

WHITE PAPER

# PCI ON GKE BLUEPRINT: REVIEW FOR PCI COMPLIANCE

DAN STOCKER | MS, CISSP, QSA  
ALLEN MAHAFFY | CISSP, QSA



C  A L F I R E .

North America | Europe

877.224.8077 | [info@coalfire.com](mailto:info@coalfire.com) | [Coalfire.com](https://Coalfire.com)

# TABLE OF CONTENTS

- Introduction ..... 3**
  - Assumptions ..... 3
  - Overview of PCI on GKE Blueprint..... 3
  - Scoping for PCI..... 5
  - Application of PCI Scoping to PCI on GKE Blueprint ..... 6
- Adoption Notes ..... 7**
  - Governance ..... 7
  - Network Security..... 7
  - Data Protection ..... 7
  - Secrets Management..... 7
  - Logical Access ..... 8
  - SDLC & Change Management ..... 9
  - Vulnerability Management ..... 9
    - Baseline Security ..... 9
    - Ongoing Vulnerability Management ..... 9
    - Anti-malware ..... 10
    - Logging, Monitoring, and Alerting..... 10
    - Scanning and Penetration Testing ..... 10
- Summary ..... 10**
- Appendix: PCI Requirements References ..... 11**

# INTRODUCTION

Google Cloud offers a GitHub project to demonstrate how to bootstrap a Payment Card Industry (PCI) compliant environment, using Google Cloud Platform (GCP) services and open source tools. Coalfire has been asked to evaluate the PCI on GKE Blueprint architecture, and as a starting point for further development, where PCI compliance is a target goal. This review is intended to outline the strengths of PCI on GKE Blueprint, provide guidance for areas where the adopter has responsibility for compliant choices, and outline the related topics essential to obtaining PCI compliance.

## ASSUMPTIONS

The PCI on GKE Blueprint should be understood as an infrastructure, into which applications can be built. It is a blueprint, not a turnkey solution for PCI compliance. In the same way that cloud computing is inherently about shared responsibility, this blueprint will require essential configuration choices by the adopter, in addition to building out of their solution (both custom infrastructure and solution stack). Google has also provided an example application (Microservices Demo) to complete the picture. Coalfire has integrated observations from that example into this review, where they offer useful context and illustrate key points.

Coalfire takes the view that compliance is an objective which must be managed to. Readers with long experience in PCI will be aware that cloud computing has not always been seen as compatible with the Data Security Standard (DSS). Thankfully, the Security Standards Council (SSC) has published [guidance](#) about cloud computing for PCI. Knowledge of the DSS will be essential to successful adoption of the PCI on GKE Blueprint in a PCI-compliant manner.

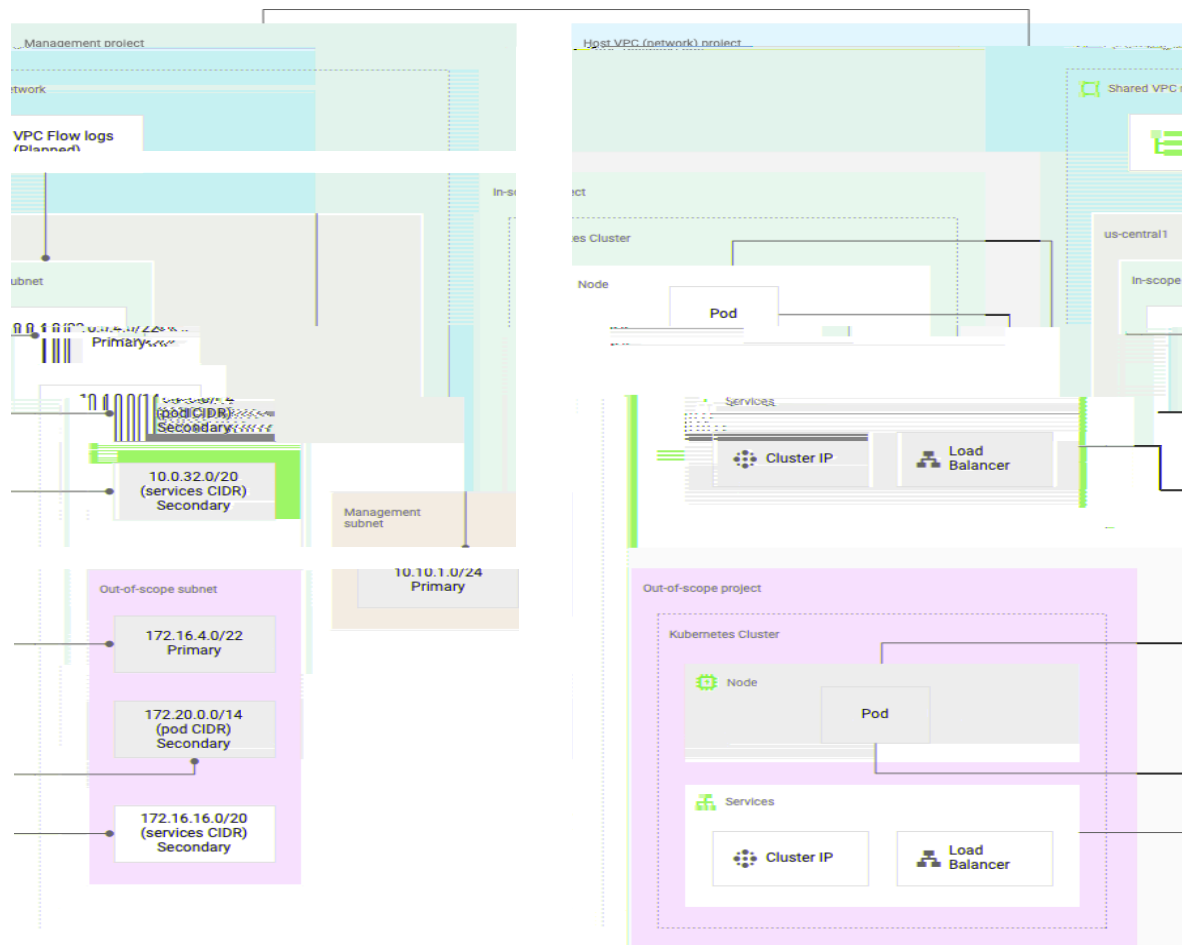
Please note that basic familiarity with the constituent technologies used to construct the PCI on GKE Blueprint will be necessary to successfully build and use it. Cloud native architecture is a dynamic, fast-moving area of cloud computing, equally adept at building production environments and supporting non-production prototyping. By the time this review is published, multiple (if not all) of the underlying technologies will have been updated. The adopter has responsibility for ensuring that the current available versions meet their need and are fit for purpose. To that end, this review will provide context for how the PCI requirements were applied, which should help facilitate ongoing understanding.

Throughout this review, where the text refers to PCI requirements, will be put in a distinct format **(3.4)** for easy recognition. Additionally, all references can be found in the Appendix, to allow correlation by requirement number.

## OVERVIEW OF PCI ON GKE BLUEPRINT

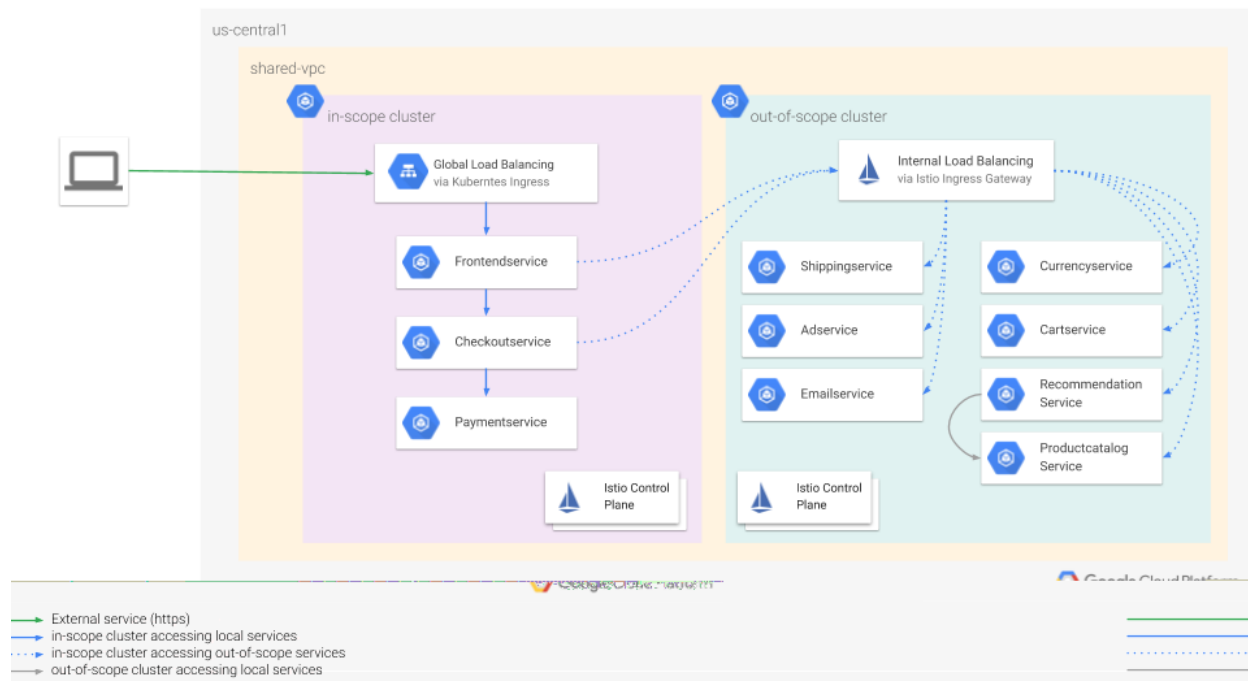
PCI on GKE Blueprint is structured as a set of Terraform scripts, along with associated resources necessary to perform an automated build of the infrastructure. The current version (*Diagram 2*) has the following high-level architecture:

- Shared *Virtual Private Cloud (VPC)*,
- Multiple subnets for the minimal set of expected purposes: in-scope, management, and out-of-scope,
- Micro-services application design, using *Istio Service Mesh* on *Google Kubernetes Engine (GKE)*, and
- Leverages *Google Identity Access Management (IAM)* to assign, align and limit access privileges.



**DIAGRAM 1: HIGH-LEVEL ARCHITECTURE**

A sample application, the “Microservices Demo” e-commerce site, was built atop PCI on GKE Blueprint to illustrate limited adoption, and facilitate this review (*Diagram 2*). Specific concerns that apply uniquely to the e-commerce site have been omitted from this review. The main reason for this is to concentrate on the architecture, and not allow example choices to distract from the central point. That said, it is important to note that the Microservices Demo application is a shell, and does not represent specific application design guidance beyond the basics.



**DIAGRAM 2: MICROSERVICES DEMO SHOP ARCHITECTURE**

There are two key areas where PCI on GKE Blueprint adopters will need to build out from the shell: data storage and security management. Google Cloud provides many options for data management and processing. These can be integrated into the PCI on GKE Blueprint design in a straightforward manner. Security management options are discussed below in the **Adoption Notes** section, in the context of the specific PCI requirements they support.

## SCOPING FOR PCI

PCI scoping is often over-simplified, due to legacy understanding from early in the history of the DSS. In a nutshell, scoping follows this process:

1. Identify the environment where cardholder data, Primary Account Number (PAN) or Sensitive Authentication Data (SAD) is stored, processed, or transmitted. This includes all networks, network devices, servers, and data stores.
2. With this set of environmental components, zoom out, to find the firewall-controlled perimeter of that environment. In a flat network, it may be everything. With careful segmentation, it may be optimal. This is the Cardholder Data Environment (CDE). It is often called Tier-1 for being the main subject of the assessment.
3. Of key importance, the CDE also includes any unrelated systems that happen to be in those networks. This is known as the scope “infection” rule, since those unrelated systems must also be assessed, as if they have cardholder data. The motivation of this rule is that their adjacency to cardholder data puts them in the same risk class.
4. There are many PCI requirements that involve functions that do not involve storing, processing, or transmitting cardholder data. These functions are considered to have a material impact on the

security of the CDE. These are often referred to as Tier-2 functions, and are not “infectious”. They can be assessed where they are, without drawing ancillary networks or systems into scope.

- The scope of a PCI assessment is the combination of Tier-1 and Tier-2.

## APPLICATION OF PCI SCOPING TO PCI ON GKE BLUEPRINT

The expected cardholder data flow for the Microservices Demo is confined to the CDE. Data flows from the CDE to the Out-of-Scope network (for fulfillment and ancillary operations) should not contain cardholder data, in order to remain out-of-scope. The adopter has choices for removing PAN data, including tokenization, hashing, and truncation. It is important to note that keeping truncated and hashed PAN together is considered equivalent to having full PAN, due to the relative cryptographic ease of reconstituting PAN from the other two. Likewise, encrypted PAN is good, but is still PAN, for scoping purposes.

The management cluster is where Tier-2 functions are housed (see the Adoption Notes section, below). Configuration for remote administrative traffic will be adopter responsibility. *Diagram 3* depicts the environment from a scoping perspective.

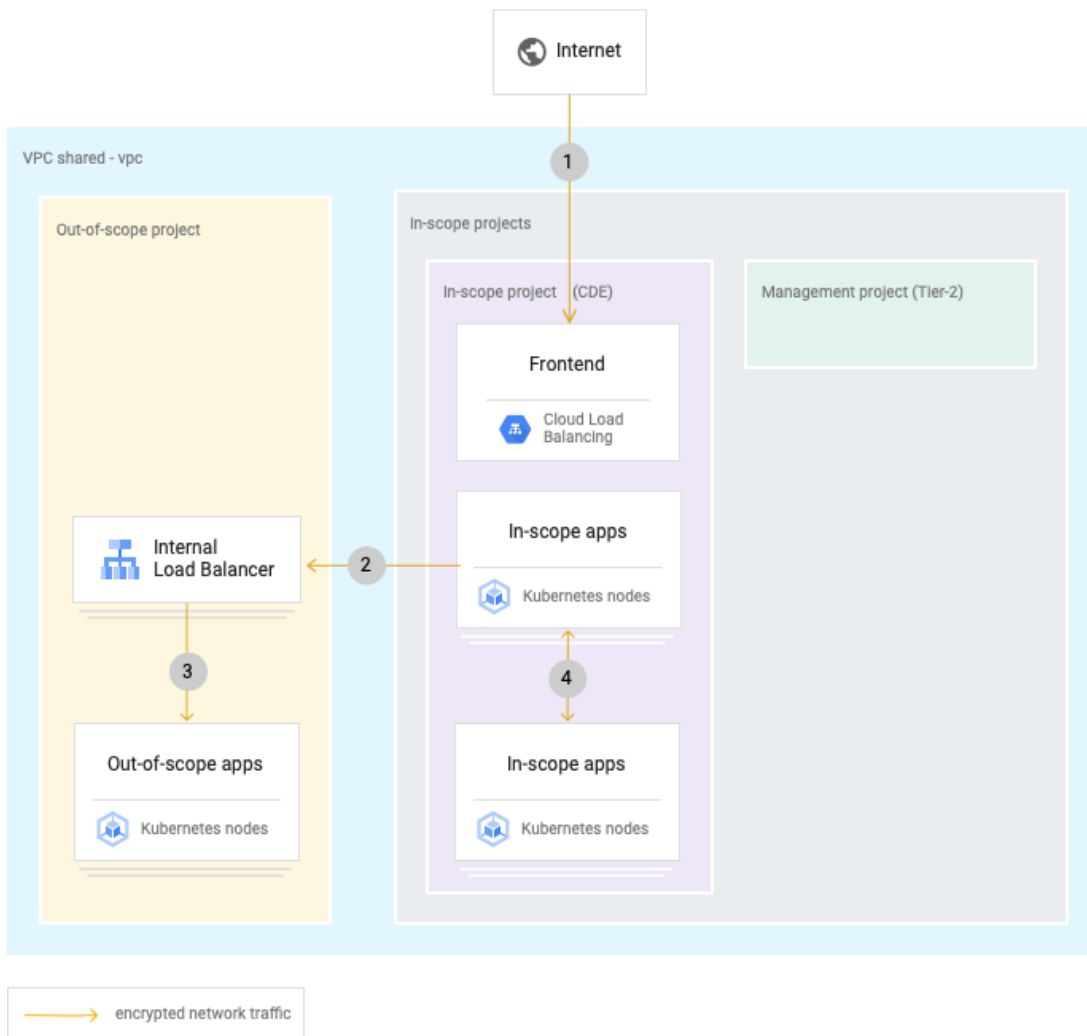


DIAGRAM 3: PCI SCOPING PERSPECTIVE

# ADOPTION NOTES

## GOVERNANCE

All PCI on GKE Blueprint adopters will require comprehensive governance for their PCI environments, including policies and procedures to support PCI topics that will eventually be assessed. The DSS embeds these throughout the DSS, adjacent to the relevant technical requirements. Established organizations with governance, risk, and compliance programs are advised to add PCI to their compliance targets, as part of the required PCI program management requirement (12.4.1). More general policies, standards, and procedures are bundled together under requirement 12. The most important governance process for PCI compliance is Change Management. This will be discussed below in the Software Development Lifecycle (SDLC) section.

Physical security for environments hosted in Google Cloud inherit PCI compliance, since Google manages those facilities and has been successfully assessed.

## NETWORK SECURITY

PCI on GKE Blueprint implements a shared VPC, with three subnets, aligned to Projects representing the CDE, Tier-2 function, and out-of-scope functions. VPC firewall rules (1.1.4) are used to limit access to and from each subnet. The CDE accepts external traffic, but none from the out-of-scope subnet. No egress rules are set, which results in no regulation by default. Per requirement 1.3.4, outbound traffic must be explicitly authorized. These rules will be adopter responsibility.

Adopters may also further segment their environment for more granular scope control. Additional subnets can be used, but with *GKE* and *Istio*, there are additional options. *GKE* use of namespaces and network policies to partition workloads within and between clusters. *Istio*'s Service Mesh concept extends this further to allow nano-segmentation via the sidecar proxy pattern. Groups of individual containers may be logically organized into cooperative sets for application-level logic purposes.

## DATA PROTECTION

PCI mandates protection of cardholder data while at-rest (3.4), and when transmitted across open, public networks (i.e., the internet). PCI on GKE Blueprint does not propose data storage as part of the blueprint, but Google Cloud offers multiple options for a variety of data types and purposes and provides detailed [Best Practice guidance](#) for their use. Adopters should understand that Google Cloud's default behavior is to [encrypt all customer data at rest](#). This requires no action on the part of the customer. Adopters with risk management (or regulatory) mandates to manage their own data encryption may do so with Google's [Cloud KMS](#) key management service (including [Cloud HSM](#)), or by using a third-party tool.

Under requirement 4, applications built on PCI on GKE Blueprint must ensure that incoming internet traffic is secured with strong cryptography (e.g., TLS1.2). The current version of the DSS (3.2.1) does not mandate encryption of traffic inside the CDE; even traffic with sensitive cardholder data. This is expected to change under the next version of the DSS (4.0), which is expected in 2021. *Istio*'s use of mutual-TLS (mTLS) among its constituent services is a future-proof design in this regard.

## SECRETS MANAGEMENT

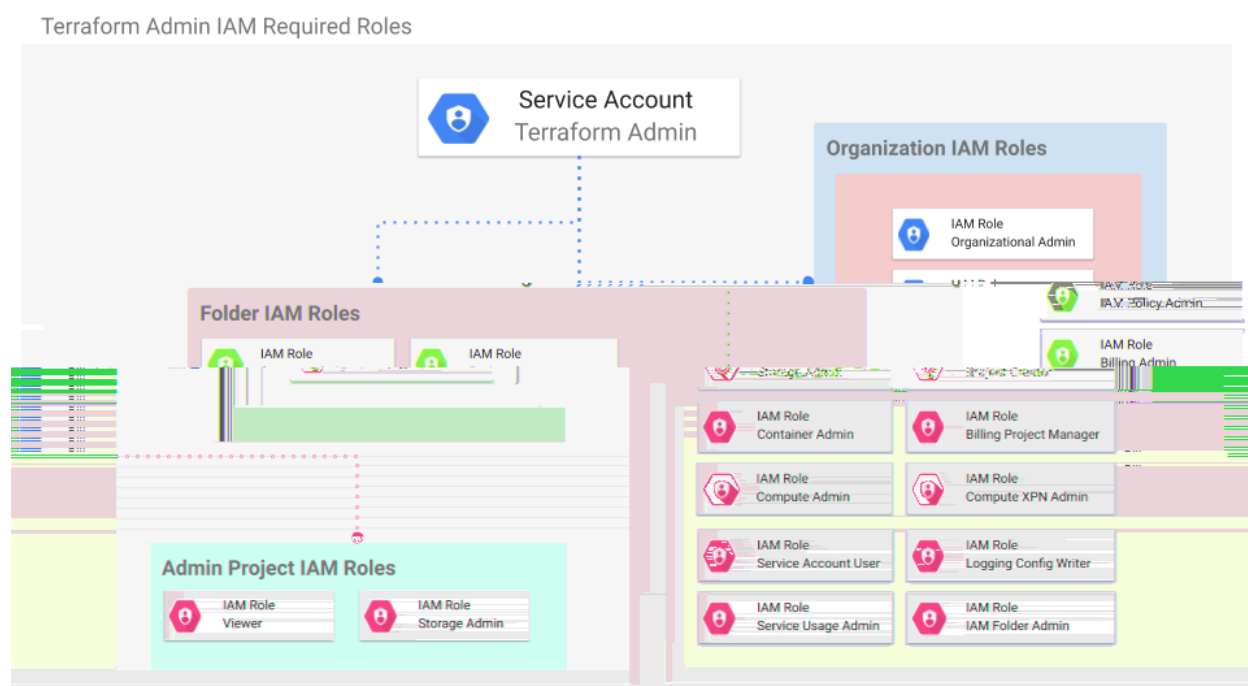
Other than cardholder data, the only secrets that must be protected are passwords (8.2.1). These must be protected in-transit, as well as at-rest. Google's *Cloud IAM* service is a PCI compliant building block for identity management, including password storage. In a cloud context, use of service accounts mandates maintaining a variety of credentials that are equivalent to passwords, but are not a concern at the present time. API keys, access tokens, and one-time passwords (OTP) for service accounts are not subject to PCI

scrutiny. Coalfire has specific guidance from the SSC that the logical access requirements (7.x and 8.x) do not apply to service accounts or machine-to-machine communication. The DSS intends to cover only interactive sessions (i.e., by humans). Coalfire notes that this limitation is expected to go away under DSS 4.0.

For others secrets, there are multiple options available. GKE native etcd [secret storage](#) is encrypted at-rest, just as all Google Cloud storage is. *GKE also* offers integration with *Google Secrets Manager*, which has a rich API for secrets management, and [Cloud KMS](#). Of course, application-level encryption of secrets is also possible, leveraging the same ecosystem used to manage encryption for data protection.

## LOGICAL ACCESS

Given the multiple abstractions in use, logical access must be managed on several levels. Starting from the base, *Cloud IAM* must be configured for use, to allow access into GCP. PCI on GKE Blueprint is built expecting a Google Cloud [Organization](#) (which can be created through *Cloud Identity* or *Gsuite*) with *Organization Admin* and *Folder Admin* permissions. The Terraform build process requires the Terraform admin account have specific *IAM* roles (Diagram 4)



**DIAGRAM 4: TERRAFORM ADMIN REQUIRED IAM ROLES**

Application build out will require further specification of users, accounts, and roles. Google provides, and Coalfire recommends, a comprehensive set of [Best Practices](#) for *Cloud IAM*. In particular the granular pre-defined IAM roles represent embedded best practices. Custom roles are possible, but should be considered carefully for unintended side-effects. At the Kubernetes level, users and *IAM* policies must be configured to allow *GKE* to work properly. PCI on GKE Blueprint sets up the default set with least privilege in mind.

In order to use *Google Service Accounts* (GSAs) within GKE, it is necessary to manage them as secrets. Adopters may find the recently released *Workload Identities* feature useful. *Workload Identities* is a GKE abstraction that allows *Kubernetes Service Accounts* (KSAs) to act as GSAs, eliminating the need for explicit secrets management.



## SDLC & CHANGE MANAGEMENT

PCI on GKE Blueprint adopters will need to establish a software development lifecycle (SDLC). Container-oriented architectures benefit most from a continuous integration / continuous deployment (CI/CD) pipeline. The signal value of “everything-as-code”, made possible in the cloud, means software management processes take on greater importance, and offer opportunities for efficiency in operational support. This is the essence of the DevOps approach.

PCI expects a series of checkpoints for code that will end up in production. The development process starts with well-trained software engineers (6.5), knowledgeable in secure coding techniques (6.5.x). Proposed changes be assessed for impact (6.4.5.4), in particular for impact on PCI compliance (6.4.5.3.b). Code must be reviewed (6.3.2) and approved (6.4.5.2), before being accepted for integration. Testing should be performed, and back-out procedures (6.4.5.4) known before pushing that code into production. These steps must be performed in a non-production environment (6.4.1). Thus, adopters will need to adapt the PCI on GKE Blueprint architecture to offer (at least) one environment that mimics production.

Beyond efficiency, further integrating security into the CI/CD pipeline (“shifting left”) offers tangible benefits for compliance, not just security. For container-based architectures, it is possible to scan release images for vulnerabilities. Coalfire recommends the best practice of considering production immutable (i.e., no manual changes are made). All changes proceed through the pipeline.

One useful side-effect of an immutable production environment comes at time of assessment. PCI allows sampling for server populations, but container architectures present practical difficulties. Some environments are not instrumented for manual access to prod. Most container workloads have ephemerality that makes interactive sessions impractical. Where prod can be credibly demonstrated to be immutable, the master images may be used as a proxy for prod. This is very useful for a variety of PCI requirements. Vulnerability scans (11.2) may be performed. Penetration testing can likewise benefit.

If the adopter has not established a CI/CD pipeline of their own, Google Cloud offers multiple services to support this foundational process. Google [Cloud Build](#) offers all of the above features. Google [Container Registry](#) plugs in to Cloud Build and offers vulnerability scanning for images ([Container Analysis Vulnerability Scanning](#)). Combined with [Binary Authorization](#) (a GKE feature), adopters can deploy with maximum confidence that production reflects the intended state of deployment.

## VULNERABILITY MANAGEMENT

### Baseline Security

Starting out on the right foot has obvious benefits. Google and the Center for Internet Security have published hardening (2.2) and benchmark guides for [Kubernetes](#) and [GKE](#) (hardening, benchmark). Adopters can also use *Cloud Security Command Center* for asset discovery and inventory purposes (2.4).

### Ongoing Vulnerability Management

Adopters have two alternatives to meet a security assurance objective for public-facing web applications (6.6): either an automated technical solution, or application vulnerability assessments before release. GCP provides three tools to help with this mandate. The Blueprint implements [Google Cloud Armor](#) web application firewall (WAF), with a default ruleset covering some of the OWASP Top 10 Risks. With rules appropriate to the architecture, use of Cloud Armor will exceed PCI requirements. If application vulnerability assessments are preferred, [Cloud Security Scanner](#) offers a platform for vulnerability scanning of release candidates, and integrates with CI/CD pipelines of various types.

PCI also requires file-integrity monitoring, which can be implemented in more than one way. Where adopters repave their environments at-least weekly, the intent of the requirement (11.5) is met. [Security](#)

[Command Center](#) offers a dashboard view of the current state of monitoring and latest results. This functionality is useful for meeting the critical security control monitoring intent of requirement [10.8](#).

## Anti-malware

The adopter will make choices for container OS, which will determine options for compliance with requirement 5. Operating systems that are not commonly-affected by malware may be exempted from this requirement ([5.1.2](#)) but not without a risk assessment. Additionally, multiple stations in the pipeline can help manage malware.

## Logging, Monitoring, and Alerting

Adopters have multiple options for setting up a logging, monitoring, and alerting architecture. Standard Google Cloud logging (Cloud Audit Logs, VPC Flow Logs, and Cloud Monitoring API) should be supplemented by applications logs to create a comprehensive audit trail. PCI on GKE Blueprint adds Istio logging to that set, and uses the OpenTelemetry API/SDK. Google's Operations Suite offers a Cloud Logging service to marshal all of these streams and perform essential monitoring and alerting. Alternative tools include Splunk and Prometheus (which has cloud/container roots). Cloud and container-based architectures often opt to leverage their SIEM for intrusion detection (IDS). Additionally, [GCP Packet Mirroring](#) can be used with a 3rd party IDS tool such as [Palo Alto Networks](#) to demonstrate PCI compliance for requirement [11.4](#).

## Scanning and Penetration Testing

To achieve PCI compliance, successful vulnerability scanning and penetration testing must be performed on a regular schedule. Scanning is expected from both the internal perspective, as well as externally. The latter must be done by an Approved Scan Vendor (ASV), which is a PCI-specific program. ASV scans cannot be accomplished with Google services, since ASVs are independent third-parties. Internal scanning can be done with traditional tools (i.e., Nessus), or by layering vulnerability scanning with Google tools (*Cloud Security Scanner* and *Container Analysis Vulnerability Scanning*). For penetration testing, Coalfire recommends engaging a vendor with experience testing cloud and container-based architectures.

## SUMMARY

Coalfire's review of PCI on GKE Blueprint finds value for the starter architecture for Google Cloud customers building an environment that will need PCI compliance. It effectively illustrates use of GKE in a segmented VPC, for control of scope. Other key Google Cloud technologies can be used to extend this disciplined approach. Adopters will have many implementation choices, and options for Google Cloud services to leverage, including *IAM*, *Cloud KMS*, *Istio*, *Cloud Storage*, and the suite of Google Cloud security services (*Cloud Armor*, *Security Command Center*, *Operations Suite*, *Cloud Scanner*, etc).

## APPENDIX: PCI REQUIREMENTS REFERENCES

**Note:** these brief summaries of PCI requirements do *not* capture all aspects of the requirements. They are intended to gather the guidance for the topic narratives, into a more easier referenced format. Needless to say, this list is not the full set of PCI requirements. When in doubt, consult a QSA for qualified advice on PCI requirements.

- 1.1.4 *Firewalls required at perimeter of CDE:* VPC Firewall rules can be used to satisfy PCI requirements for stateful inspection of traffic.
- 1.3.4 *Outbound traffic must be explicitly authorized:* Google Cloud customers are responsible for setting, and managing, appropriate outgoing traffic VPC Firewall rules.
- 2.2 *Hardening standards for all in-scope systems:* Google Cloud offers multiple guides, which customers must adapt to their particular purposes.
- 2.4 *Keep an inventory for all in-scope systems:* Cloud Security Command Center can help Google Cloud customers identify and manage their inventories.
- 3.4 *Data protection for Primary Account Numbers (PAN):* Google Cloud encrypts all at-rest customer data by default. Customers also have options for explicit management of data protection with *Cloud KMS* and other tools.
- 4 *Cardholder data transmitted over open, public networks must be protected by strong cryptography:* Google Cloud offers customers the option to use TLS 1.2. Doing so is the customer's responsibility.
- 5.1.2 *Exception for anti-malware on operating systems not commonly-affected by malware:* adopters may exercise this option, if supported by a risk assessment.
- 6.3.2 *Performing code reviews for all custom-developed software* is the responsibility of the adopter.
- 6.4.1 *Development and test environments must be separate from production:* Google Cloud customers must design and implement separate environments for development and testing purposes. These must be isolated from production.
- 6.4.5.2 *Authorized approval for all changes* is the responsibility of the adopter.
- 6.4.5.3.b *PCI compliance testing for all changes* is the responsibility of the adopter.
- 6.4.5.4 *Preparing backout procedures for all changes* is the responsibility of the adopter.
- 6.5 *Ongoing training in security software development* is the responsibility of the adopter.

- 6.5.x** *Knowledge of secure coding techniques* is the responsibility of the adopter.
- 6.6** *Web-application firewall, or Application Vulnerability Assessments:* Google Cloud customers may either use *Cloud Armor*, or integrate vulnerability scanning of release images into their CI/CD pipeline.
- 7.x** *Authorization best practices* are implemented in *Cloud IAM*, but must be used properly by Google Cloud customers.
- 8.x** *Specific minimum authentication standards* are implemented in *Cloud IAM*, but must be used properly by Google Cloud customers.
- 8.2.1** *Passwords must be protected by strong cryptography in transmission:* *Cloud IAM's* PCI compliance supports Google Cloud customers. This mandate is not applicable to service accounts.
- 10.8** *Timely detection of critical control failures:* *Cloud Security Command Center* can help Google Cloud customers identify control failures with a dashboard view.
- 11.2** *Internal and external vulnerability scanning:* Google Cloud customers may layer use of *Cloud Security Scanner* and *Container Analysis Vulnerability Scanning* to help satisfy the internal scanning mandate. External scanning must be performed by an ASV.
- 11.3** Penetration testing is an adopter responsibility.
- 11.4** *An Intrusion Detection System* can be assembled by marshaling the many sources of log data available from Google Cloud, along with application-level logs, into *Operations Suite* and/or *BigQuery* for monitoring and alerting.
- 11.5** *File Integrity Monitoring* can be accomplished with weekly repaving, and monitoring with *Security Command Center*.
- 12.4.1** *Establishing PCI Compliance Program* is adopter responsibility, and key to effective and efficient management of PCI compliance.

## ABOUT THE AUTHORS

### **Dan Stocker** | Director, Cloud Advisory

Dan established the Cloud Advisory practice at Coalfire, which grew out of his work advising and assessing the major cloud service providers (AWS, Azure, Google, Salesforce, IBM, and Oracle). Methodology developed in that work has been applied to multiple verticals and to extend security and privacy compliance understanding to leading-edge cloud technology (e.g., containers).

### **Allen Mahaffy** | Principal, Cloud Advisory

Allen works in the Cloud Advisory practice performing advisory and assessments for major cloud service providers and large enterprises. His other experience in cloud include assessing and analyzing cloud micro-service architectures, container orchestration, security and compliance automation, and various emerging cloud technologies.

Published April 2020.

## ABOUT COALFIRE

Coalfire is the trusted cybersecurity advisor that helps private and public sector organizations avert threats, close gaps, and effectively manage risk. By providing independent and tailored advice, assessments, technical testing, and cyber engineering services, we help clients develop scalable programs that improve their security posture, achieve their business objectives, and fuel their continued success. Coalfire has been a cybersecurity thought leader for nearly 20 years and has offices throughout the United States and Europe. For more information, visit [Coalfire.com](https://www.coalfire.com).

Copyright © 2020 Coalfire Systems, Inc. All Rights Reserved. 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.

April 2020