HITRUST COMPLIANCE REVIEW

MICROSOFT AZURE SECURITY AND COMPLIANCE BLUEPRINT – HITRUST HEALTH DATA AND ARTIFICIAL INTELLIGENCE (AI)

REVIEW AND GUIDANCE FOR IMPLEMENTATION

MICHAEL T. WILLIAMS
ALISON HEIDT

COALFIRE

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

TABLE OF CONTENTS

TABLE OF FIGURES

EXECUTIVE SUMMARY

The Microsoft Azure Security and Compliance Blueprint – HITRUST Health Data and Artificial Intelligence (AI) encompasses services, components, and operations necessary to marshal Azure cloud resources for data analytics and predictive learning purposes. When the data of interest entails Protected Health Information (PHI), the Blueprint can implement security control measures to preserve the confidentiality and integrity of data during transmission, processing, and storage. Further, the Blueprint can support a customer's regulatory requirements concerning PHI; for example, its usage can be complementary with a customer's existing HITRUST Common Security Framework (CSF) certification. Microsoft Azure is a HITRUST-certified cloud service provider (CSP), with a diverse catalog of certified service offerings. The Blueprint is founded on a core set of certified services, and thus engineered with the HITRUST CSF in mind.

The Blueprint can be used to deploy private analytics solutions specific to a customer's internal needs, or public solutions serving a particular customer constituency. It is intended to be a flexible architecture adaptable to a wide array of analytical use cases.

This whitepaper constitutes a review of the Blueprint architecture and functionality with respect to HITRUST-certified customer environments, examining how specifically it can satisfy HITRUST CSF security requirements. The whitepaper also considers the interplay of Azure- and customer-oriented responsibilities for Blueprint deployment, configuration, and management in a manner consistent with the HITRUST CSF. Last, the whitepaper presents an illustrative deployment use case to help the reader visualize the Blueprint architecture in action.

Coalfire Systems would like to thank Microsoft Azure for the opportunity to review the Blueprint in advance of production release, and to compose this whitepaper on its behalf.

INTRODUCTION

Without continual growth and progress, such words as improvement, achievement, and success have no meaning. - Benjamin Franklin

We live in an information security age, wherein data is a precious commodity underlying social and business exchange, and requires constant protection from loss, theft, and exploitation. The consequences of inadequate protection can be expensive and difficult to bear. It is essential for organizations to understand information security measures, and implement them with respect to technology, operations, and management. Industry and government entities publish and mandate a dizzying complex of security requirements, in the interest of promoting security compliance within organizations; they are challenged to update requirements as technologies and business models continue to evolve. In turn, organizations are challenged to stay abreast of the compliance landscape, tactically implement security within budgetary limits, and confidently assure customers that the confidentiality and integrity of their data is preserved at all times. This applies to virtually every industry— healthcare is no exception. In fact, the healthcare industry especially embodies the complexity and ineluctability of security compliance, because of the intricacies of privacy regulation (at the local, state, and federal levels), and the high exploitation value of Protected Health Information (PHI).

Out of this state of affairs arose the creation of the Health Information Trust (HITRUST) Alliance. Formed as a consortium of leading healthcare and information security organizations, HITRUST developed the Common Security Framework (CSF) as a certifiable controls-based risk management methodology, grounded in ongoing due diligence toward healthcare information security and privacy. The HITRUST CSF offers a harmonized approach to security compliance and risk management, grounded in the Security and Privacy Rules of the Health Insurance Portability and Accountability Act (HIPAA) of 1996, but extending well beyond these to incorporate requirements from a diverse range of frameworks (e.g., Payment Card Industry Data Security Standard (PCI DSS), NIST Cybersecurity Framework, and International Standards Organization (ISO) 27001:2013), allowing organizations to leverage the HITRUST CSF across their entire compliance landscape[1]. The HITRUST CSF is updated regularly (generally, twice annually) by the Alliance, to maintain its relevance to new regulations, technologies, and industry practices. Presently, HITRUST CSF Version 9.0 is publicly available for newly certifying organizations.

The HITRUST CSF comprises a catalog of one hundred thirty-five (135) security and privacy *controls*, each an extensive set of related requirement statements. For a given certifying organization, a subset of requirement statements will be drawn from the control catalog, and organized into an *assessment object* encompassing nineteen (19) information security *domains*. Where a control represents a tightly knit collection of requirements (e.g., control 10.f concerns cryptography requirements), a domain represents cross-related requirements from an assortment of controls (e.g., the Network Security domain includes requirements from the cryptography, service provider security, network segregation, security monitoring, and asset management controls). The exact composition of requirements within an assessment object depend on the organization's specific *scoping factors*, which are organizational, technological, geographical, and regulatory characteristics.[2] The number and rigor of requirements within the assessment

[1] *Cf. HITRUST CSF Standards and Regulations Cross-Reference* and *An Introduction to the HITRUST Common Security Framework (CSF).*

[2] *Organizational Factors* concern the number of PHI records retained by the organization. *Geographical factors* consider whether the organization operates in a given US state, multiple states, or offshore. *System factors* consider accessibility to the in-scope environment (by third parties, via the Internet, and/or from public locations); use of mobile devices in the environment; and the number system users, interfaces, and daily PHI transactions. *Regulatory factors* concern applicability to particular frameworks covered by the

object scale with the number and level of scoping factors. As such, the object is tailored to the specific organization

Requirements are scored in terms of a five-tier maturity system (derived from the National Institute of Standards and Technology (NIST) Program Review for Information Security Management Assistance (PRISMA) model). The maturity tiers are the following: Policy, Process, Implemented, Measured, and Managed.[3] For each of these, a requirement may rate as one of five graduated scores, ranging from not mature (i.e., Not Compliant) to fully mature (i.e., Fully Compliant).[4] This approach allows compliance to be assessed along a continuum, whereby the organization can finely gauge how well it addresses each requirement. It also fully accounts for both the *documented* and *implemented* security posture of the organization.

The HITRUST CSF is a repeatable methodology. An organization under HITRUST must recertify annually, alternating between a full *validated assessment* (entailing all applicable requirements) and a milder *interim assessment* (i.e., a pulse-check validation of a small subset of requirements). The organization works with a HITRUST-certified assessor firm to review the organization's documentation and implementation evidence, draft corresponding findings, and submit both to HITRUST for certification review. HITRUST maintains rigorous certification standards not only for assessed organizations, but also for assessor firms. Additional information regarding HITRUST can be found at https://hitrustalliance.net/.

Microsoft Azure has retained a HITRUST CSF certification since 2015, covering a diverse collection of its cloud service offerings. Azure recognizes the importance of the certification to its healthcare customer base, given its status as a Business Associate (BA). Thus, Azure has an assessment object, undertakes a full assessment every year (opting out of the interim assessment), and offers certified services under its Business Associate Agreement (BAA) with healthcare customers. The BAA[5] specifies the eligible services, provisioning and usage obligations, data protection, and breach notification terms between service provider and customer. At this time, Azure services are certified under Version 8.1. Azure is actively pursuing certification under Version 9.0. This whitepaper considers Azure's six hundred seventy-two (672) requirements under Version 9.0.

The Microsoft Azure Security and Compliance – HITRUST Health Data and Artificial Intelligence (AI) Blueprint is architected from a core set of certified services. As such, when properly configured and deployed, the Blueprint can provide an analytics platform consistent with the HITRUST CSF. The Blueprint was developed from Day One with PHI protection and HITRUST compliance in mind; however, *how* the Blueprint is used (with respect to given organizational premises and assets, within the context of an organization's security governance program) is a large area of responsibility that the customer bears. Azure's HITRUST certification does not relieve the customer of its own duty to achieve (and maintain) certification-- it is not a rubber stamp. But within a model of *shared* responsibility, accounting for service *provision* and *consumption* as complementary operations, Azure and customer together can contribute to the deployment of a Blueprint with full compliance assurance. We will delve into the Blueprint architecture, and the shared responsibility model, later in the whitepaper.

HITRUST CSF, such PCI DSS, Federal Trade Commission (FTC) Red Flags, and Texas Health and Safety Codes.

[3] *Policy* considers the documentation of security requirements, while *Process* considers the documentation of how requirements are met by the organization. *Implemented* considers the actual security mechanisms and methods in action. *Measured* examines how security is evaluated for effectiveness, while *Managed* examines how effectively the organization's management addresses security deficiency.

[4] Cf. *Risk Analysis Guide for HITRUST Organizations and Assessors*

[5] The BAA is available at http://aka.ms/BAA

The commonly acknowledged benefits of cloud service usage include reduced capital and operational expenditure (Capex and Opex) by organizations, metered service costs, on-demand resource scalability, and access to advanced technology without the burden of ownership. These certainly apply to the Blueprint. But one more benefit can be cited— reliance on a CSP's compliance posture, potentially lowering (but not eliminating) an organization's level of attestation effort.

THE BLUEPRINT ARCHITECTURE

OVERVIEW

The Blueprint is intended to be deployed within a customer's virtual enclave within Azure, accessible from the customer's premises. It is bound to a customer's Azure subscription. Fundamentally, it is a Platform as a Service (PaaS) offering that combines a core set of Azure services into a generalized data analytics architecture. As such, the services represent the minimum requisite set of functions for effective data analysis. However, the Blueprint also functions as a Software as a Service (SaaS) offering, in that its architectural framework can be customized to address specific end-user analytical use cases. From the end-user perspective, it is a software solution that requires little to no development, engineering, administration, or technical support responsibility. Put another way, the Blueprint is a *solutions platform*, whereby individual deployments are case-specific specializations of the general architecture.

The Blueprint is also intended to obviate *deidentification* of data as a condition of analysis. Deidentification is the process of obscuration or anonymization of identifying elements in a data set, such that association of the resultant set with specific individuals is statistically highly improbable. The Department of Health and Human Services (DHHS) Office of Civil Rights (OCR) has asserted that when properly deidentified (demonstrable through statistical analysis), PHI is exempt from HIPAA Security and Privacy Rule requirements (equivalently, it ceases to fall under the rubric of *protected* health information). However, achieving statistically validated deidentification can be a challenging and expensive undertaking; also, it can severely limit the varieties of analysis one may wish to perform. The Blueprint provides a sample architecture and design architecture can safely, securely process data without prior deidentification, because the architecture is designed with HITRUST CSF-certified services that can be configured to preserve PHI confidentiality and integrity.

The generalized Blueprint architecture is depicted in Figure 1. Specific Azure services are arrayed in six (6) functional tiers that govern the importation, processing, analysis, storage, and security management of data within a deployment.

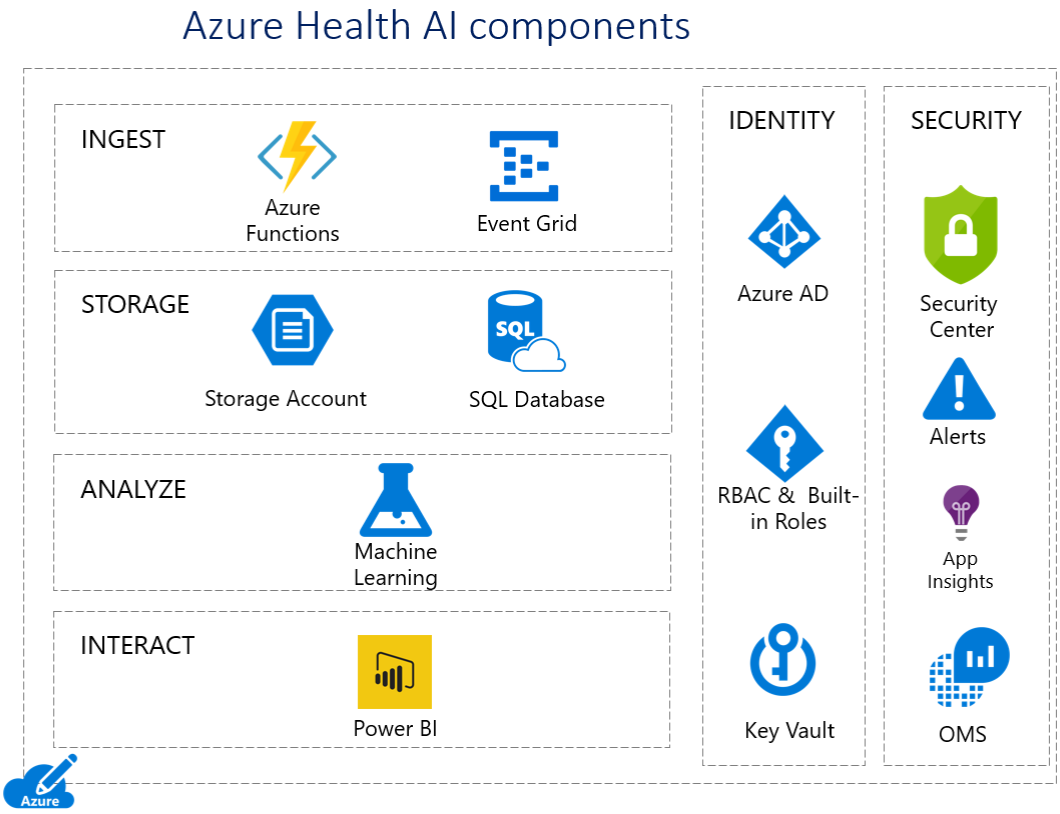

Figure 1: Blueprint Core Achitecture

The Ingest tier is designed to illustrate initial bulk importation of a sample set of fictitious raw PHI data into the deployment (establishing a basis for analysis), and transactional data ingestion over time (evolving the analysis basis). This is achieved via the Azure Functions and Event Grid services, using an event-trigger paradigm. Event Grid pairs specific data sources (i.e., event *publishers*) with Azure Functions as an event *handler.* Whenever a data event is cued within Event Grid, an sample Azure Function triggers an authentication token request from Azure Active Directory. The Azure Function then passes the received token to Azure KeyVault, to request a key for establishing a Transport Layer Security (TLS)-based connection with a SQL Database instance. KeyVault validates the request against its access policies; if successful, the key request is authorized, and a key meeting defined criteria is returned. The sample Azure Function can subsequently initiate a database connection to import the data. A Separate sample Azure Functions are used for bulk data importation and transactional data intake. Effectively, Azure Functions orchestrate data flow within the Blueprint, with Event Grid serving as a subscription relationship between data source and destination. Figure 2 depicts the general process.

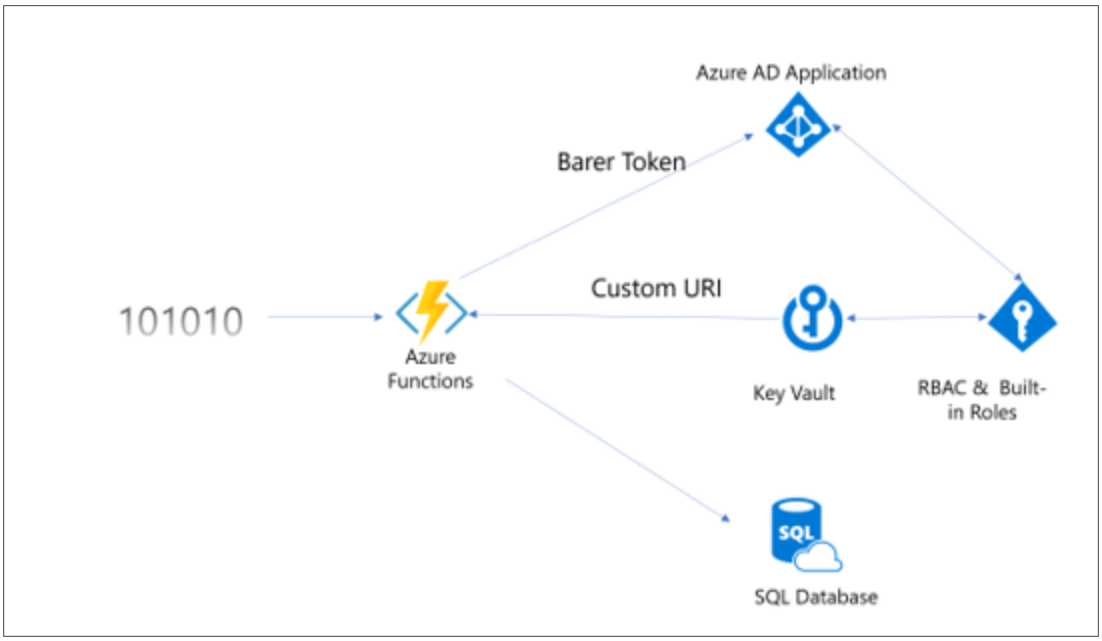

Figure 2: Blueprint Data Ingestion

The Storage tier addresses the management of ingested data within the Blueprint. A Structured Query Language (SQL) Database instance serves as a relational database repository, wherein data is normalized in accordance with a particular schema. The customer can implement industry-standard schemas for healthcare information, such as Fast Healthcare Interoperability Resources (FHIR) for patient data management, Digital Imaging and Communications in Medicine (DICOM) for medical image management, and International Classification of Diseases (ICD) code sets for medical procedure and diagnosis identification. Schemas support the efficient execution of SQL queries to process data for Machine Learning use. The sample SQL Database instance is integrated with Azure Active Directory for access control, and can implement Transparent Data Encryption (TDE) to cryptographically secure data within database structures (e.g., tables, views, and stored procedures). Underlying the SQL Database instance is an Azure Storage instance that can support file, block, or Binary Large Object (BLOB) data types. The instance also can encrypt data outside the database, via the Storage Service Encryption (SSE) function. An Azure Storage container can also be configured as an event publisher, to trigger ingestion into the database.

A sample Analysis tier comprises a sample Azure Machine Learning service, an artificial intelligence engine that can analyze and "learn" statistical patterns from data, in order to make behavioral predictions and decisions without human intervention. Machine Learning employs an extensive library of analytical algorithms (e.g., Linear Regression, Decision Tree, Random Forest, Naïve Bayes, and K Nearest Neighbors) and the R statistical language to create one or more *training models*, which in turn are refined via experiments with data derived from the SQL Database. The ultimate goal is a continually fine-tuned model that can accurately determine an outcome with high precision, classify an undifferentiated data population, or execute a decision based on learned knowledge. Azure Machine Learning can glean insights that would otherwise elude manual human analysis.

The Interact tier processes Machine Learning output for human review, using the Power Business Intelligence (PowerBI) service to visualize output in the form of highly customizable graphs, charts, reports, and dashboards. Visualizations are updated in real-time, and allow for incisive slice-and-dice functions to view predictions or decisions across different dimensions. Visualizations can also be embedded into

external web services and applications for extensibility. PowerBI is essentially the end-user analysis interface of the Blueprint.

The Identity tier addresses access control, authentication, and authorization of Blueprint users and interoperating components. Its foundation is the Azure Active Directory (AAD) service, a cloud analogue of the familiar Windows Active Directory domain infrastructure (based on the Lightweight Directory Access Protocol (LDAP) and Kerberos authentication). All Blueprint components are AAD-integrated, allowing for seamless authentication, identity monitoring, group policy enforcement, and a uniform role-based access control (RBAC) model throughout the deployment. Closely allied to AAD is the Azure KeyVault service for key management. KeyVault is the repository for cryptographic secrets used within the deployment for encryption operations (e.g., TDE, SSE, TLS, AAD password hashes), and governs key generation, dissemination, rotation, revocation, archivation, and escrow. Dissemination requires request validation against AAD to vet the requestor identity, and ensure it has authorization for key receipt and usage.

Last, the Security tier allows for real-time situational awareness of Blueprint operations. At the tier's heart is the Azure Security Center, a service that functions as a configurable security information and event management (SIEM) platform. Security Center employs Security Content Automation Protocol (SCAP) and Open Vulnerability Assessment Language (OVAL) based checklists and taxonomies to automatically identify, assess, and report on vulnerabilities; also, to provide actionable remediation recommendations. The Azure Applications Insights service functions as a configurable performance baseline analyzer, monitoring service telemetry to detect patterns in user and data interactions, resource consumption, and operational behaviors. Application Insights can provide real-time notification, root cause determination, operational impact measurement, and actionable recommendations for identified performance issues. The Azure Operations Management Suite (OMS) service automates operational security functions within the Blueprint deployment, including mandatory configuration monitoring and enforcement, antimalware protection administration, component patch management, and deployment availability management. Like Security Center and Application Insights, OMS also analyzes audit log and telemetry data to monitor the operational health of the deployment. The Azure Alerts service provides a unified notification layer for all Security tier services.

The overall Blueprint dataflow can be visualized in Figure 3 below:

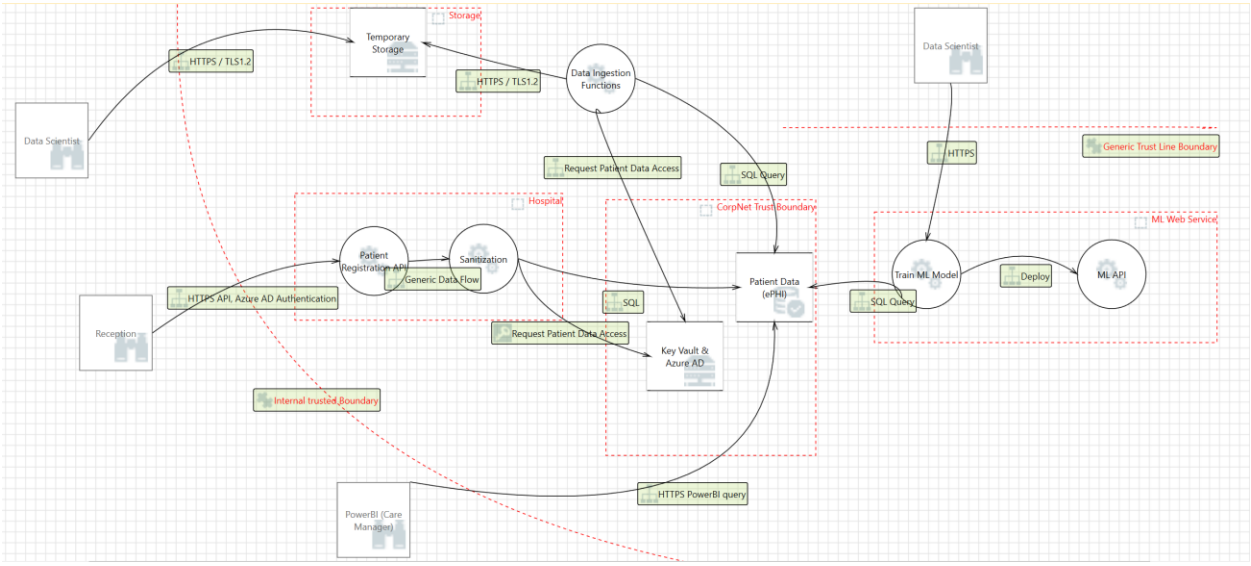

Figure 3: High-Level Blueprint Data Flow

In addition to these tiers, the Blueprint architecture includes resources:

- **Azure Resource Manager (ARM) Templates**: JavaScript Object Notation (JSON) files to marshal the instantiation of Blueprint components in a controlled, repeatable, and portable fashion. The templates are defined with standard Blueprint configuration metadata, but can be adapted to support instantiation of new objects and configurations.
- **Automated PowerShell and Global Administrator Scripts**: these orchestrate ARM templates to facilitate correct Blueprint deployment.
- **Microsoft Threat Model**: a comprehensive, Blueprint-specific threat model (consumable as a .tm7 file by the Microsoft Threat Model Tool) to identify, classify, and rank security risks across Blueprint components, data flows, and trust boundaries. Particularly, the model can help customers assess risks when modifying the architecture for specific use cases.
- **Azure Customer Responsibility Matrix (CRM)**: a spreadsheet-style document to depict Azure HITRUST CSF Version 9.0 requirements and corresponding security responsibilities for Azure and the customer.
- **Microsoft Azure Security and Compliance – HITRUST Health Data & Artificial Intelligence (AI) Blueprint**: the present whitepaper, serving as a security compliance guide for deploying and using the Blueprint architecture.

THE SHARED RESPONSIBILITY MODEL

OVERVIEW

As a rule, security does not happen automatically by itself. It must be engineered into an information system, applied thoughtfully across the system's technical, operational, and management dimensions; also, it must account for interactions with users and external parties (e.g., service providers and downstream Business Associates). Microsoft Azure incorporates this principle into its cloud service offerings, including the Blueprint, and provides corresponding guidance to its customers.

In Azure, security responsibility is a continuum between the CSP and customer. At one end of the continuum, Azure is responsible for security ***of*** the cloud; namely, the secure configuration and deployment of infrastructure used to provision services. This includes physical hardware and core software underlying compute, storage, database, and networking services (e.g., bare-metal hypervisors, software-defined network (SDN) components, and storage media); also, the physical facilities supporting service availability.

At the other end, the customer is responsible for security ***in*** the cloud; that is, security of logical components residing atop the service provision infrastructure. This includes the configuration and deployment of virtual resource instances (e.g., virtual machines, applications, and storage containers), customer data, and user access (e.g., identity management, authentication and authorization, and virtual network methods). Also included are customer-owned assets used to consume services (e.g., on-premise endpoints, web browser software, and local network infrastructure facilitating customer access to the Azure Management Portal). In this, customer responsibilities are similar to a traditional infrastructure datacenter / platform hosting scenario (cf. Figure 2).

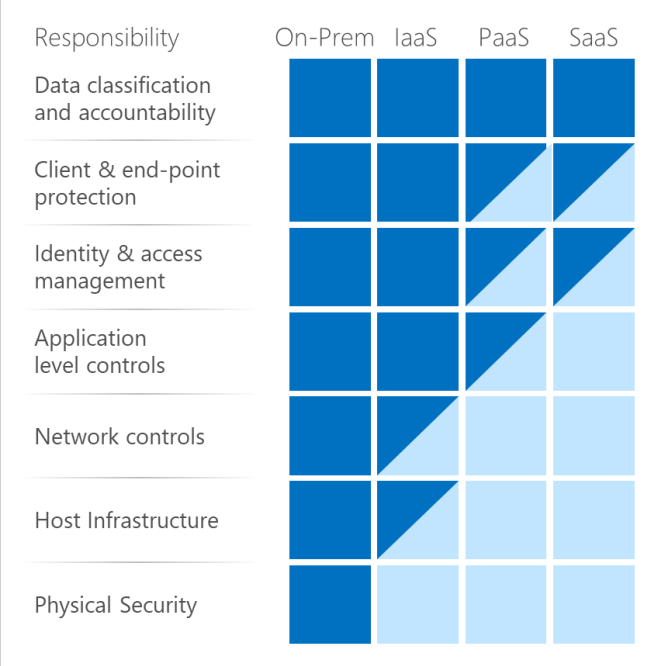

Figure 4: Azure Shared Responsibility Model

An equivalent way of viewing the security responsibility continuum is to consider that cloud service usage embraces both service *provision* (by the Azure) and service *consumption* (by the customer). Both provision and consumption must be performed securely to ensure the security of service usage. They are

complementary sides of the cloud security coin. Thus, Azure and customer both own degrees of security responsibility for cloud service usage. In the context of the Blueprint and the HITRUST CSF, this generally translates into shared responsibilities across domain requirements, to account for Blueprint provisioning and consumption. It also means that when using the Blueprint, customers leverage Azure's secure provisioning infrastructure, and can cite Azure's HITRUST CSF certification in their own assessment activities.[6] This can effectively save the customer the effort of Azure testing and validation, but it does not obviate the customer from testing and validating their secure consumption of the Blueprint. The complete picture of HITRUST-compliant Blueprint usage involves Azure's certification *plus* the customer's own.

It should be noted that Azure's HITRUST CSF certification is based on its scoping factors, which may differ from a customer's own. Consequently, the customer's assessment object may have a different mix of applicable requirements. It is possible that some customer requirements will not exist in Azure's assessment object. In such cases, the customer assumes full responsibility for satisfying the requirement.

When it comes to its share of security responsibility for a Blueprint deployment, the customer can consider three possible options:

- Security implementation hosted on an Azure VM (e.g., use of a third-party software solution for managing cryptographic keys)
- On-premise security implementation (e.g., use of a cryptographic module or library on-premises)
- Leverage an Azure service offering (e.g., the Blueprint already incorporates the KeyVault service)

The choice is ultimately the customer's own to make.

RESPONSIBILITIES BY HITRUST CSF DOMAIN

To further clarify the nature of shared security responsibility, the following sections summarize Azure and customer obligations at the domain level.

Information Protection Program: The customer is responsible for the implementation, configuration, management, and monitoring of an information protection program covering its information technology assets and uses. The program should be based on an industry-accepted risk management and security compliance framework (like the HITRUST CSF), with explicit executive management support, budgetary backing, and designated senior-level personnel responsible for program management, execution, review, and development. The program should entail annual independent security assessment, remediation, and reporting activities, as well as processes for communicating and enforcing security requirements to users. Azure, per its HITRUST CSF certification, meets these requirements for its service provision environment and infrastructure, via its own program (patterned after the ISO 27000:2013 Information Security Management System (ISMS) methodology).

Endpoint Protection: The customer is responsible for security configuration of its virtual and on-premise endpoints that interact with the Blueprint. This includes implementation of antimalware, antispam, software patching, application black/whitelisting, host-based firewall, and file integrity monitoring mechanisms. Azure satisfies these requirements with respect its service provision environment; e.g., through the installation of

[6] In HITRUST CSF terms, this is called *control reliance*. The customer can only rely upon requirements that exist in both a provider's assessment object and its own. For a given requirement during assessment, the provider's findings and scores would be incorporated into the customer's own, factoring in the degree and circumstances under which requirement satisfaction is achieved by both parties. Note that while control reliance can reduce attestation level of effort, it can also expose the customer to deficiencies in the provider's portion of responsibility.

the Azure Security Pack suite of security tools on infrastructure and platform components, and provision of secure baseline Virtual Machine Images (VMIs) to customers for VM instantiation. Note that customer VM instances are not part of the Blueprint architecture; from an endpoint security perspective, they are no different in customer responsibility from on-premise endpoints (cf. “Deployment Considerations for the Blueprint” section of this whitepaper).

Portable Media Security: The customer is responsible for the security configuration and usage of portable media (e.g., thumb drives, external hard drives) it permits to interact with the Blueprint. This includes implementation of appropriate measures for portable media access control and usage restriction, encryption, transport and storage, sanitization and disposal; also, implementation of a data loss prevention (DLP) solutions to safeguard against data exfiltration. Portable media usage is not permitted within Azure datacenter facilities, where the Blueprint can be hosted.

Mobile Device Security: The customer is responsible for the security configuration and usage of mobile devices (e.g., smartphones and laptops) it permits to interact with the Blueprint. This includes implementation of appropriate measures for access control and usage restriction, device endpoint security, device physical security, and employment of a mobile device management (MDM) and data loss prevention (DLP) solutions to centrally manage and monitor devices. The customer is also responsible for telework security, whereby mobile devices may be used outside of organizationally controlled spaces; and for Bring Your Own Device (BYOD) policy, whereby non-organizational (i.e., personal) devices may be used. Mobile device usage, telework, and BYOD are not permitted within Azure datacenter facilities, where the Blueprint can be hosted.

Wireless Security: The customer is responsible for the security configuration and usage of wireless access methods with the Blueprint. This includes implementation of appropriate measures for access control and usage restriction, secure authentication, transmission encryption, wireless access point (WAP) monitoring and scanning. Wireless access is not permitted within Azure datacenter facilities, where the Blueprint can be hosted.

Configuration Management: The customer is responsible for the security configuration and management of the Blueprint, whereby the customer oversees specific Blueprint configuration settings during deployment. The OMS service can help facilitate enforcement of a mandatory baseline configuration for the deployment. The customer is also responsible for the security configuration of owned assets that interact with the Blueprint. The customer should implement a configuration management plan and process for defining baseline configurations and mandatory settings, and the controlled execution of configuration changes (involving change review, security testing, adjudication, approval, and deployment). Azure, per its HITRUST CSF certification, meets these requirements for its service provision environment and infrastructure, via the Microsoft Software Development Lifecycle (SDL) methodology.

Vulnerability Management: The customer is responsible for the institution and implementation of a technical vulnerability management program that addresses the identification, classification, risk rating, monitoring, and remediation of vulnerabilities that may affect its Blueprint deployment (including on-premise and hosted customer assets used to interact with the deployment). The vulnerability management program should be integrated into the risk management and information protection programs, and include vulnerability scanning and periodic penetration testing. Azure meets HITRUST CSF requirements for vulnerability management of its service provision infrastructure via its program (based on ISO 27000:2013 and NIST SP 800-137).

Network Protection: The customer is responsible for security configuration, monitoring, and management of network access to its Blueprint deployment. While situated within Azure datacenters, user access to a deployment via the web-based Azure Management Portal necessarily entails the customer’s local network

infrastructure (i.e., switches, routers, firewalls, load balancers, addressing schemes, and network services such Domain Name Service (DNS) and Dynamic Host Configuration Protocol (DHCP)); thus, security of this infrastructure is solely the customer's responsibility. Processes for network protection should be integrated into audit monitoring, access control, authentication, transmission protection, incident management, vulnerability management, and configuration management processes; indeed, the HITRUST CSF requires the implementation of *continuous monitoring* methodology that involves security process integration, periodic measurement, and ongoing remediation (as a complement to formal annual assessment activities). Within a Blueprint deployment, network protection is largely accomplished via integration of KeyVault, Azure Active Directory, and Security Center services; however, additional Azure network services can be integrated into the deployment (cf. "Deployment Considerations for the Blueprint" section of this whitepaper). Azure satisfies HITRUST CSF requirements for network protection of its service provision infrastructure via control mechanisms implemented at the core physical network and SDN tiers.

Transmission Protection: The customer is responsible for the establishment and execution of a cryptographic key management process with respect to its Blueprint deployment. The Azure KeyVault service can be employed to manage cryptographic keys consumed within the deployment, but the customer must define the associated key management criteria (e.g., key length and strength, lifespan, rotation schedule, dissemination, revocation, archivation, and escrow). Note that HITRUST CSF also requires the identification, monitoring, and management of remote access to the deployment by third parties. Key management practices should support the use of transmission encryption methods such as Hypertext Transfer Protocol Secure (HTTPS), TLS, and Internet Protocol Secure (IPSec). Azure also employs KeyVault (alongside internal, purpose-built cryptographic key stores) for key management pertinent to its service provision infrastructure.

Password Management: The customer is responsible for authenticator and identity definition, management, and protection with respect to a Blueprint deployment. This includes authenticator complexity, lifespan, refreshment, and reuse requirements. The customer can facilitate this largely through the Azure Active Directory service. Note that multifactor authentication (MFA) is not inherently part of the Blueprint, but can be integrated into the deployment (cf. "Deployment Considerations for the Blueprint" section of this whitepaper). Azure satisfies HITRUST CSF requirements for authentication management for its service provision infrastructure, via MFA involving an Active Directory-managed username/password credential pair and a KeyVault-managed cryptographic keycard.

Access Control: The customer is responsible for account management, access control, and authorization of its users of a Blueprint deployment. This includes management and monitoring of account authorization, creation, modification, disablement, and termination events; also, session management and monitoring (e.g., controlling session duration, lockout, termination, and reestablishment). The customer can facilitate this largely through the Azure Active Directory service. Note that the HITRUST CSF requires periodic review and validation of accounts for continued need and authorization levels, user acknowledgement of Acceptable User Policy (AUP) or Rules of Behavior (RoB) requirements governing interaction with organizational assets (including the customer's Blueprint deployment), and specific procedures for physical and logical access termination when personnel leave the organization (e.g., retrieval of access badges, authentication hard tokens, and equipment; disablement of accounts within twenty-four (24) hours of personnel departure). Azure satisfies HITRUST CSF requirements for access control of its service provision infrastructure, via processes and mechanisms for tightly restricting and monitoring access by service development, engineering, and administration personnel.

Audit, Logging, and Monitoring: The customer is responsible for the security configuration and management of audit monitoring of its Blueprint deployment. This include the enablement of audit logging features (like the Azure Security Center), definition of auditable events and log content, log data filtration and report

generation, control of access to log data, and audit log review and analysis. The audit monitoring process should dovetail with the incident management program, as security incidents typically will be identified through log review and analysis. It should also dovetail with the security awareness and training program, as users should be apprised of monitoring of their activities within the deployment. Azure satisfies HITRUST CSF requirements for audit monitoring of the service provision infrastructure, largely facilitated through the use of the proprietary Monitoring and Diagnostics Service (MDS) and Security Logging Auditing and Monitoring (SLAM) enterprise tools.

Education, Training, & Awareness: The customer is responsible for the institution and maintenance of a security awareness and training program for its organizational members, including those who interact with a Blueprint deployment. The program should include measures for initial training provision (during personnel onboarding), annual refresher training, recordation of training completion, and monitoring to ensure training completion. Note that HITRUST requires initial training to occur no later than sixty (60) days after personnel hiring, and training records to be retained for a minimum of five (5) years. Training should address both general information security topics, role-specific subject matter (e.g., role-based security training for network administrators, incident responders, and data analysts), and healthcare security/privacy requirements. The customers is also responsible for personnel acknowledgement of its AUP/RoB requirements; this can be enfolded into security awareness training material, and noted through recordation of training completion. Azure implements a security awareness and training program for its personnel, including those who support the provisioning infrastructure of the Blueprint.

Third Party Assurance: The customer is responsible for management and monitoring of third-party service providers that may interact with its Blueprint deployment. This includes performance of security due diligence and evaluation on providers (prior to service performance or delivery), and contractual formulation of terms addressing service definitions, delivery, performance levels, changes, information security, compliance monitoring, termination, penalty and remedy conditions. The customer is expected to monitor providers against these terms on an annual basis. Azure bears the same responsibilities for its service provisioning infrastructure, holding its providers (principally, datacenter vendors) to the terms of the Microsoft Master Services Agreement. It worth noting that Azure does not outsource service development or management activities; these are performed in-house by Azure personnel.

Incident Management: The customer is responsible for the establishment and exercise of a formal incident management program that accounts for the identification, classification, risk rating, reporting, monitoring, handling, and resolution of incidents that threaten the confidentiality, integrity, or availability of its data (including data within a Blueprint deployment). The incident management program includes coverage of PHI breach handling and notification procedures, per HIPAA and HITECH (e.g., breach notification within sixty (60) days after discovery; maintenance of a log of unauthorized disclosures for DHHS submission). As explained below (cf. “Data Protection and Privacy” paragraph), Azure has limited capacity to provide breach reporting to customers. Azure datacenters satisfy these requirements for its service provision infrastructure through georeplication, redundant telecommunications, and environmental security controls (e.g., fire suppression, leak detection, and temperature/humidity monitoring).

Business Continuity & Disaster Recovery: The customer is responsible for instituting a business continuity process and contingency plan that accounts for the recovery, reconstitution, and recovery of its Blueprint deployment. This includes determination of required capacity, identification of critical business functions, definition of recovery point/time objectives, identification of key roles and responsibilities, and annual plan training and testing activities. Blueprint deployments can be georeplicated across multiple Azure datacenters, but note that data backup services are not part of the Blueprint architecture. Azure datacenters satisfy HITRUST CSF requirements for business continuity and disaster recovery through georeplication, redundant telecommunications, and environmental security controls.

Risk Management: The customer is responsible for the establishment of a risk management program that encompasses risk identification, classification, risk rating, and remediation activities in a controlled manner (employing a formal methodology) at planned intervals, and whenever significant organizational, technological, architectural, infrastructural, or security changes occur. The program ties into the customer's configuration management plan, as changes should be analyzed for potential risk and security impact. Risk management should cover the customer's Blueprint deployment usage, as well as security controls pertinent to customer assets that access and interact with the deployment. The Blueprint core services satisfy HITRUST CSF requirements for risk management, via an enterprise program based on the ISO 31000 methodology.

Physical & Environmental Security: The Blueprint is hosted physically with Azure datacenter facilities, which satisfy HITRUST CSF requirements for physical access control, surveillance, physical maintenance, and environmental security. As a general rule, use of the Blueprint should not require the customer to interact directly with datacenter facilities or personnel. However, the customer is responsible for physical and environmental security of its own facilities, from which its Blueprint deployment may be accessed and used. For example, physical security and maintenance of customer endpoints interacting with the deployment is solely a customer responsibility.

Data Protection & Privacy: The customer should observe that when using the Blueprint, its data is hosted within the Azure service environment; however, Azure explicitly and strictly has no visibility into its specific nature or location. Azure storage resources employ data obscuration and randomization techniques to safeguard data from brute-force exfiltration and facilitate high data availability; however, these render the data accessible only to respective customers (and their authorized users). Thus, regarding PHI privacy, Azure does not assume responsibility for (1) reporting on a customer's authorized disclosures, (2) monitoring the specific locations of customer PHI, (3) monitoring specific user access to customer PHI, (4) establishing and abiding specific retention schedules for PHI, or (5) specifically encrypting PHI on behalf of the customer. In these cases, the customer assumes full responsibility. Azure does assume responsibility for breach reporting (per HIPAA §164.410); however, it can provide only limited information on affected identities and data types. All hosted customer data remains within Azure permanently unless (1) deleted by the customer, or (2) the customer terminates its Azure subscription (in which case, the data will remain available for sixty (60) days after termination, to allow graceful data removal by the customer). These points are communicated in the Microsoft BAA.

DEPLOYMENT CONSIDERATIONS FOR THE BLUEPRINT

Given the aforementioned shared responsibility model, one can appreciate that customer responsibilities for security come into play when deploying the Blueprint. The following sections provide deployment and configuration guidance to help the customer implement and use the Blueprint in a manner consistent with the HITRUST CSF.

NETWORK SECURITY

The HITRUST CSF requires end-to-end encrypted transmission of PHI. At no point along the transmission pathway should a malicious interloper be able to intercept, redirect, or tamper with transmitted PHI. One should note that while the Blueprint encrypts internal transmissions (e.g., between Azure Functions and Azure Storage), it does not inherently address transmissions *to* the deployment (e.g., from a customer endpoint to deployment boundary). It is intended that the customer provisions transmission encryption for this portion of the network pathway. A virtual private network (VPN) between the deployment and the customer premises can provide a dedicated tunneled connection to securely transmit PHI through the Internet, using TLS or IPSec encryption methods.

The customer can implement its own VPN, or may consider using the Azure VPN Gateway and Virtual Network services (covered under the Microsoft BAA, but not included in the Blueprint) in tandem with its deployment. In the latter case, the deployment would be bounded within a Virtual Network instance, which would constitute a termination perimeter for VPN Gateway connections with the customer's network. In both cases, the customer should ensure that the encryption method employs (at minimum) the AES 256 cryptographic algorithm. This is in keeping with FIPS and OCR guidance for protecting PHI.[7] Older algorithms, particularly the Triple Data Encryption Standard (3DES), are no longer considered adequate for PHI. Per NIST guidance, the customer should also consider avoiding the use of Secure Socket Layer (SSL) based VPN solutions; SSL has been subject to a crop of pervasive exploits (e.g., POODLE, BEAST, and Heartbleed) in recent years, and is generally considered less secure than TLS (especially, TLS 1.2 and later).

The Blueprint also does not address management of network traffic leading to the deployment boundary. The solution focuses on using identity-based segmentation (authenticating all connection endpoints), as the solution is PaaS based it has forgone NSG's or V-net network segmentation for identity isolation using Authenticated and authorized connections, with log trails for all connections. As of this writing Microsoft has introduced SQL Virtual Network service endpoint segmentation, but this Vnet segmentation was not used in the Blueprint.

This is where the customer's on-premise firewall and load balancer topologies would come into play, wherein the customer has configured specific rulesets, access control lists (ACLs), and policies to enable only the minimum necessary set of ports and protocols to connect to Azure. Irrespective of method, an important consideration is how segmenting is applied in the Blueprint deployment from non-PHI environments, such that deployment PHI does not find its way to untrusted environments and unauthorized users.

ACCESS CONTROL AND AUTHENTICATION

[7] The OCR has asserted that unauthorized disclosure of strongly encrypted PHI does not constitute a reportable breach event. While important for transmitted PHI, this consideration is critical for at-rest PHI stored on mobile devices that can be easily stolen.

The Blueprint features powerful, adaptable access control methods throughout all deployment tiers. Azure Active Directory is the access control heart of the deployment. While the Blueprint establishes a default set accounts and roles during instantiation, the customer should ensure that all intended identities and their accounts observe the principles of *least privilege* (i.e., each has only the minimal set of permissions and rights necessary to perform authorized actions), *separation of duties* (i.e., identities only minimally overlap in terms of permissions and rights), and *role membership* (i.e., permissions and rights are attach to roles, and assigned to identities through membership and inheritance in one or more roles). The customer also should ensure that identities observe established account management practices (e.g., account disablement after ninety (90) days of inactivity; session termination after thirty (30) minutes of inactivity; separate accounts for privileged and non-privileged functions). All of this can be accomplished via AAD.

Regarding authentication, AAD can manage password management policy and enforce appropriate constraints (e.g., uniqueness, complexity, duration, reuse, and refreshment). AAD stores passwords as Secure Hash Authentication (SHA) 2 encrypted values, and employs Kerberos-based challenge handshake validation of identities and services. The Blueprint does not deploy MFA capability by default, but it can be enabled during deployment; in which case, MFA interoperates with the KeyVault service to derive cryptographic authenticators (as keys or certificates) to complement AAD username/password credential pairs.

DATABASE SECURITY

The Azure SQL Database service within the Blueprint implements TDE to cryptographically secure data (in either the database schema or the individual table column). TDE supports usage of the AES 256 and 3DES algorithms; as aforementioned, AES 256 is strongly preferred by FIPS and the OCR. TDE generates a database encryption key (DEK) for cryptographic functions. The customer can elect to store the DEK within the SQL Database instance, but should consider storage within KeyVault instead. Note the DEK will be unique to each SQL Database instance (though shared amongst databases within a given instance), and therefore also unique to each Blueprint deployment. If the DEK is lost or destroyed, the instance data is effectively *crypto-shredded*; that is, it has been cryptographically rendered unrecoverable and unusable. DEK storage within KeyVault ensures the key can be escrowed for recovery purposes.[8]

The SQL Database instance also implements dynamic data masking, to selectively obfuscate PHI data elements on the fly to users. Masking can visually replace elements with dummy default, random number, string mask, or null values while leaving the underlying PHI in its original state. The customer may wish to use dynamic data masking for certain kinds of end-user reporting, delimited to certain AAD identities within the deployment. This can complement AAD-based access control and authorization restrictions within the SQL Database instance.

The Blueprint enables firewall functionality at the SQL Database instance. Rules are defined at two levels-- server (to restrict access to the SQL Database instance) and database (to further restrict access to individual databases within the instance). The principal advantage here is a defense-in-depth strategy. If the perimeter network boundary of the Blueprint deployment were compromised, the SQL Database firewall layers would still protect stored PHI (if provisioned with appropriate rulesets) within the instance. Also, firewall rules can be highly customized to specific databases, to further delimit access to specific identities and services.

[8] As a side note, KeyVault features a *soft-delete* recovery option can be enabled, whereby deleted keys, secrets, and even entire vault instances are recoverable within ninety (90) days of deletion. The option effectively retains the deleted items in non-deallocated memory for the recovery period, with automatic deallocation (i.e., permanent deletion) once the items have exceeded the period.

DATA STORAGE AND BACKUP

In addition to encrypting SQL Database instances, the Blueprint encrypts the Azure Storage instance of its deployment. This is accomplished via enablement of the Storage Security Encryption (SSE) instance setting. SSE derives an AES 256 key from KeyVault to encrypt all stored data. Azure Storage also supports TLS-based data transfer, so that PHI can be encrypted from origin to destination, while also supporting real-time inline decryption for processing purposes. There is no need to pre-encrypt data prior to introduction into the deployment.

The customer should note that the Azure Storage instance can be georeplicated across Azure datacenters. This is an inherent data availability and redundancy feature of the service provisioning environment. Indeed, the entire deployment itself can be georeplicated. The customer should carefully consider the regulatory ramifications of storing PHI in different geographic regions. It is recommended that PHI pertaining to United States citizens remain within US regions; conversely, PHI belonging to foreign nationals should remain within the respective national borders. Different nations implement widely divergent privacy regulations, ownership conditions, breach definitions, and penalties; thus, PHI storage in different international regions can incur significant organizational risk. Even within the US, different states implement different privacy regulations (for example, Texas, Vermont, Massachusetts, and California have distinct privacy regulation programs that apply to organizations operating within their jurisdictions, even if the organizations physically exist outside of the states); thus, the customer should take this into consideration as well.

Note that the data on Azure Storage instances may be subject to customer obligations for legalistic electronic discovery, record retention (e.g., under National Archives and Record Administration (NARA) guidelines), and data owner access. The customer should evaluate these obligations, and ensure deployment operations can support them. Azure Storage instances are based on scalable solid-state media; they can support the permanent storage of an unlimited volume of data. Conversely, the customer should also consider its obligations to dispose data (e.g., extracts of PHI records) after a defined time period.

Should the customer integrate VMs into the deployment, it can also employ the Azure Backup service (certified under the HITRUST CSF, but not included in the Blueprint) to automate VM backup for disaster recovery. Note that the customer should incorporate its Blueprint deployment into its disaster recovery and contingency plan processes (including annual test exercises or simulations to ensure that the deployment can be restored within acceptable recovery point and time objectives (RTO/RPO), by trained personnel who can expeditiously respond to a contingency event).

AUDITING AND MONITORING

The customer should carefully consider the selection of auditable events and data elements to capture via Security Center, OMS, and Application Insights. These services can aggregate a daunting array and volume of audit data; however, not all data may be immediately relevant to the customer's security or performance concerns. Log analysis and review activities should be tailored to events of concern, for efficiency and efficacy of incident response. This tailoring may require modification over time, as the customer's deployment evolves in functionality. The customer should also consider ensuring that specific PHI elements are not captured in log data. While not specifically a HITRUST CSF compliance issue, it is a best practice to curtail the proliferation of PHI to the minimal necessary set of components, identities, and environments. For example, it may be desirable that only data analysts can view PHI via PowerBI; thus, security analysts should not be able to see PHI in Security Center logs.

Per the HITRUST CSF, the customer should plan to retain all audit log data for a minimum of ninety (90) days online, and twelve (12) months offline (e.g., on dedicated Azure Storage instances). Audit logs, like data on Azure Storage instances, may also be subject to the customer's legalistic electronic discovery,

record retention, and data owner access obligations; it may need to be retained for a far longer period (on the order of several years to permanently).

The customer should plan to incorporate its Blueprint deployment into its incident management program, to address security incidents arising in relation to the deployment. However, Azure will alert customers to incidents affecting the security of the underlying service provision infrastructure. Note that the Security Center is informed by the Microsoft Global Threat Intelligence capability for real-time awareness and advisory of persistent and emergent security exploits.

VULNERABILITY MANAGEMENT

The customer should plan to incorporate its Blueprint deployment into its vulnerability management program. This would entail monthly scanning of the deployment to identify and assess vulnerabilities for remediation. The Security Center and OMS services can be core elements of this process. Further, the Security Center can support real-time situational awareness of security compliance state and real-time automated remediation; in essence, it can also be a critical component of the customer's continuous monitoring program. Note that the HITRUST CSF requires the performance of annual penetration testing by an external, independent party. Azure performs such testing for its service provision environment, and supports customer testing of deployments (so long as the test assays and rules of engagement do not extend beyond the deployment boundaries). The customer should consult with Azure to validate test performance and scheduling ahead of time.

ENDPOINT PROTECTION

The customer should note that the Blueprint does not rely upon Azure VM instances. The customer may choose to use on-premise endpoints or hosted VMs to interact with its deployment; in either case, the customer should carefully consider the extent to which the endpoints directly engage with PHI. If endpoints may transmit, process, or store deployment PHI, then they should implement similar confidentiality and integrity safeguards as the deployment. In particular, they should encrypt transmitted and stored PHI, restrict access to the minimum necessary set of identities, and monitor access behaviors. The OMS service can be extended to the VMs for change tracking of system and data files. When monitored files are modified, OMS can detect and alert on the change, providing detailed information on the change nature and actors. Thus, it can serve as a file integrity monitoring (FIM) solution.

OMS can also be extended to monitor VM configurations, and enforce mandatory settings. The deployment does not perform this function by default. Similarly, the Security Center and Application Insights services can also extend to the VMs for security and performance monitoring purposes. This is especially important for seamless antimalware, software update, flaw remediation, and configuration baseline management of the instances. Azure guarantees the availability of security-hardened VMIs, from which the customer instantiates VMs. The customer should be prepared, however, to subsequently harden its VMs on a continual basis.

A BLUEPRINT USE CASE

PREDICTING HOSPITAL LENGTH OF STAY

The use case deploys the Blueprint to predict a newly admitted patient's length of stay (LOS), based on Machine Learning-driven analysis of the patient's intake data against a historical data aggregation. ContosoClinic is a fictitious hospital network featured in the use case for illustration.[9] ContosoClinic's network administrators would like to predict patient LOS to optimize operational efficiency and enhance quality of care. ContosoClinic is certified under HITRUST CSF Version 9.0.

The basic Blueprint architecture of this use case is depicted in Figure 5.

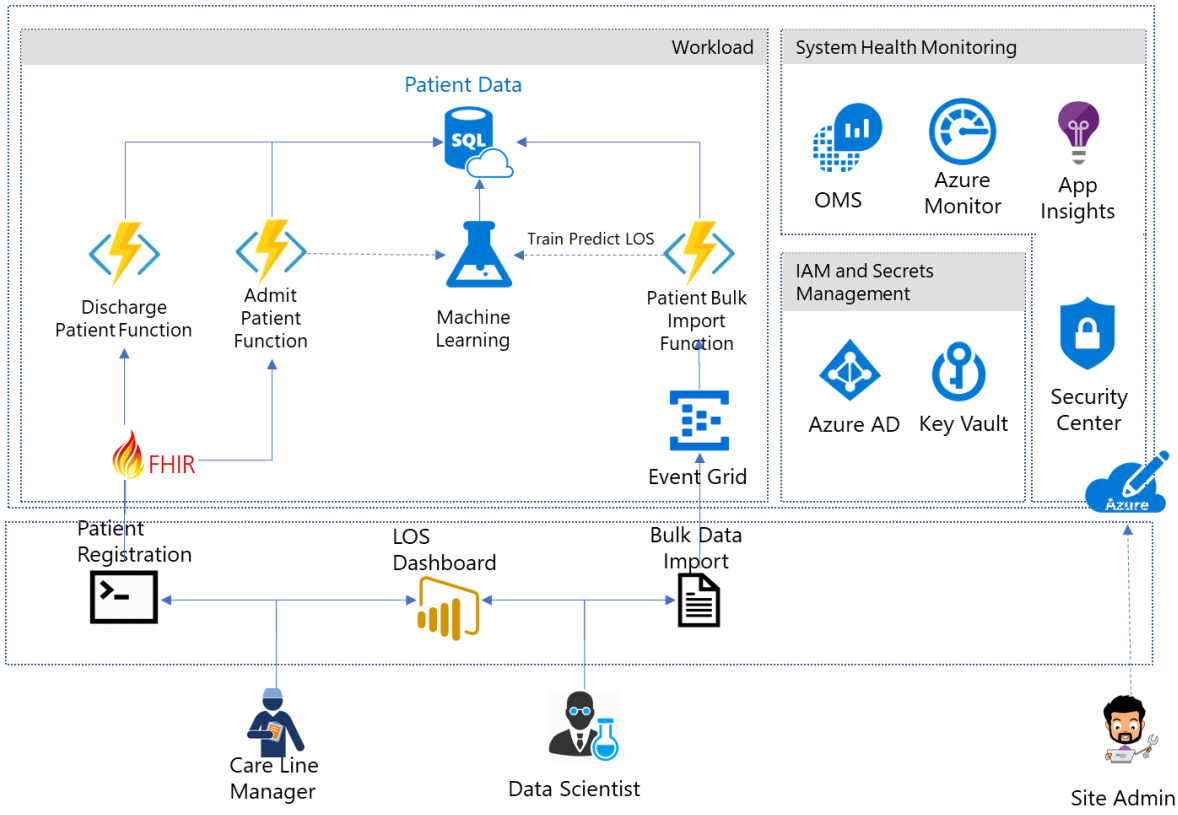

Figure 5: Blueprint Use Case Architecture

USE CASE ACTORS AND ROLES

9 In an actual production deployment, customers would use their own patient data to train Machine Learning models for LOS prediction.

The ContosoClinic use case features the following actors with specific Blueprint-related roles:

Site Administrator (Alex)

- Alex evaluates technologies that can reduce the IT overhead and cost of managing a hospital network. Alex has begun to evaluate Microsoft Azure service offerings, but struggles to understand how to configure specific services per ContosoClinic's HIPAA requirements. Alex recently selected the Blueprint as a compliance-ready turnkey solution. He will be responsible for the initial establishment, configuration, and administration of the deployment, as well as ongoing security management; but will have no access to PHI within it.

Data Scientist (Debra)

- Debra develops analytical models to glean actionable care insights from patient data. She is versed in SQL and R programming within an integrated development environment (IDE). She will use Machine Learning within the Blueprint to create, train, and optimize her LOS models. Debra will have the ability to import, export, manipulate, and report on data within Machine Learning; but this level of access does not extend to the SQL Database instance.

Database Analyst (Danny)

- Danny is the database engineer of ContosoClinic, skilled as a Microsoft SQL Server database architect, developer, and administrator. His expertise will nicely translate to Azure SQL Database (which is driven by a SQL Server database engine). He will develop relational database schema to support Debra's analytical models, administrate database performance and security, and potentially extend database capabilities through Transact-SQL (T-SQL) programming. He will have no visibility into PHI stored in the database instance.[10]

Care Line Manager (Chris)

- Chris is directly responsible for overseeing patient admission and discharge at ContosoClinic. He needs to ensure that facility staffing supports patient throughput at all times. Chris will interact with LOS dashboards in PowerBI to review LOS predictions; this will help him calibrate facility staffing in real time, and ensure that patient processing does not suffer diminished quality of service. He will have limited access to patient PHI for this purpose.

Chief Medical Information Officer (Caroline)

- Caroline, in collaboration with Chris and Debra, determines the dominant factors in patient length of stay, and adjusts facility resource allocation to support maximum quality of care. Caroline will use LOS dashboards in PowerBI to view predictions and accordingly manage resources in the hospital network. For instance, she can correlate long-term patient LOS with intake pathology trends, and allocate the right care resources (e.g., pulmonary, cardiovascular, neuropathological, or oncological) to the right facilities. Caroline's work can also help Debra further refine LOS prediction models. Caroline will have access to patient PHI.

Auditor (Han)

10 Though Danny will have DBA-level privileges within SQL Database, he cannot bypass TDE to view plaintext PHI because his access is mediated through an Azure Function that returns only the cryptotext. This was an intentional Blueprint design choice.

- Han is a HITRUST CSF certified auditor who is also experienced in ISO and SOC compliance. He was hired to review Contosoclinc's network. Han will review the Blueprint Customer Responsibility Matrix, as well as ContosoClinic's documented policies and procedures; interview technical, operational, and management personnel; and assess the Blueprint deployment to determine its degree of security compliance. He may perform technical testing on some aspects of the Blueprint deployment, or may observe some deployment operations involving PHI; however, he will have no access to PHI.

OPERATIONS AND SECURITY CONFIGURATION

This section details the default Blueprint configurations and security measures for ContosoClinic's use. The organization relies upon its own network infrastructure and on-premise endpoints for access to the deployment in Azure; thus, the section does not discuss use of non-Blueprint Azure services.

AAD:

- Centralized, LDAP-based access control, authentication, authorization, and domain security governing all identities and services.
- AAD Identity Protection to detect, notify, and investigate potential vulnerabilities affecting ContosoClinic identities.
- MFA via the `-enableMFA` configuration switch.
- Password expiration after 60 days via the `-enableADDomainPasswordPolicy` configuration switch.
- Built-in domain roles to assign permissions to identities via membership.
- Domain auditing of identity and service behaviors.

Event Grid and Azure Functions:

- AAD integration for access control, authentication, authorization.
- Enablement of event and access logging.
- All data flow between components transacted by authenticated Azure Function calls

SQL Database:

- AAD integration for access control, authentication, authorization.
- Transparent Data Encryption (TDE) with column-level encryption to restrict access to patient PHI elements (first and last names, date of birth, social security number, etc.). DEKs are AES 256-based and stored in KeyVault.
- SQL Vulnerability Assessment to discover, track, and remediate potential database vulnerabilities.
- SQL Database Threat Detection to detect, notify, and counteract potential database exploits in progress (analogous to AAD Identity Protection).
- SQL Database Auditing of identity and service behaviors.
- SQL Database performance counters to quantitatively monitor quality of service.
- Dynamic data masking to restrict PHI viewing ability to Debra, Caroline, and Chris.

- Azure Function restriction of Danny's PHI visibility.
- Server- and database-level firewalling with default deny-all/allow-exception rulesets.

Azure Storage:

- AAD integration for access control, authentication, authorization. Authentication request data is validated by AAD and KeyVault.
- HTTPS endpoints to transmit data to storage via TLS.
- Anonymous access is disallowed for storage containers.
- Auditing of storage access by identities and services, with automated alerting of anonymous access.
- SSE to encrypt all data in containers.
- No storage quotas are set.

Machine Learning:

- AAD integration for access control, authentication, authorization.
- Auditing of Machine Learning web service access by identities and services.

Key Vault:

- Key Vault storage of:
    - Application Insights key
    - TLS certificates
    - Machine Learning web service endpoint key
    - Machine Learning service API key
    - SQL Database DEKs
    - TLS certificates
    - Patient data storage access key
    - Patient data connection string
    - Patient data table name
- AAD integration for access control, authentication, authorization.
- Key management policy for key strength, lifespan, rotation, revocation, escrow, and archivation.
- All keys expire within 12 months of inception.
- All keys are protected within HSMs.
- Applications have unique keys (unless interoperation at runtime requires shared secrets).
- Auditing of KeyVault web service access by identities and services.
- Permitted cryptographic operations delimited to those strictly required for Blueprint operation.

Security Center:

- AAD integration for access control, authentication, authorization.
- Microsoft Global Threat Intelligence support.
- Real-time security anomaly monitoring.
- Event notification via Azure Alerts.

Application Insights:

- AAD integration for access control, authentication, authorization.
- Interoperability with OMS and Machine Learning for predictive performance analysis.
- Interoperability with DevOps to support continuous configuration management.
- Real-time performance telemetry monitoring.
- Event notification via Azure Alerts.

OMS:

- AAD integration for access control, authentication, authorization.
- Workspace is enabled for Security Center and Workload Monitoring.
- Workload Monitoring is enabled for:
    - Activity Log Analytics
    - Azure WebApp Analytics
    - Change Tracking
    - Identity and Access
    - Key Vault Analytics
    - Security and Audit
    - SQL DB Analytics
- Application Insights Connector is enabled.
- Event notification via Azure Alerts.

Data classification and FHIR:

- All sensitive data in the Blueprint is tagged as electronic protected health information (ePHI), as follows:
    - *dataProfile* => “ePHI”
    - *owner* => <Side Admin UPN>
    - *environment* => “Pilot”
    - *department* => “Global Ecosystem”

    - *tier* => API | Application | DataStore | Operations
- The Blueprint implements the following FHIR resources for data ingestion and output:
    - Condition
    - Encounter
    - Observation
    - Patient
- The Blueprint can be extended to add support for additional resources as needed.

CONCLUSION

The Microsoft Azure Security and Compliance – HITRUST Health Data and Artificial Intelligence (AI) Blueprint is a robust foundation for developing health analytics solutions-- flexible in application, powerful in capability, and elegant in design. Additionally, it embodies information security principles in architecture, deployment, and operation. As such, customers can deploy the Blueprint within their existing HITRUST CSF compliance posture, comfortable that the confidentiality and integrity of PHI can be preserved.

However, this comfortability rests upon a shared responsibility model, wherein both Azure and customer play critical security assurance roles. The Blueprint is securely engineered and provisioned, using Azure cloud infrastructure and platform resources. Secure configuration, implementation, and usage resides with the customer. A complete picture of HITRUST CSF compliance for the Blueprint encompasses the certification of both parties, demonstrating appropriate due diligence for both service provision and consumption.

APPENDIX A: TABLE OF ACRONYMS

Acronym	Full Term
3DES	Triple Data Encryption Standard
AAD	Azure Active Directory
ACL	Access Control List
AES	Advanced Encryption Standard
AI	Artificial Intelligence
AUP	Acceptable User Policy
BA	Business Associate
BAA	Business Associate Agreement
BI	Business Intelligence
BLOB	Binary Large Object
BYOD	Bring Your Own Device
Capex	Capital Expenditure
CE	Covered Entity
CMS	Center for Medicare and Medicaid Services
CRM	Customer Responsibility Matrix
CSF	Common Security Framework
CSP	Cloud Service Provider
DHHS	Department of Health and Human Services
DICOM	Digital Imaging and Communications in Medicine
DMZ	Demilitarized Zone
FedRAMP	Federal Risk and Authorization Management Program
FHIR	Fast Healthcare Interoperability Resources
FIM	File Integrity Monitoring
FIPS	Federal information Processing Standard
FISMA	Federal Information Security Management Act
GPO	Group Policy Object
HIPAA	Health Insurance Portability and Accountability Act
HITRUST	Health Information Trust
HSM	Hardware Security Module
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IaaS	Infrastructure as a Service
ICD	International Classification of Diseases

IDE	Integrated Development Environment
IDPS	Intrusion Detection/Prevention System
IP	Internet Protocol
IPSec	Internet Protocol Secure
ISMS	Information Security Management System
ISO	International Standards Organization
IT	Information Technology
LDAP	Lightweight Directory Access Protocol
LOS	Length of Stay
MDM	Mobile Device Management
MFA	Multifactor Authentication
NACL	Network Access Control List
NARA	National Archives and Records Administration
NIST	National Institute of Standards and Technology
NSG	Network Security Group
OCR	Office of Civil Rights
OMS	Operations Management Suite
Opex	Operating Expense
OVAL	Open Vulnerability Assessment Language
OWASP	Open Web Application Security Project
PaaS	Platform as a Service
PCI DSS	Payment Card Industry Data Security Standard
PHI	Protected Health Information
PRISMA	Program Review for Information Security Management Assistance
RBAC	Role Based Access Control
RoB	Rules of Behavior
RSA	Rivest Shamir Adelman
SaaS	Software as a Service
SAS	Shared Access Signature
SCAP	Security Content Automation Protocol
SDL	Software Development Lifecycle
SHA	Secure Hash Algorithm
SIEM	Security Information and Event Management
SP	Special Publication
SME	Subject Matter Expert
SQL	Structured Query Language

SSE	Storage Service Encryption
SSL	Secure Socket Layer
TDE	Transparent Data Encryption
TLS	Transport Layer Encryption
URI	Uniform Resource Identifier
URL	Universal Resource Locator
VM	Virtual Machine
VMI	Virtual Machine Image
VPN	Virtual Private Network

APPENDIX B: TABLE OF REFERENCES

Tabulated below are materials referenced in this whitepaper. The reader may wish to further explore these for additional background, guidance, and insight.

Health Insurance Portability And Accountability Act (HIPAA)	*HIPAA Security Rule* (CFR Title 45, Volume 1, Part 164 Subpart C)
	HIPAA Privacy Rule (CFR Title 45, Volume 1, Part 164 Subpart E)
	HIPAA Final Rule (CFR Title 45, Volume 1, Parts 160 and 164)
	HITECH Interim Final Rule (Public Law 111-5, Title XIII)
	An Introduction Resource Guide for Implementing the HIPAA Security Rule (NIST SP 800-66)
HITRUST Alliance	*An Introduction to the HITRUST Common Security Framework (CSF)*
	HITRUST Common Security Framework (CSF) Version 9
	HITRUST Common Security Framework (CSF) Summary of Changes
	HITRUST CSF Standards and Regulations Cross-reference
	Risk Analysis Guide for HITRUST Organizations and Assessors
Microsoft Azure	*Microsoft Azure Security and Compliance – HITRUST Health Data and Artificial Intelligence (AI) Blueprint Demo* (https://github.com/Azure/Health-Data-and-AI-Blueprint)
	Microsoft Azure Customer Responsibility Matrix (CRM) (http://aka.ms/HealthCRMBlueprint)
	Microsoft Business Associate Agreement (BAA) (http://aka.ms/BAA)
	Microsoft Trust Center (http://www.Microsoft.com/TrustCenter)

ABOUT THE AUTHORS

Michael T. Williams | Senior Consultant | Healthcare & Life Sciences

Michael Williams is a Senior Consultant in the Coalfire Healthcare and Life Sciences practice, wherein he serves a diverse array of healthcare organizations on HIPAA/HITECH and HITRUST projects. These range from advisory engagements for clients exploring their capacities to achieve healthcare security compliance, to assessment engagements to help clients maintain existing certifications.

Alison Heidt | Associate Consultant | Healthcare & Life Sciences

Alison Heidt is an Associate Security Consultant for the Coalfire Healthcare and Life Sciences practice. As an Associate Consultant, she actively supports a broad assortment of Coalfire clients, collaborating with engagement teams to deliver timely, high-quality advisory and assessment services.

Published February 2018.

ABOUT COALFIRE

Coalfire is the 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 more than 16 years and has offices throughout the United States and Europe.