of OpenShift, RHEL, and Atomic Host. It is understood that the payment entity must address every PCI DSS v3.2 requirement; however, for this opinion paper, Coalfire has primarily considered technical controls pertinent to and in alignment with OpenShift, RHEL and Atomic Host. Coalfire identified and aligned PCI DSS v3.2 technical requirements from requirement 1, requirement 2, requirement 6, requirement 7, requirement 8, and requirement 10 for applicability and/or supportability by the reviewed products.

COALFIRE OPINION

Security controls, features, and functionality that are built into OpenShift, RHEL, and Atomic Host are capable of supporting and/or addressing relevant technical PCI DSS v3.2 requirements as outlined in the compliance applicability detail section of this paper. It is also Coalfire's opinion that there are sufficient additional Red Hat and third-party solutions available in the market today that satisfy PCI DSS technical requirements where OpenShift and/or RHEL or Atomic Host do not alone support or meet the PCI DSS v3.2 technical requirements.

In general, the use of software-defined networking through OpenShift’s implementation of Open vSwitch along with the OpenShift routing tier are best aligned with requirement 1. The use of OpenShift’s software-defined networking capabilities may be combined in coordination with the payment entity’s overall network security architecture to sufficiently address control requirements. Additionally, placement of OpenShift nodes on the entity’s enterprise network should take into consideration the role that the containers hosted on them may play in the CDE. Strategic network placement and use may better facilitate security, control, and segmentation objectives for the payment entity.

OpenShift, RHEL, and Atomic Host are capable of supporting configuration parameters respectful to requirements specific to the use of vendor default passwords and other default security parameters aligned with PCI DSS requirement 2. The scaled-down Atomic Host helps payment entities align with least function controls, also in support of requirement 2. OpenShift provides mechanisms for application lifecycle management as well as the isolation of test and development environments in support of requirement 6.

Access to the OpenShift platform is limited through the API and/or web console of the OpenShift Master, and role-based access controls provide support for requirement 7. Likewise, RHEL and Atomic Host are capable of supporting role-based access controls as defined by the payment entity for the administration and maintenance of the underlying OS. Through integration with third-party identity and access control sources, OpenShift provides partial support for requirement 8. Finally, because logs are generated by OpenShift, RHEL, Atomic Host, and their individual components, these can provide partial support for requirement 10.

In the compliance applicability detail section of this document, Coalfire suggests consideration for additional non-technical requirements that may be relevant to the use of OpenShift, RHEL, and Atomic Host for PCI DSS v3.2 regulated workloads. Also included are considerations for organizations using segmentation to minimize the scope of PCI DSS assessments and to reduce risk to cardholder data (CHD) in the environment. In the absence of current guidelines for containerization, this opinion is formed based on existing guidance from the Payment Card Industry Security Standards Council (PCI SSC) Special Interest Group (SIG)s for virtualization and cloud computing. The opinion attempts to align with expectations that an assessor may have for addressing a payment entity's infrastructure. The opinion may be useful to an assessor desiring to better understand the use of OpenShift, RHEL, and Atomic Host in PCI DSS scope. Without specific guidance for the use of containerization technologies in infrastructures supporting a CDE, the payment entity risks that an assessor may find the segmentation controls and PCI DSS requirement controls insufficient.

Finally, no one product, technology, or solution is capable of fully addressing security and compliance requirements. Security is a design principle that must be addressed through carefully planned and implemented strategies. Entities seeking compliance are best able to obtain it through a governance, risk, and compliance (GRC) program. For this reason, the introduction of new technologies such as OpenShift and Atomic Host should include payment entity-defined security design principles to reduce risk and maintain or improve security. While Coalfire disclaims the generic suitability of any product for regulatory compliance, Coalfire can confirm that with careful planning and implementation, OpenShift on RHEL or Atomic Host could be included as part of a payment entity's infrastructure in support of a CDE.

INTRODUCING PCI DSS V3.2

The Payment Card Industry Data Security Standard version 3.2 (PCI DSS v3.2) is a proprietary information security standard that was created to reduce credit card fraud by stipulating a series of requirements regulating the use of information systems that handle CHD and sensitive authentication data (SAD). PCI DSS is not an optional standard. As stated, all entities who process, store, or transmit CHD and/or SAD, regardless of geographic location, must comply with the standard or they can be fined and refused access to the card brand's payment system. In addition to the PCI DSS v3.2, the PCI SSC provides additional guidance relevant to understanding and implementing the standard.

UNDERSTANDING PCI DSS SCOPE

A reliable approach for determining where PCI DSS is required to be applied begins with identification and definition of scope for review. Per the PCI DSS Requirements and Security Assessment Procedures document, "PCI DSS security requirements apply to all system components included in or connected to the CDE. The CDE is comprised of people, processes, and technologies that store, process, or transmit cardholder data and sensitive authentication data." (PCI Security Standards Council, 2016). PCI DSS recommends that an assessed entity confirm the accuracy of their PCI DSS scope at least annually or prior to the annual assessment. To help identify scope, the payment entity should evaluate their systems to identify all locations and flows of CHD and identify all systems that are connected to or, if compromised, could impact the CDE. These systems should be included in scope for review.

For this assessment, the PCI DSS v3.2 requirements were limited primarily to technical requirements pertaining to the Red Hat OpenShift Container Platform (OpenShift) on Red Hat Enterprise Linux (RHEL) and/or Red Hat Enterprise Linux Atomic Host (Atomic Host), including "network devices, servers, computing devices, and applications." (PCI Security Standards Council, 2016) Some additional consideration was made for PCI DSS v3.2 operational requirements, inasmuch as these requirements may need to address nuances of the technical platform. The technical evaluation assumed the use of the OpenShift either on RHEL or Atomic Host for developing, testing, building, managing, monitoring, orchestrating, deploying, and hosting a payment entity's CDE applications. Categorizing the infrastructure and application components helps to identify the role that systems play within the CDE with respect to the CHD. When a system is directly engaged with storing, processing, or transmitting CHD or SAD, it is considered a Tier 1 system. Tier 1 systems may include payment applications, host systems that host payment applications, networks that provide transport for CHD or SAD, and so forth. Systems that are adjacent to Tier 1 systems that do not have any direct interactions with CHD or SAD but may provide an avenue to gain access to Tier 1 systems are considered Tier 2 systems. Tier 2 systems, as an example, may include infrastructure monitoring, management, logging, identity, and authentication systems. At a minimum, the OpenShift control plane, when used in this manner, may be considered a Tier 2 system, while OpenShift cluster application nodes may be considered a Tier 1 system when used to host containers that handle CDE. Also included as a Tier 1 system would be any system that provides software-defined networking (SDN) capabilities for the transmission of CHD or SAD where this data is transmitted through the device providing SDN services.

Network Segmentation

Network segmentation, per PCI DSS v3.2, states that isolating (segmenting) the CDE from the remainder of an entity's network is not a PCI DSS requirement; however, it is strongly recommended as a method that may reduce:

- The scope of the PCI DSS assessment
- The cost of the PCI DSS assessment
- The cost and difficulty of implementing and maintaining PCI DSS controls
- The risk to an organization (reduced by consolidating CHD into fewer, more controlled locations) (PCI SSC, 2016)

To evaluate the possibility for the use of segmentation to reduce scope, the payment entity must have a clear understanding of business needs and processes related to the storage, processing, or transmission of CHD. There is an assumption that risk to an organization is reduced through reduction and consolidation of CHD into fewer and more controlled locations as it pertains to storage, processing, and transmission. This is a fair assumption given that the reduction in scope decreases the surface area for attack and minimizes the number of entry points needed to be controlled.

The payment entity will need to consider the presence and placement of OpenShift internal components, nodes, application pods, and containers within their infrastructure when defining segmentation, establishing DMZs, and isolating secure internal zone assets. Deployment of CDE and non-CDE containers on a single cluster node (host) may impact the payment entity's designed segmentation by making all hosted and connected workloads in scope for assessment, regardless of network segmentation. Each payment entity can benefit from understanding controls that are present within OpenShift on RHEL or Atomic Host for isolation of workloads to determine the risk associated with commingling CDE and non-CDE. Current guidance from PCI SSC Special Interest Group (SIG) on cloud and virtualization strongly recommends separation of workloads representing differing zones of trust onto separate hosts.

Wireless Technology

Where used to store, process, or transmit cardholder data, the wireless network and everything attached to it would be considered in scope for application of PCI DSS v3.2 requirements and testing procedures. OpenShift is designed to be a part of the enterprise data center infrastructure and would typically be deployed on a wired network.

Use of Third-Party Service Providers/Outsourcing

PCI DSS describes where scope may be impacted using third-party service providers or payment entity outsourcing of services. Specifically, a payment entity that uses a third-party service provider that is used to store, process, or transmit cardholder data on the payment entity's behalf or to manage components that are part of the payment entity's infrastructure, including routers, firewalls, databases, physical security, and/or servers must consider the impact that the use of the service provider may have on the security of the CDE. The expectation is that there will be clear understanding of responsibilities from each party for the protection of CHD. While OpenShift supports multi-tenancy and may be used by service providers for delivery of services to payment entities, the applicability of PCI DSS v3.2 to OpenShift pertaining to delivery by third-party service or hosting providers was not thoroughly evaluated for this document. However, some of the findings in this document may be useful to third-party service or hosting providers for architecture and design of implementation of OpenShift on either RHEL or Atomic Host implementations in support of PCI DSS regulated workloads.

PCI DSS REQUIREMENTS

The PCI DSS standard is comprised of six "control objectives" with twelve "requirements". The following is a listing of control objectives and their associated requirements. These technical and operational requirements were evaluated for applicability with the use of OpenShift on either RHEL or Atomic Host by a payment entity in a PCI DSS regulated environment. Ultimately, the burden for implementing PCI DSS requirements is on the payment entity, regardless of present capabilities of the technology to provide control.

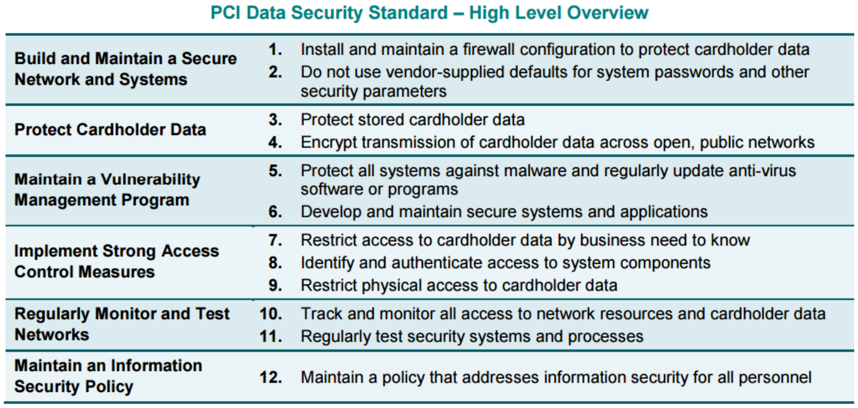

PCI Data Security Standard – High Level Overview

Build and Maintain a Secure Network and Systems	**1.** Install and maintain a firewall configuration to protect cardholder data **2.** Do not use vendor-supplied defaults for system passwords and other security parameters
Protect Cardholder Data	**3.** Protect stored cardholder data **4.** Encrypt transmission of cardholder data across open, public networks
Maintain a Vulnerability Management Program	**5.** Protect all systems against malware and regularly update anti-virus software or programs **6.** Develop and maintain secure systems and applications
Implement Strong Access Control Measures	**7.** Restrict access to cardholder data by business need to know **8.** Identify and authenticate access to system components **9.** Restrict physical access to cardholder data
Regularly Monitor and Test Networks	**10.** Track and monitor all access to network resources and cardholder data **11.** Regularly test security systems and processes
Maintain an Information Security Policy	**12.** Maintain a policy that addresses information security for all personnel

Figure 1: PCI DSS Security Standard High Level Overview (PCI SSC, 2016)

NEW AND EMERGING TECHNOLOGY CHALLENGES TO COMPLIANCE

Many of the new and emerging technologies, such as containerization, software defined data center, software defined networking, software-defined storage, and others, provide more efficient, scalable, extensible, and expedient means to support and deliver services for business and their customers. However, often the challenge with early adoption of new technologies is to understand and properly address the impact that a new technology has on security and compliance. PCI SSC SIGs are sometimes formed to provide guidance on the use of a new technology; however, new technologies often emerge and evolve faster than regulating bodies are capable to address them. For organizations wishing to gain early benefits of a new technology, the level of risk must be evaluated and mitigated, especially as it may impact compliance requirements and the security of protected data.

INTRODUCING OPENSHIFT CONTAINER PLATFORM

OpenShift is an open source container application platform by Red Hat based on open source container technologies and the Kubernetes container cluster manager. It is an integrated platform to run, orchestrate, monitor, and scale containers. It provides build automation, deployment automation, support for a multi-tenant subscriber experience, and a graphical user experience. Much of the information found in this section can be found on Red Hat's website with the product documentation for OpenShift.

OPENSHIFT ARCHITECTURE

To better understand OpenShift, it is useful to explain the components that make up the container platform. The following is a listing of components and the roles that the components play starting with underlying components that support OpenShift.

Operating System - OpenShift can be deployed on either RHEL or Atomic Host. Atomic Host is a container-optimized, minimal footprint OS powered by Red Hat Enterprise Linux. This pre-hardened OS assists organizations with meeting requirements for least functionality due to its lightweight, purpose-built nature as it only includes necessary features, functions, and services to host containers.

Operating Environment - OpenShift can be deployed on bare-metal physical hardware or on virtual infrastructure. It can be deployed on private or certified public cloud environments depending on an organization's specific use cases.

OCI runtime - OpenShift uses an Open Container Initiative (OCI)-compatible runtime for the execution of Linux containers.

Kubernetes – Kubernetes provides orchestration for complex multi-container services. Kubernetes also provides scheduling for services across a container host cluster. OpenShift adds, to Kubernetes, developer and operations-centric tools that enable rapid application development, easy deployment and scaling, and long term life-cycle maintenance for teams and applications. OpenShift also leverages integrated components from Kubernetes to automate application builds, deployments, scaling, health management, and more. Included in the automation capabilities of OpenShift is the ability to configure and deploy Kubernetes container host clusters.

OpenShift Components

OpenShift Nodes – Nodes are instances of RHEL or Atomic Host with the OpenShift software installed. Nodes are where end-user applications are run in containers. Nodes will contain the necessary OpenShift node daemon, the container runtime, and other necessary services to support the hosting of containers.

Containers – Containers are end-user application instances and/or application components run in Linux containers. This Open Container Initiative (OCI)-compatible container technology provides an open source software development and delivery platform that allows applications to be packaged for portability. The container only includes the necessary libraries, functions, elements, and code required to run the application.

Pod –While application components run in containers, OpenShift orchestrates and manages pods. A pod is an orchestrated unit in OpenShift made up of one or more containers. OpenShift will schedule and run all containers in a pod together. Generally, a pod should only contain a single function such as app server or web server and should not include multiple functions such as database and app server.

OpenShift Master – The Master is the control plane for OpenShift. The Master maintains and understands the state of the environment and orchestrates all activity that occurs on the nodes. Just like the nodes, the OpenShift Master is run on either RHEL or Atomic Host. While the Master is technically also a node and can participate in the software-defined network, for separation of function, the OpenShift master should NOT be scheduled to run application instances (pods). The following are the four functions of the OpenShift Master:

> **API and Authentication** – The Master provides the single API that all tooling and systems interact with. Everything that interacts with the OpenShift environment must go through this API. All API requests are SSL-encrypted and must be authenticated. Authorizations are handled by fine-grained role-based access control (RBAC). It is recommended to tie the Master to an external identity and access management system using LDAP or OAuth providers. The Master evaluates requests for both authentication (AuthN) and authorization (AuthZ). Users of the OpenShift environment who have been granted access can be authorized to work with specific projects.
>
> **Desired and Current State** – The state of the OpenShift environment is held in the OpenShift data store. The data store uses etcd, a distributed key-value store. Other information about the OpenShift environment and pertaining to the Master are held in the data store including user account information and the RBAC rules. The OpenShift environment state, including application environment information and non-application user data, is kept in the data store. Important environment variables, secrets data, and other information are also held in the data store.
>
> **Scheduler** – The scheduler determines pod placement within the OpenShift environment. It uses a combination of configuration and environment state (CPU, memory, and other environmental factors) to determine the best fit for running pods across the nodes in the environment. The scheduler is configured with a simple JSON file in combination with node labels to carve up the OpenShift environment. This allows placement of pods within the OpenShift environment to be based on the real-world topology, making use of concepts such as regions, zones, or other constructs relevant to the enterprise. These factors also can contribute to the scheduled placement of pods in the environment and can ensure that pods run on appropriate nodes associated with their function.
>
> **Health and Scaling** – The Master is also responsible for monitoring the health of pods and scaling the pods as desired to handle additional load. The Master executes liveness and readiness tests using probes that are defined by users. The Master can detect failed pods and remediate failures as they occur.

Service Layer – The OpenShift Service Layer allows for application components to easily communicate with one another. For instance, a front-end web service containing multiple web servers would connect to database instances by communication with the database service. OpenShift automatically handles load balancing across the services' instances. In conjunction with probes, OpenShift's service layer ensures that traffic is only directed toward healthy pods, which helps to maintain component availability.

Persistent Storage – Linux containers are natively ephemeral and only maintain data for as long as they are running. Applications and/or application components may require access to a long-term persistent storage repository, such as may be required for a database engine. OpenShift provides the means to connect pods to external real-world storage, which allows for stateful applications to be used on the platform. Persistent storage types that are usable include iSCSI, Fiber Channel, and NFS, as well as cloud-type storage options.

OpenShift Router – The routing layer provides external access to applications hosted in the OpenShift environment. The routing layer operates in partnership with the service layer and provides automated load balancing to pods for external clients. The OpenShift router runs in pods on the platform, but receives traffic from the outside world and proxies the traffic to the appropriate pods. The router uses the service endpoint information to determine where to route and load balance traffic; however, it does not route traffic through the service layer.

OpenShift SDN – The OpenShift software-defined network is a unified cluster network that enables communication between pods across the OpenShift Enterprise cluster. The OpenShift SDN configures an overlay network that uses Open vSwitch (OVS). Red Hat currently provides two SDN plug-ins for use with OpenShift. The ovs-subnet plug-in provides a "flat" pod network where every pod can communicate with every other pod and service. The ovs-multitenant plug-in provides project-level isolation for pods and services. Each project receives a unique Virtual Network ID (VNID) that identifies traffic from pods assigned to the project. Pods from different projects cannot send packets to or receive packets from pods and services of a different project. Administrators of OpenShift can join or isolate projects as required.

OpenShift Registry – The OpenShift Registry provides integrated storage and management for sharing container images.

Figure 2 is a high-level illustration of the OpenShift components.

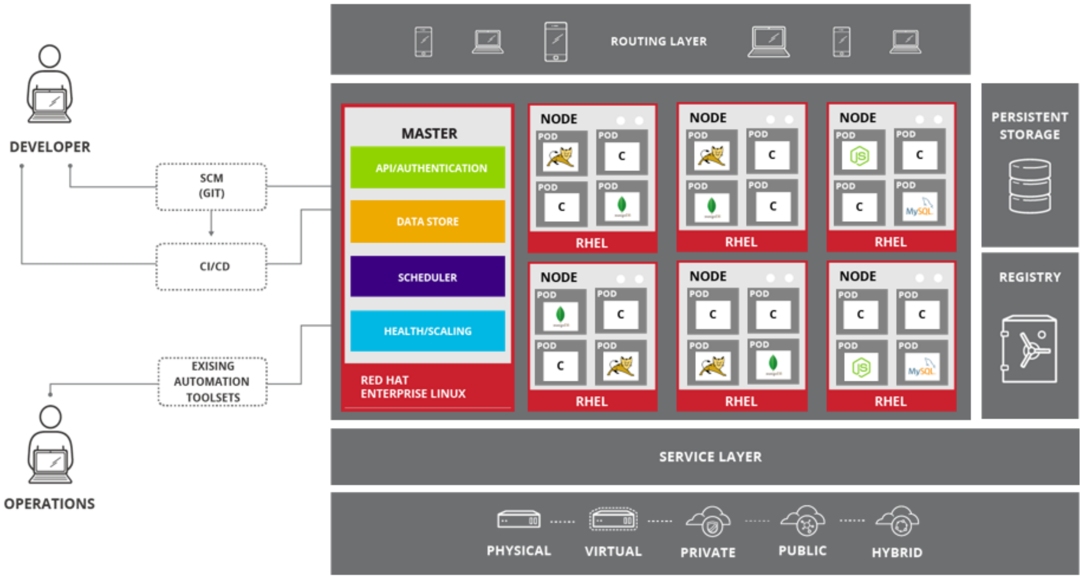

Figure 2: High Level OpenShift Architecture

Users – User (operators, developers, application administrators) access to OpenShift is provided through standard interfaces including the Web UI, CLI, and IDEs. These interfaces go through the authenticated and RBAC-controlled API. Users do not require system-level access to any of the OpenShift hosts, even for complicated application debugging and troubleshooting tasks.

There are three types of user accounts that can exist in an OpenShift environment: regular users, system users, and service accounts.

Regular users are created automatically in the system upon first logon or via the API. Most interactive OpenShift users, including operators, developers, and application administrators, will be represented by this type of user account.

System users are typically created when the infrastructure is first defined and are primarily purposed for enabling the infrastructure to interact with the API securely. System users typically include a cluster administrator (access to everything), a per-node user, users used by routers and registries, and various others. There is also an anonymous system user that is used by default for unauthenticated requests.

Service accounts are special system users associated with projects. These accounts are usually created when a project is first initiated. Project administrators can create more service accounts for defining access to the contents of each project.

Projects – A project is a Kubernetes namespace with additional OpenShift annotations and metadata. It is the central vehicle by which access to resources for regular users is managed. A project allows a community of users to organize and manage their content in isolation from other communities.

For more information on OpenShift concepts, features, and functions, please refer to Red Hat's product documentation.

OPENSHIFT SECURITY AND COMPLIANCE RELEVANT FEATURES, FUNCTIONS, AND DESIGN

OpenShift runs on RHEL or Atomic Host and makes use of existing security features built into these operating systems. Red Hat manages the RHEL and Atomic Host packages and provides a trusted distribution. Red Hat is committed to responsive action to security vulnerabilities. With a similar approach, Red Hat proactively manages the OpenShift platform. The security of OpenShift includes and utilizes hardened technologies such as SELinux; process, network, and storage separation; proactive monitoring of capacity limits (CPU, disk, memory, etc.); and encrypted communications for infrastructure transport including SSH, SSL, etc. Additionally, OpenShift provides secure authentication and authorization options to support organization compliance requirements.

For initial implementation and continuing maintenance, OpenShift components, including OS-level packages, can be acquired from Red Hat-controlled distribution points, whether with signed RPMs or through the Red Hat Container Catalog. Red Hat provides enterprise support for Red Hat Enterprise Linux, Atomic Host, and OpenShift Container Platform supported versions including vulnerability patches and updates.

SCOPE AND APPROACH FOR REVIEW

The understanding of OpenShift on RHEL and Atomic Host and their combined capabilities was gained through product specification, installation, configuration, administration, and integration documentation provided by Red Hat and generally made available from Red Hat's public facing web site. Coalfire further conducted interviews and engaged in live product demonstrations with Red Hat personnel. For live product demonstration purposes, OpenShift was also implemented on RHEL in the Coalfire lab environment to provide hands-on testing and and analysis of the system's capabilities to support compliance.

Coalfire's review of OpenShift on RHEL and Atomic Host began with a general alignment of the applicability of the technology against the high-level PCI DSS control objectives. This was further narrowed down to specific requirements that were considered applicable to either OpenShift, RHEL, or Atomic Host. An analysis of capability for the reviewed technology to address the applicable requirements was then conducted. This analysis primarily focused on what a PCI DSS QSA might review when following the PCI DSS testing procedures during an assessment of applicable requirements.

SCOPE OF TECHNOLOGY AND STANDARD TO REVIEW

Coalfire was tasked by Red Hat to review OpenShift as deployed on either RHEL or on Atomic Host. The primary focus of review included the components, features, and functionality of OpenShift along with the supporting underlying OS features and functionality when the OpenShift components are deployed on RHEL or Atomic Host. Coalfire did not include in the assessment application pods or containers that a payment entity may deploy on OpenShift. Containers that were deployed in the lab environment were strictly used for the purposes of demonstrating the platform's orchestration, deployment, and management capabilities. Furthermore, Coalfire did not make an assessment of publicly available image registries or repositories that may be used for acquiring applications, services, dependencies, or other elements to be hosted on or used for the setup OpenShift.

For this review, Coalfire included requirements from PCI DSS Requirements and Security Assessment Procedures Version 3.2 April 2016 publication available from https://www.pcisecuritystandards.org. For a broader understanding of the requirements and their applicability to technical solution implementation, Coalfire also reviewed supporting documentation provided by the PCI SSC including the approach to assessment guidance, a sample report on compliance form, and self-assessment questionnaires. Applied understanding of PCI DSS requirements and recommendations was also made available from PCI DSS guidance documents on relevant topics such as scoping and segmentation, best practices for maintaining PCI DSS compliance, risk assessment guidance, cloud computing, and virtualization.

COALFIRE EVALUATION METHODOLOGY

Coalfire initially examined the PCI DSS requirements and identified them as either procedural or technical. Qualification of a requirement as procedural or technical was based on a review of the requirement narrative, testing procedures, and guidance.

"Non-technical", operational requirements that include definition and documentation of policies, procedures, and standards were not considered applicable to the technical solution. Likewise, "non-technical" requirements including operational procedures that describe manual processes were not assessed against the technology. Examples of this type of "non-technical" requirement included maintenance of facility visitor logs, verification of an individual's identity prior to granting physical or logical access, performance of periodic physical asset inventories, or generation of network topology or flow diagrams.

Technical requirements were then assessed to determine applicability to the solution and/or solution components. Where achievement of the requirement objectives was more likely to be met using an external and non-adjacent mechanism, the requirement was determined to be "not applicable" to the assessed technology. Examples of requirements that Coalfire determined not to be applicable to OpenShift, RHEL, or Atomic Host included the use of encryption key management, wireless networking, technical physical access controls, and antivirus solutions.

Where the requirement was qualified as applicable, Coalfire further assessed the capability of the solution to address the requirement. For applicable requirements, Coalfire designated a qualitative category of capability including whether the solution was determined to fully support the requirement, partially support the requirement, or was unable to support the requirement. In cases where the requirement was determined to be applicable but unsupported, additional thought for the use of third-party solutions could be considered.

Each applicable requirement is described in the table in the following section. This table includes the findings of applicability along with a short narrative describing the capability. The table uses Harvey Balls to describe the qualitative applicability of the technology to the requirement.
(see: https://en.wikipedia.org/wiki/Harvey_Balls)

The following key describes the meaning for the Harvey Balls.

DESIGNATION	KEY	MEANING
Fully Supported	●	The evaluated technology has the means, through configuration or by design, to fully support the requirement.
Partially Supported	◐	The evaluated technology has the means, through configuration or by design, to partially support the requirement. Additional technology or organizational policy, procedures, or standards may be required to fully meet the requirement objective.
Not Supported	○	The evaluated technology does not have the means, through configuration or by design, to support or meet the requirement objective. The technology relies on an external technology, system, or organizational procedure to meet the requirement objective.
Does Not Apply	N/A	The requirement, though it may be technical in nature, is not applicable to the evaluated technology. Meeting the requirement is completely achieved independently of the evaluated technology.
Non-Tech	Non-Tech	The requirement is purely operational requiring a defined and documented policy, procedure, or standard. This may also include operational, non-technical procedures such as the implementation of training, maintaining of paper logs, or performance of personnel identity verification.

Table 1: Applicability Key

RED HAT OPENSHIFT APPLICABILITY TO PCI DSS V3.2

For the sake of brevity, requirements that were designated as not applicable or non-technical were removed from the following detail.

PCI DSS 3.2 COMPLIANCE APPLICABILITY DETAIL

ID	REQUIREMENT	SUPPORT	NARRATIVE
2.1	Always change vendor-supplied defaults and remove or disable unnecessary default accounts before installing a system on the network. This applies to ALL default passwords, including but not limited to those used by operating systems, software that provides security services, application and system accounts, point-of-sale (POS) terminals, payment applications, Simple Network Management Protocol (SNMP) community strings, etc.	●	OpenShift, RHEL, and Atomic Host do not come preconfigured with a default password. An initial administrator account must be created for UI or API access to OpenShift once it is deployed. The password for the host's root account must be initially set at the time of deployment. There is a default OpenShift administrator account, only accessible via a shell session for the root account on either RHEL or Atomic Host systems acting as Masters. This root account is a necessary account for OS operations as well as some OpenShift host interaction, and, as stated previously, has no default password. SNMP community strings for the host OS can be modified as necessary to support the payment entity's standard.

ID	REQUIREMENT	SUPPORT	NARRATIVE
2.2.1	Implement only one primary function per server to prevent functions that require different security levels from co-existing on the same server. (For example, web servers, database servers, and DNS should be implemented on separate servers.) Note: Where virtualization technologies are in use, implement only one primary function per virtual system component.	●	OpenShift provides the means to separate application workloads within the OpenShift cluster onto separate hosts within the cluster. This allows the payment entity to separate CDE from non-CDE pods as well as DMZ and internal pods to separate hosts. Additionally, the OpenShift etcd data store can be hosted separately from the OpenShift Master. Where virtualization platforms are used to host OpenShift, additional consideration may be necessary for placement of OpenShift nodes in accordance with requirement 2.2.1
2.2.2	Enable only necessary services, protocols, daemons, etc., as required for the function of the system.	●	When using RHEL as the base OS for the OpenShift, system hardening must be performed by the payment entity to ensure that only the necessary services, protocols, daemons, etc. are installed and running according to the entity's standards. Atomic Host is purpose-built to support container-based workloads and comes pre-configured with only the necessary elements required to run containers. During the setup process, OpenShift only installs what is necessary to support the platform.
2.2.3	Implement additional security features for any required services, protocols, or daemons that are considered to be insecure. Note: Where SSL/early TLS is used, the requirements in Appendix A2 must be completed.	●	Communication at the Management and Control planes is performed using TLS 1.2 encryption. However, the OpenShift software-defined networking (SDN) does not natively encrypt inter-pod/container communication between nodes. Where this level of security may be desired, the use of a third-party SDN solution may be required to support cross-container communication.
2.2.4	Configure system security parameters to prevent misuse.	◐	The payment entity must determine if their documented standards for security parameters are in alignment with OpenShift's capabilities. This will be determined on a payment entity by payment entity basis. The configuration of RHEL and Atomic Host is capable of being configured and hardened per best practices. For OpenShift, the payment entity may be required to adjust standards in support of OpenShift's requirements, such as the use of SELinux in enforcing mode on Linux hosts and the underlying use of the iptables firewall and/or firewalld.

ID	REQUIREMENT	SUPPORT	NARRATIVE
2.2.5	Remove all unnecessary functionality, such as scripts, drivers, features, subsystems, file systems, and unnecessary web servers.	◐	RHEL is configurable to remove unnecessary functionality per industry standards and best practices. This includes the implementation of configuration hardening standards and strategies that are prescribed by the payment entity and/or available from other sources such as CIS. Atomic Host comes pre-hardened with the minimum necessary functionality to support running containers. All additional functionality available from OpenShift as layered on either RHEL or Atomic Host is necessary for the function of the platform and is described in detail in product documentation.
2.3	Encrypt all non-console administrative access using strong cryptography. Note: Where SSL/early TLS is used, the requirements in Appendix A2 must be completed.	●	The only way to access the OpenShift Master or other OpenShift cluster hosts for non-console administrative access is through SSH. Telnet is not enabled. Administrative access to the API or the Web User Interface occurs over SSL using TLS 1.2.
2.6	Shared hosting providers must protect each entity's hosted environment and cardholder data. These providers must meet specific requirements as detailed in Appendix A1: Additional PCI DSS Requirements for Shared Hosting Providers.	◐	Accommodations are in place with OpenShift to support the use of the platform by shared hosting providers. This allows the shared hosting provider to protect each entity's hosted environment and cardholder data. There are multiple considerations for separating tenants in a multi-tenant environment including: separating tenants to their own OpenShift projects, the use of multiple networks in support of multi-tenancy, and the use of a multi-segmented network. The latter involves the creation of separate nodes for each tenant with blocking of the SDN between tenant hosts to prevent workload communication across tenant boundaries.
6.4.1	Separate development/test environments from production environments, and enforce the separation with access controls.	●	Like tenant separation, OpenShift can be implemented in a way to allow for separation of development/test environments from production environments including the use of role-based and attribute-based access controls (RBAC/ABAC) to enforce separation. Access controls can be applied to projects where separate projects can be created for each of development and test environments.
6.4.2	Separation of duties between development/test and production environments	●	Access controls for the administration and management of OpenShift can be established to support the separation of duties; different levels of access can be granted to development and test projects and users as opposed to production projects and users.

ID	REQUIREMENT	SUPPORT	NARRATIVE
7.2	Establish an access control system(s) for systems components that restricts access based on a user's need to know and is set to "deny all" unless specifically allowed. This access control system(s) must include the following:	●	OpenShift, either on RHEL or Atomic Host, is capable of being configured to support the payment entity's standards for access control based on minimum access requirements. RBAC and ABAC controls can be implemented to establish authorization with a granular level of control for interaction with and/or administration of OpenShift and the underlying OS. No initial access is granted, except for access initiated during setup. It is recommended that quick setup or a custom setup be performed and anonymous access not be utilized. Though anonymous access does not authorize an individual to perform any actions, it is recommended to use a strict identity and access implementation.
7.2.1	Coverage of all system components.	●	Supports coverage of all systems components up to the container itself. Control for the application or service running in the container is not covered by OpenShift and would be the responsibility of the payment entity or application developer to implement.
7.2.2	Assignment of privileges to individuals based on job classification and function.	●	Roles can be established that are in alignment with the payment entity's standards and definition for job classifications and functions.
7.2.3	Default "deny-all" setting.	●	The solution can be implemented to deny-all access except for that which is explicitly permitted during the setup.
8.1.4	Remove/disable inactive user accounts within 90 days.	○	For OpenShift, it is recommended to use a third-party, external identity provider to manage user accounts for authentication to OpenShift. OpenShift can use secure LDAP or OAuth to integrate with third-party identity and authentication providers. All console, remote or local, access is handled by OS accounts. Accommodations exist with the OS to enable enforcement options for user accounts; however, the payment entity may consider integration with outside identity and access providers to support central management and improved access controls.
8.1.6	Limit repeated access attempts by locking out the user ID after no more than six attempts.	○	Recommended to use a third-party, external identity provider for identity and authentication management.
8.1.7	Set the lockout duration to a minimum of 30 minutes or until an administrator enables the user ID.	○	Recommended to use a third-party, external identity provider for identity and authentication management.

ID	REQUIREMENT	SUPPORT	NARRATIVE
8.1.8	If a session has been idle for more than 15 minutes, require the user to re-authenticate to re-activate the terminal or session.	○	OpenShift only requires re-auth when the token expires. The default lifetime of an access token is 24 hours. While the lifetime of an access token can be reduced through configuration to fifteen (15) minutes, doing so may prove exceedingly burdensome to the user of the system as this would require reauthentication upon token expiration regardless of session idle state. Authorization tokens expire every five (5) minutes, which requires authorization settings to be updated to reflect any changes in authorization that may be set during the session. As there is not currently a sufficient mechanism to support idle session timeout, it is recommended to set idle limits for authentication to the terminal or client (user) device. This could prevent someone from walking up to an unoccupied terminal and hijacking the user's open session. Outside of this commonly implemented control, the payment entity may choose to implement an API gateway to provide more control to API sessions.
8.2	In addition to assigning a unique ID, ensure proper user-authentication management for non-consumer users and administrators on all system components by employing at least one of the following methods to authenticate all users: Something you know, such as a password or passphrase; something you have, such as a token device or smart card; or something you are, such as a biometric.	○	Recommended to use a third-party, external identity provider for identity and authentication management.
8.2.1	Using strong cryptography, render all authentication credentials (such as passwords/phrases) unreadable during transmission and storage on all system components.	◐	Along with the third-party identity and authentication provider, OpenShift, RHEL, and/or Atomic Host provide the means to support encryption of the authentication process to the system component.
8.2.3	Passwords/passphrases must meet the following: Require a minimum length of at least seven characters. Contain both numeric and alphabetic characters. Alternatively, the passwords/ passphrases must have complexity and strength at least equivalent to the parameters specified above.	○	Recommended to use a third-party, external identity provider for identity and authentication management.
8.2.4	Change user passwords/passphrases at least once every 90 days.	○	Recommended to use a third-party, external identity provider for identity and authentication management.

ID	REQUIREMENT	SUPPORT	NARRATIVE
8.2.5	Do not allow an individual to submit a new password/passphrase that is the same as any of the last four passwords/passphrases he or she has used.	○	Recommended to use a third-party, external identity provider for identity and authentication management.
8.2.6	Set passwords/passphrases for first-time use and upon reset to a unique value for each user, and change immediately after the first use.	○	Recommended to use a third-party, external identity provider for identity and authentication management.
8.3	Secure all individual non-console administrative access and all remote access to the CDE using multi-factor authentication. Note: Multi-factor authentication requires that a minimum of two of the three authentication methods (see requirement 8.2 for descriptions of authentication methods) be used for authentication. Using one factor twice (for example, using two separate passwords) is not considered multi-factor authentication.	○	Recommended to use a third-party, external identity provider for identity and authentication management.
8.3.1	Incorporate multi-factor authentication for all non-console access into the CDE for personnel with administrative access. Note: This requirement is a best practice until January 31, 2018, after which it becomes a requirement.	○	Recommended to use a third-party, external identity provider for identity and authentication management.
8.3.2	Incorporate multi-factor authentication for all remote network access (both user and administrator, and including third-party access for support or maintenance) originating from outside the entity's network.	○	Recommended to use a third-party, external identity provider for identity and authentication management.
8.5	Do not use group, shared, or generic IDs, passwords, or other authentication methods as follows: Generic user IDs are disabled or removed. Shared user IDs do not exist for system administration and other critical functions. Shared and generic user IDs are not used to administer any system components.	○	Recommended to use a third-party, external identity provider for identity and authentication management. It is further recommended that the root account credentials for the host operating systems be vaulted to prevent use. Furthermore, all activities that are performed by "root" should be logged and monitored for anomalies, which may indicate a possible compromise.
10.1	Implement audit trails to link all access to system components to each individual user.	●	All access to the API or web console for OpenShift is logged, including the HTTP method being invoked. All console or non-console access to the underlying OS, whether using RHEL or Atomic Host, can be logged as well. The host will need to be configured for logging to support this requirement. Both support the implementation of an audit trail to link access to individual users, where users are supplied with unique identities.

ID	REQUIREMENT	SUPPORT	NARRATIVE
10.2	Implement automated audit trails for all system components to reconstruct the following events:	◐	See sub-requirements below.
10.2.1	All individual user accesses to cardholder data.	●	All access to the API or web console for OpenShift is logged. All console or non-console access to the underlying OS, whether RHEL or Atomic Host, can be logged. Typically, CHD would be contained and isolated within the operations of the pod and would not likely be accessible through the OpenShift API or web interface. Where persistent storage is used to support extended storage of CHD, access to the storage should be logged as well. OpenShift does not directly provide access to persistent storage, it merely provides a conduit to connect persistent storage volumes to associated containers.
10.2.2	All actions taken by any individual with root or administrative privileges.	●	All access to the API is logged, including access through the web interface. For RHEL and Atomic Host, actions taken by root or accounts with administrative privilege either with remote or direct console access can be logged.
10.2.3	Access to all audit trails.	◐	With OpenShift, a log file or journal can be specified for the destination of audit logs. Access to the logs would be made available from the OS level either with a root or administrator account. The OS could be configured to log access to the journal or log file. For better protection of audit trails, it is recommended to configure the system to direct logs to a qualified central syslog server or SIEM solution.
10.2.4	Invalid logical access attempts.	◐	Invalid logical access attempts, as they pertain to incorrect input of credentials for access to OpenShift, would most likely be logged by the identity and authentication provider, external to OpenShift. Unauthorized attempts to access system components, run unauthorized commands, or execute unauthorized privileged commands would be logged by the OS.
10.2.5	Use of and changes to identification and authentication mechanisms—including but not limited to creation of new accounts and elevation of privileges—and all changes, additions, or deletions to accounts with root or administrative privileges.	◐	For OpenShift, this will be primarily handled by the third-party external identity and authentication system that is integrated with LDAP or OAuth capabilities of OpenShift. Local OS accounts that are granted access to applications, root, or administrative access would be logged at the OS level.

ID	REQUIREMENT	SUPPORT	NARRATIVE
10.2.6	Initialization, stopping, or pausing of the audit logs.	●	Initializing, stopping, or pausing of audit logs generated by OpenShift are handled by the Master; however, stopping or pausing the logs requires stopping the OpenShift Master itself. Thus, auditing cannot be reconfigured or stopped without reconfiguring OpenShift, which will itself be logged.
10.2.7	Creation and deletion of system- level objects.	●	Creation and deletion of system level objects is logged by OpenShift (for OpenShift objects) and by RHEL or Atomic Host for OS objects.
10.3	Record at least the following audit trail entries for all system components for each event:	●	See below.
10.3.1	User identification	●	User identification is present as a logged artifact.
10.3.2	Type of event	●	The type of event is present as a logged artifact.
10.3.3	Date and time	●	Date and time is present as a logged artifact.
10.3.4	Success or failure indication	●	Success or failure of the action taken is a logged artifact.
10.3.5	Origination of event	●	The origination or source of the event is a logged artifact.
10.3.6	Identity or name of affected data, system component, or resource.	●	The name of the affected data, system component, or resource is a logged artifact.
10.4	Using time-synchronization technology, synchronize all critical system clocks and times and ensure that the following is implemented for acquiring, distributing, and storing time. Note: One example of time synchronization technology is Network Time Protocol (NTP).	●	Time-synchronization for OpenShift uses the system time. The NTP client can be installed and configured for the underlying host(s) OS during the OpenShift setup.
10.4.1	Critical systems have the correct and consistent time.	●	The use of time synchronization with the industry accepted NTP server(s) as a centralized and trusted time source helps to ensure that the systems have the correct and consistent time. The payment entity may want to use an outside monitoring solution to provide alerting and notification of drift.
10.4.2	Time data is protected.	●	Time keeping is configured at the OS layer. Access to the OS, can be strictly limited, controlled and logged to minimize risk and the possibility of compromise to time data.

ID	REQUIREMENT	SUPPORT	NARRATIVE
10.4.3	Time settings are received from industry-accepted time sources.	●	NTP can be configured to use an industry-accepted time source. Where the system is unable to directly access a time source from the Internet, the payment entity can implement an internal NTP server that can communicate with the external source. Internal protected devices can be configured to use the internal NTP server as a time source.
10.5	Secure audit trails so they cannot be altered.	●	Securing of audit trails occurs at the OS layer, where minimal access is granted to log or journal files. All access to these files is also logged. Coalfire additionally recommends the use of a centralized syslog or SIEM solution to provide greater protection of the critical audit trails.
10.5.1	Limit viewing of audit trails to those with a job-related need.	●	Using RBAC, view access to audit trails can be limited to those with a job-related need. OpenShift’s audit trails are stored in the OS filesystem. Access to audit trails would be handled at the OS level. Similar RBAC can be implemented at the OS level to limit privileged access. Coalfire also recommends the use of a centralized syslog server or SIEM solution dedicated to log management, security, analysis, and reporting.
10.5.2	Protect audit trail files from unauthorized modifications.	●	Modification of the audit trail would require filesystem-level access at the OS layer, either with RHEL or Atomic Host depending on the choice of OS. Limitation for access to the audit trails can be handled by the OS RBAC. Again, a central syslog or SIEM solution is recommended.
10.5.3	Promptly back up audit trail files to a centralized log server or media that is difficult to alter.	●	The system provides the means to integrate with a centralized log server for log shipping.
10.5.4	Write logs for external-facing technologies onto a secure, centralized, internal log server or media device.	●	For all system components, logs can be directed to a centralized, internal log server or media device.
10.5.5	Use file-integrity monitoring or change-detection software on logs to ensure that existing log data cannot be changed without generating alerts (although new data being added should not cause an alert).	○	Coalfire recommends using an external syslog server or SIEM solution with built in or integrated file-integrity monitoring and change detection.
10.7	Retain audit trail history for at least one year, with a minimum of three months immediately available for analysis (for example, online, archived, or restorable from backup).	○	Coalfire recommends using an external syslog server or SIEM solution for log management. These solutions are more apt to provide the level of retention necessary to support the payment entity’s analysis and reporting requirements as well as its retention requirements.

ID	REQUIREMENT	SUPPORT	NARRATIVE
11.5	Deploy a change-detection mechanism (for example, file-integrity monitoring tools) to alert personnel to unauthorized modification (including changes, additions, and deletions) of critical system files, configuration files, or content files, and configure the software to perform critical file comparisons at least weekly. Note: For change-detection purposes, critical files are usually those that do not regularly change, but the modification of which could indicate a system compromise or risk of compromise. Change-detection mechanisms such as file-integrity monitoring products usually come pre-configured with critical files for the related operating system. Other critical files, such as those for custom applications, must be evaluated and defined by the entity (that is, the merchant or service provider).	○	Coalfire recommends using an external change management solution configurable to identify and protect critical system configuration and component files and content and to alert when alterations occur.

Table 2: Compliance Applicability of Red Hat OpenShift on RHEL or Atomic Host

Regarding requirement 1, “install and maintain a firewall configuration to protect cardholder data”, it is expected that the payment entity will implement firewalls and routers at the edge of the network to protect the trusted network from untrusted external networks. Additionally, the payment entity may consider the use of additional routers, firewalls, and/or layer 3 switching to provide additional segmentation or isolation of networks within the trusted realm of the entity’s network. The design of the network should take into consideration the placement and use of OpenShift nodes. This includes the placement of cluster nodes with respect for the intended hosted pods and/or containers. Pods and/or containers may represent elements of services that typically exist within the DMZ such as a web service, portal, or web site. Other pods and/or containers may represent elements that traditionally exist on internal networks such as application or database. The placement and use of these nodes will be important regarding ensuring proper boundary protection between the untrusted networks, the DMZ, and trusted networks including separation of CDE and non-CDE where segmentation is applied for scope reduction.

COALFIRE OPINION

It is the opinion of Coalfire that Red Hat OpenShift Container Platform hosted on either Red Hat Enterprise Linux or Red Hat Enterprise Linux Atomic Host as reviewed can be effective in providing support for the outlined objectives and requirements of PCI DSS v3.2. Through proper implementation and integration into the payment entity’s infrastructure, OpenShift and the underlying Red Hat OS may be usable in support of CDE. The OpenShift control plane, including the Master, can be viewed as a Tier 2 system for scoping, while the OpenShift nodes that host CDE, including payment applications or other applications that process, transmit, or store CHD, would be considered as a Tier 1 system for scoping. The inclusion for assessment scoping for PCI DSS is conditional upon clearly defining network segmentation to separate CDE from non-CDE. OpenShift either on RHEL or Atomic Host provides process, network, and storage isolation for containers in the environment.

Coalfire’s opinion is based on observations and analysis of the provided documentation, interviews with Red Hat personnel, and hands on engagement with a lab environment. The provided opinion is dependent upon several underlying presumptions or caveats. These caveats include adherence to vendor best practices and hardening of configuration as supported by the system components. This solution should be implemented in alignment with the payment entity's mission, values, business objectives, general approach to security, and with the overall compliance program. Inclusion into the entity's overall compliance program includes considerations for supporting network infrastructure (build and maintain a secure network), segmentation efforts, physical security, personnel security, vulnerability testing and penetration testing, and ongoing risk and compliance evaluation.

ADDITIONAL INFORMATION, RESOURCES, AND REFERENCES

More information about PCI DSS including information about the PCI SSC, additional guidance, and assessment process can be found by visiting https://www.pcisecuritystandards.org.

More information about Red Hat OpenShift Container Platform including, architecture, implementation, and administration can be found on the OpenShift Container Platform documentation pages at https://docs.openshift.com/container-platform/latest/welcome/index.html .

More information about Red Hat Enterprise Linux including Release Notes, Installation, Migration, Systems Administration, Networking, and more can be found at https://access.redhat.com/documentation/en/red-hat-enterprise-linux/

More information about Red Hat Enterprise Linux Atomic Host including an overview of Atomic Host, the use of container on Red Hat platforms, release notes, installation and configuration guides, and other reference guides can be found at https://access.redhat.com/documentation/en/red-hat-enterprise-linux-atomic-host/

Hardening guides for Red Hat Enterprise Linux can be found on the Center for Internet Security web site, Red Hat's web site, and others.

BIBLIOGRAPHY

PCI Security Standards Council. (2016). *Payment Card Industry (PCI) Data Security Standard Requirements and Security Assessment Procedures Version 3.2.* (L. PCI Security Standards Council, Ed.) Retrieved from Official PCI Security Standards Council Site: https://www.pcisecuritystandards.org/documents/PCI_DSS_v3-2.pdf?agreement=true&time=1492112201359

PCI SSC. (2016, April). *PCI DSS v3.2.* Retrieved from www.pcisecuritystandards.org: https://www.pcisecuritystandards.org/documents/PCI_DSS_v3-2.pdf?agreement=true&time=1471887203375

ABOUT THE AUTHORS

Jason Macallister | Senior Consultant, Security Engineering, Coalfire Systems, Inc.

Mr. Macallister consults on Information Security and regulatory compliance topics as they relate to advanced infrastructure, emerging technology, and cloud solutions.

Chris Krueger | Principle, Security Engineering, Coalfire Systems, Inc.

As Principle, Mr. Krueger contributes as an author and thought leader on information security and regulatory compliance topics for Coalfire's clientele in the "new and emerging" technology areas.

Published June 2017

ABOUT COALFIRE

As a trusted advisor and leader in cybersecurity, Coalfire has more than 15 years in IT security services. We empower organizations to reduce risk and simplify compliance, while minimizing business disruptions. Our professionals are renowned for their technical expertise and unbiased assessments and advice. We recommend solutions to meet each client's specific challenges and build long-term strategies that can help them identify, prevent, respond, and recover from security breaches and data theft. Coalfire has offices throughout the United States and Europe. www.coalfire.com

Coalfire is solely responsible for the contents of this document as of the date of publication. The contents of this document are subject to change at any time based on revisions to the applicable regulations and standards (HIPAA, PCI-DSS et.al). Consequently, any forward-looking statements are not predictions and are subject to change without notice. While Coalfire has endeavored to ensure that the information contained in this document has been obtained from reliable sources, there may be regulatory, compliance, or other reasons that prevent us from doing so. Consequently, Coalfire is not responsible for any errors or omissions, or for the results obtained from the use of this information. Coalfire reserves the right to revise any or all of this document to reflect an accurate representation of the content relative to the current technology landscape. In order to maintain contextual accuracy of this document, all references to this document must explicitly reference the entirety of the document inclusive of the title and publication date; neither party will publish a press release referring to the other party or excerpting highlights from the document without prior written approval of the other party. If you have questions with regard to any legal or compliance matters referenced herein you should consult legal counsel, your security advisor and/or your relevant standard authority.

WP_Red Hat OpenShift Container Platform Product Applicability to PCI DSS 3.2_June 2017