



HIPAA and the Data Warehouse Built for the Cloud

HOW HEALTHCARE BUSINESS ASSOCIATES CAN MEET HIPAA
SECURITY REQUIREMENTS AND TECHNICAL SAFEGUARDS





Contents

3	Safeguarding Protected Healthcare Information (PHI)
5	Is your organization a HIPAA business associate?
7	How does the cloud data warehouse fit into a BA's HIPAA computing requirements?
9	An overview of HIPAA Technical Safeguards
11	A checklist of technologies to meet HIPAA safeguards
12	Encryption everywhere: Key management
14	More best practices in encryption key management
15	Encryption everywhere: End-to-end encryption
16	Advanced access controls
18	Ensuring HIPAA data integrity: Time travel
19	Additional supporting technologies
20	The fast path to a HIPAA-ready cloud data warehouse
21	About Snowflake

Safeguarding Protected Healthcare Information (PHI)

Defenses against data breaches

The ten largest healthcare data breaches of 2016 included nine providers and one business associate (BA), affecting nearly 13 million individuals.¹ The cyber security incident of the business associate involved the personal health information of 3.3 million individuals, making it the third largest breach of the year.

Even more alarming, the number of cyber attacks on the healthcare industry has more than doubled since 2010. Attacks on the millions of healthcare business associates will continue to increase. Many are using ageing computer systems that do

not use the latest security technology. In addition, they may be unaware of the security measures needed to safeguard protected health information (PHI).

To cyber criminals, the continued allure of PHI contained in medical records remains clear: Patient names, addresses, medical record numbers, birthdays and Social Security numbers can be used to fraudulently obtain credit cards or loans and commit tax or insurance fraud. A single medical record is worth 10 to 20 times more than a US credit card number.



COMPLIANCE WITH THE HIPAA SECURITY RULE

To keep PHI out of the hands of criminals, business associates are subject to the HIPAA Security Rule, a national set of security standards for protecting certain health information that is held or transferred in electronic form. Complying with the standards set forth in the HIPAA Security Rule helps BAs to accomplish three critical objectives:

- Significantly mitigate the risk of data breach
- Minimize the loss of readable records in the event of a breach
- Avoid penalties for non-compliance, as enforced by the Office for Civil Rights of the U.S. Department of Health & Human Services.

A CHECKLIST FOR PHI SECURITY

This ebook is specifically designed to help business associates understand how the HIPAA Security Rule translates into requirements for technology environments used to manage PHI data. These environments increasingly center on cloud data warehouses for the data-intensive analytics and optimization services that many BAs offer.

This ebook provides a short summary of the benefits of cloud data warehousing and an overview of the HIPAA Security Rule and its Technical Safeguards. From there, you'll get a checklist of security requirements that BAs can use to assess the appropriateness and business value of cloud data warehouse options. Armed with this information, you can champion the highest security measures available to protect your organization and the PHI of millions of people you've been tasked to keep secure.

Is Your Organization a HIPAA Business Associate?

An expanded definition under the HIPAA Omnibus Rule

The U.S. Department of Health and Human Services (HHS) estimates there are up to two million HIPAA business associates. Yet, many organizations that are BAs may not be aware of their status. Thus, you may be unaware of your responsibility to protect PHI.

A DEFINITION THAT HAS EVOLVED OVER TIME

HHS defines a business associate as: “a person or entity that performs certain functions or activities that involve the use or disclosure of protected health information on behalf of, or provides services to, a covered entity....”

“Business associate functions and activities include: claims processing or administration; data analysis, processing or administration; utilization review; quality assurance; billing; benefit management; practice management; and repricing. Business associate services are: legal; actuarial; accounting; consulting; data aggregation; management; administrative; accreditation; and financial.”

However, in its white paper, “Who is a HIPAA business associate?”, the law firm McDonald Hopkins notes:

The Omnibus Rule [an amendment to HIPAA] added the following new categories of business associates:

- Those who store or otherwise maintain PHI.
- Health Information Organizations (HIOs), e-prescribing gateways and others that provide data transmission services to a covered entity and require routine access to PHI.
- Anyone who offers a personal health record to individuals on behalf of a covered entity.
- Subcontractors of business associates, if (i) the business associate delegates to the subcontractor a function, activity or service that the business associate has agreed to perform for the covered entity, or for another business associate, and (ii) any of the delegated functions, activities or services involve the creation, receipt, maintenance or transmission of PHI.

UNWITTINGLY ENSNARED IN HIPAA REQUIREMENTS

However, under the Omnibus rule, McDonald Hopkins notes: “HIPAA obligations and potential liability can extend to subcontractors who have no direct connection or relationship with any covered entity, no matter how far the PHI flows down the chain from business associate to subcontractors or how little the subcontractor knows about the relationship with the covered entity.”

The graphic below captures a typical scenario:



In its report the firm states:

In these circumstances, business associate agreements would be required between:

- The covered entity and business associate A
- business associate A and subcontractor B
- Subcontractor B and subcontractor C
- Subcontractor C and subcontractor D

The extension of business associate status to subcontractors can ensnare unsuspecting individuals and organizations because, prior to the Omnibus Rule, subcontractors were untouched by the HIPAA Rules. Many could still be unaware that they are performing functions for covered entities or dealing with PHI.

The bottom line: If in doubt, any organization handling PHI, through direct or indirect engagement with a covered entity or other business associate, should take the steps necessary to meet HIPAA data security requirements.



How Does the Cloud Data Warehouse Fit into a BA's HIPAA Computing Requirements?

A focus on core competencies

HIPAA business associates provide a wide range of analytic and administrative services, yet all share a common trait: BAs' core strengths lie in data-related services, not in IT. As a result, BAs are increasingly moving the data warehouses that store PHI to the cloud, saving the significant expense of buying, maintaining and securing on-premises systems.

In addition to cost savings on:

- Hardware, software licenses and maintenance
- IT infrastructure and data centers
- The salaries of the highly specialized people who support these technology operations

... the right data warehouse, built for the cloud, can deliver benefits including:

- **A HIPAA-ready, secure environment:** As detailed in the balance of this ebook, BAs must maintain an extensive security infrastructure to address five HIPAA Technical Safeguards: Access Control, Audit Control, Integrity, Person or Entity Authentication, and Transmission Security.

Properly managed, the security measures provided by a cloud data warehouse can be a much more effective and less expensive option than BAs attempting to manage the security infrastructure on their own. The right cloud data

warehouse provider is expert in HIPAA-ready security, provides this technology to thousands of customers, is always up to date with best-of-breed security practices and can assist BAs with their HIPAA-related security concerns.

- **Faster deployment:** A data warehouse built for the cloud can go live in weeks or just a few months. BAs can see benefits much sooner with a fraction of the cost of the upfront investment than with an on-premises solution.
- **Software upgrades:** On-premises and "cloud-washed", on-premises solutions take a standard "waterfall" development approach to functionality updates. To enable the annual or biannual update, IT must usually take the system down or place it in maintenance mode – more lost time and money. To avoid this, IT may anchor to a specific version of the software, which creates another set of challenges.

With a modern cloud data warehouse built for the cloud, the upgrades should originate from an agile DevOps approach – incremental updates every month that avoid any disruption to customers.

- **Lower staffing costs:** As noted above, the number of people who maintain an on-premises, enterprise data warehouse and supporting infrastructure can be an enormous expense. Depending on the cloud alternative, BAs can significantly reduce or nearly eliminate this expense, depending on the level of functionality, automation and vendor management of the solution.
- **Pay only for what you use:** An on-premises data warehouse forces BAs to buy enough storage space and compute horsepower to handle peak demand. Peak demand represents only a few days a month, or a few weeks at the end of a financial year. With the right cloud data warehouse, BAs can pay only for the compute and storage resources actually used. As an additional factor to consider, the cost of the actual storage and compute resources should be significantly lower with a cloud solution due to cloud economies of scale.

Altogether, these benefits are driving many business associates to champion cloud data warehouse technology within their organizations. With the right security capabilities in place, BAs can adopt built-for-the-cloud data warehousing as a high-performance, and highly secure, HIPAA-ready platform for storing and working with PHI.



An Overview of HIPAA Technical Safeguards

Broad guidelines for protecting PHI

Personal health information stored and analyzed in cloud data warehouses is subject to HIPAA security requirements:

The Security Rule defines technical safeguards in § 164.304 as “the technology and the policy and procedures for its use that protect electronic protected health information and control access to it.”

– HIPAA Security Series, Centers for Medicare & Medicaid Services, U.S. Department of Health & Human Services

The HIPAA Technical Safeguards do not entail specific products or technologies. Instead, they deliver five categories of broad security measures that BAs must incorporate into their HIPAA technology environments, including cloud data warehouses. These five categories are:

1. **Access Control:** Official HIPAA language describes access control as, “The ability or the means necessary to read, write, modify, or communicate data/information or otherwise use any system resource.” Further, the Access Control standard requires a BA to “implement technical policies and procedures for electronic information systems that maintain electronic protected health information to allow access only to those persons or software programs that have been granted access rights as specified in § 164.308(a)(4)[Information Access Management].”

Four implementation specifications are associated with the Access Controls standard.

- I. **Unique User Identification (Required)**
 - II. **Emergency Access Procedure (Required)**
 - III. **Automatic Logoff (Addressable)**
 - IV. **Encryption and Decryption (Addressable)**
2. **Audit Control:** The official language states business associates must “implement hardware, software, and/or procedural mechanisms that record and examine activity in information systems that contain or use electronic protected health information.” The U.S. Department of Health & Human Services notes, “Most information systems provide some level of audit controls with a reporting method, such as audit reports. These controls are useful for recording and examining information system activity, especially when determining if a security violation occurred.”

3. **Integrity:** Integrity is defined in the Security Rule, at § 164.304, as “the property that data or information have not been altered or destroyed in an unauthorized manner.” Clearly, protecting the integrity of electronic PHI is a primary goal of the Security Rule. Here, the integrity standard requires business associates to “implement policies and procedures to protect electronic protected health information from improper alteration or destruction.”
4. **Person or Entity Authentication:** HIPAA language states that BAs must “implement procedures to verify that a person or entity seeking access to electronic protected health information is the one claimed.”

5. **Transmission Security:** Finally, BAs must “implement technical security measures to guard against unauthorized access to electronic protected health information that is being transmitted over an electronic communications network.” This standard has two implementation specifications:

- I. **Integrity Controls (Addressable)**
- II. **Encryption (Addressable)**

Where this implementation specification is a reasonable and appropriate safeguard for a covered entity, the covered entity must “implement a mechanism to encrypt electronic protected health information whenever deemed appropriate.”

How do these requirements translate into a checklist of technologies that BAs should look for in the cloud data warehousing services they may use? Flip to the next page to find out.



A Checklist of Technologies to Meet HIPAA Safeguards

How BAs can lighten their security burden

Building and maintaining a technology environment to meet HIPAA requirements is a highly complex and costly task. However, the right cloud data warehouse platform can provide critical ongoing support for HIPAA Technical Safeguards, allowing BAs to fully focus on their core competencies.

The next few pages describe these critical security technologies, and why they deliver significant protection, value and peace of mind for business associates.

The table below shows how specific, modern cloud data warehousing features can satisfy HIPAA technical requirements. Each feature is explained in the following pages.

HIPAA REQ.	TRANSMISSION SECURITY	ACCESS CONTROL	PERSON OR ENTITY AUTHENTICATION	INTEGRITY	AUDIT CONTROL
FEATURES	OCSP	Advanced Access Controls <ul style="list-style-type: none"> • Multi-factor authentication • Federated authentication and SSO • IP whitelisting 		Time Travel	Logging & Reporting
	Encryption Everywhere				

WHY IT MATTERS: A HIPAA-READY DATA WAREHOUSE BUILT FOR THE CLOUD

There are many cloud data warehouses but most are not built for the cloud. Most offerings are actually “cloud-washed” versions of warehouses originally designed for on-premises enterprise use.

When it comes to HIPAA security, the “built for the cloud” distinction is extremely important because it determines whether a business associate assumes none, some or all of responsibility to create and manage security features. Specifically, only a vendor offering a built-for-the-cloud SaaS (software as a service) data warehouse can deliver a HIPAA-ready environment that incorporates all of the necessary security capabilities. Here, the BA is not responsible for building or maintaining any of these features since they are already built into the SaaS cloud data warehouse offering.

In contrast, cloud data warehouses offered as IaaS (infrastructure-as-a-service) or PaaS (platform-as-a-service) leave security up to the customer to build and/or manage.

For any BA considering using a cloud data warehouse service, it’s important to find out exactly which security features the vendor provides and which are the responsibility of the BA.

Encryption Everywhere: Key Management

Helping to ensure the security of stored PHI

Encryption is at the heart of all data security, whether the data is stored on premises or in the cloud. Today's best practices for securing stored PHI include AES-256 – a data encryption method used to prevent unauthorized access to files. AES-256 uses strong 256-bit encryption and is widely considered best in class. Of AES, the National Institute of Standards and Technology (NIST) states, "Whenever possible, AES (Advanced Encryption Standard) should be used for the encryption algorithm because of its strength and speed."

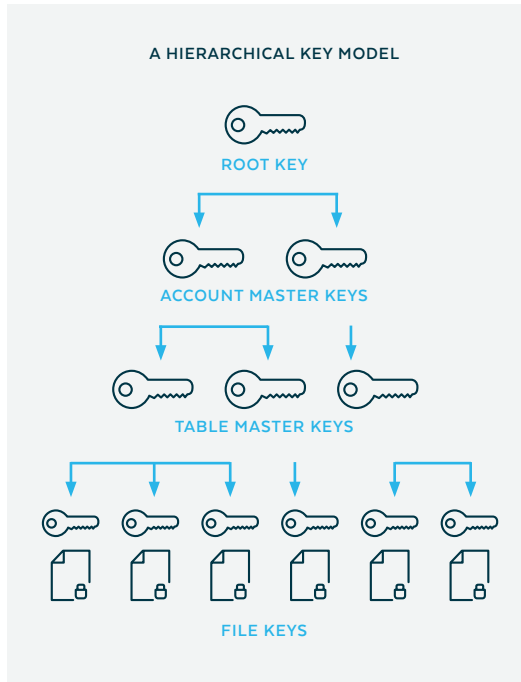
KEY MANAGEMENT IS ... KEY!

As a best practice, all PHI data in a cloud data warehouse should be encrypted by default, using AES 256-bit encryption. But encryption of the PHI data is just the first step. Therefore, how the encryption keys are managed is of paramount importance to help ensure only authorized data access.

Keys should be managed in a hierarchical model and automatically rotated on a regular basis by the cloud data warehouse provider. This allows data to be automatically re-encrypted ("rekeyed") on a regular basis. Data encryption and key management should be entirely transparent to the BA, requiring no configuration or management.

BEST PRACTICES FOR ENCRYPTION KEY MANAGEMENT

A hierarchical key model is the cornerstone of best-practice encryption key management. A key hierarchy has several layers of keys, where each layer of keys (the parent keys) encrypts the layer below (the child keys). When a key encrypts another key, security experts refer to it as "wrapping." In other words, a parent key in a key hierarchy wraps all of its child keys.



A hierarchical key model provides tighter security in a multi-tenant cloud service – the standard delivery model for cloud-based data warehouses. With such a model, all customer accounts are isolated from each other because each account has a separate key hierarchy. As a best practice, this layer of isolation should be used in addition to other access control techniques such as multi-factor authentication, federated services and IP whitelisting, which are described later in this ebook.

A hierarchical key model ideally consists of four levels of keys: the root key, account master keys, table master keys, and file keys. Each account master key corresponds to one customer account in the cloud data warehouse. Each table master key corresponds to one table in a database. That means that every account and every table is encrypted with a separate key. Similarly, every single data file is encrypted with a separate key and wrapped by one of the table master keys.



More Best Practices in Encryption Key Management

Important follow-through to keep PHI safe

Encryption key rotation manages the keys through different states during their life cycle. In the active state, the key is used to encrypt data and is available for use by the originator. In the retired state, the key is used only to decrypt data and is available to the recipient. When a key is discarded, it's used for neither encryption nor decryption.

Key rotation ensures that keys go from the active state to the retired state in this life cycle during a limited period of time. In addition, with key rotation, new data gets fresh keys.

Rekeying is the process of re-encrypting data with new keys. After a specific time period, new encryption keys replace old encryption keys. This is independent and completely orthogonal to key rotation. While key rotation ensures that a key is transferred from its active state (originator usage) to the retired state (recipient usage), rekeying ensures that a key is transferred from its retired state to being destroyed.

In other words:

KEY ROTATION = "new data gets fresh keys"

REKEYING = "old data gets fresh keys"

Rekeying, therefore, completes the life cycle of keys by ensuring that keys can be destroyed.

Encryption Everywhere: End-to-End Encryption

Reduce the vulnerability to attack

End-to-end encryption covers all data states:

- Data in flight to the cloud data warehouse
- Data moving inside the cloud data warehouse
- Data at rest

As such, end-to-end encryption is a technology best practice to ensure that only the business associate and the cloud data warehouse service can read the data—no one else. This includes web service providers such as Amazon, Microsoft Azure and others. End-to-end encryption is thus the most secure way for BAs to communicate with their cloud data warehouse.

End-to-end encryption is important because it minimizes the attack surface of any data BAs may have stored in the cloud data warehouse. In the case of a security breach of any third party (for example, of Amazon's cloud-based Simple Storage Service [S3]), PHI and other data are protected because they're always encrypted, regardless of whether the breach is due to:

- The exposure of access credentials to the underlying cloud storage provider's infrastructure.
- The exposure of data files directly by the cloud storage provider.
- Whether the breach was perpetrated by an insider or by an external attacker.
- Whether the breach was inadvertent or intentional.

As a best practice, the encryption keys should be only in the custody of the BA and the cloud data warehouse provider to minimize vulnerability.

READ THE FINE PRINT

Not all cloud data warehouses provide end-to-end encryption, a hierarchical key model, key rotation and rekeying. It's important for BAs to carefully read prospective cloud services providers' Business Associate Agreement (BAA) to understand which security measures the provider will deliver and which are expected of the BA. Again, this varies widely based on whether the cloud data warehouse is delivered as SaaS, IaaS or PaaS.

Advanced Access Controls

Seamless protections to guard against unauthorized access

MULTI-FACTOR AUTHENTICATION

Multi-factor authentication (MFA) is a security mechanism that requires users to be authenticated through more than validation procedure. MFA is built from a combination of physical, logical and biometric validation techniques, and, as a best practice, should be used to secure business associates' access to their cloud data warehouse.

To gain access to a secured location or system, MFA requires security measures to validate and authenticate a user based on what they know, such as a unique passcode, and what they have – a device such as a mobile phone or security token from which to receive a one-time passcode.

FEDERATED AUTHENTICATION AND SINGLE SIGN-ON

In a cloud data warehouse environment, the use of federated authentication services means that users don't gain access to the data warehouse with a direct login. The login process requires authentication through the business associate's server. In other words, the passwords and policies don't exist in the cloud data warehouse, they exist only in the BA's systems.

Here's how the process works:

- Federated authentication does not validate the user's actual password on the cloud data warehouse platform. Instead, the platform receives a SAML assertion in an HTTP POST request.
- The SAML assertion has a limited validity period, contains a unique identifier, and is digitally signed.

- If the assertion is still within its validity period, has an identifier that has not been used before, and has a valid signature from a trusted identity provider, the user is granted access to the application.
- If the assertion fails validation for any reason, the user is informed that their credentials are invalid.

Federated authentication is typically used in conjunction with single sign-on (SSO), which allows a single authentication credential—such as a user ID and password, smart card, one-time password token or a biometric device—to access multiple or different systems within a single organization. For a BA, these systems can include cloud applications and systems that are separate from the cloud data warehouse.

Why it matters: In addition to saving users from remembering and entering multiple authentication credentials, federated authentication and SSO make it easy for system administrators to revoke access privileges when employees leave the BA or as necessary, eliminating opportunities for unauthorized access.

IP WHITELISTING

A whitelist is a list of IP addresses that are granted access to the BA's cloud data warehouse. When an IP whitelist is used, all other entities are denied access. As a best practice, the cloud data warehouse provider can offer more granular access options. These include allowing people to log in from certain IP addresses, such as users working at home, to gain access to only certain data sets. For other, more sensitive data, access can be restricted only to in-office environments.

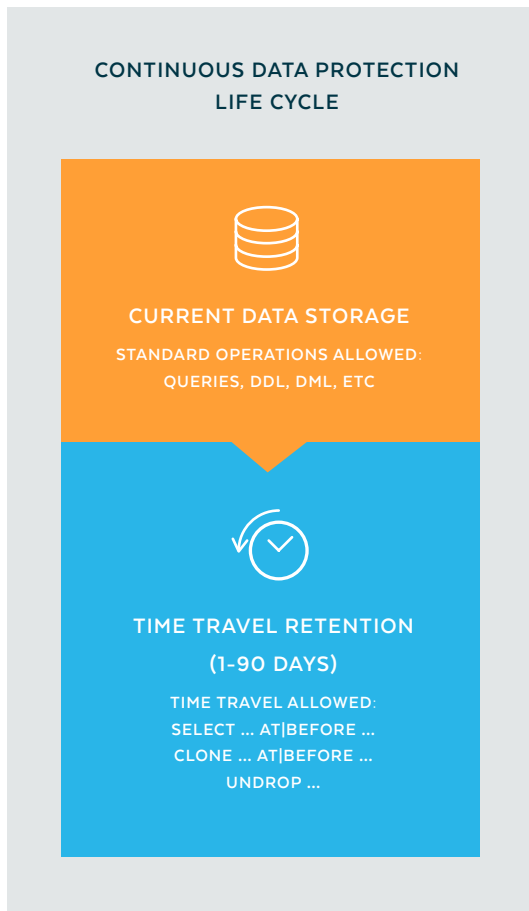
Why it matters: When used in conjunction with encryption, MFA, federated authentication and SSO, IP whitelisting is a powerful security tool in restricting access to PHI to authorized users only.



Ensuring HIPAA Data Integrity: Time Travel

Restoring a cloud data warehouse to a previous state, with queries intact

Data breaches aren't the only threat to BAs working with PHI. Data destruction, either accidental or intentional, is also a serious risk. As a best practice, a cloud data warehouse should make it easy for BAs to recover from data corruption or loss with time travel, a rollback feature that restores previous versions of data and the queries made against that data.



In the most advanced cloud data warehouse systems, a time travel feature is available for up to 90 days. With it, users can perform multiple actions within a defined time period within the 90-day window:

- Use saved or new queries to analyze past data that has since been updated or deleted.
- Create clones of entire tables, schemas and databases at or before specific points in the past.
- Restore tables, schemas and databases that have been dropped.

Time travel goes far beyond HIPAA's integrity safeguard of implementing "policies and procedures" to protect PHI. Accidents happen. Or, hackers aren't motivated to steal your data – they see pleasure in simply deleting it. When any of these disastrous events happen, a time travel feature can play an invaluable role in assuring the integrity of PHI.

Additional Supporting Technologies

A roundup of best-practice capabilities

ONLINE CERTIFICATE STATUS PROTOCOL (OCSP)

OCSP is a common scheme for maintaining the security of cloud resources, and is important in ensuring the safe transmission of PHI. When a user attempts to access a cloud data warehouse, using a secure protocol such as HTTPS, OCSP is used to send a request for certificate status information. The cloud data warehouse receives back a response of “current,” “expired,” or “unknown.”

Why it matters: Browsers and other apps will generally warn the user about certificate revocation and prevent them from proceeding. If the user was allowed to visit a site with a revoked certificate, it's possible the user would not be communicating with the intended site but with a malicious man-in-the-middle (MITM). OCSP plays a critical role in thwarting MITM attacks.

LOGGING AND REPORTING

As a best practice, the cloud data warehouse provider should provide logging and reporting mechanisms that allow business associates to see up to 100 complete previous entries:

- Who logged in
- Which data they accessed
- Which queries were run
- The duration of the user's session

The cloud data warehouse service should also let a BA's administrators create alerts to let them know:

- When a specific user has logged in
- Where an IP address came from
- What the user tried to access and if they were successful

As a best practice, the cloud data warehouse should send a daily report containing all of this information.

Why it matters: The ability to see who's accessed PHI and how users have manipulated protected data is central to HIPAA's Audit Control Technical Safeguard.

The Fast Path to a HIPAA-Ready Cloud Data Warehouse

Reducing risk with a best-practice checklist

Having the right technology in place is more important than ever. The HHS Office for Civil Rights (OCR) conducts periodic audits of a BA's compliance with HIPAA Privacy, Security, and Breach Notification Rules. Building HIPAA-compliant security capabilities requires significant technology expertise, time and investment, which diverts resources and focus away from a BA's core business.

Today, more and more BAs are deploying a best-in-class, SaaS, cloud data warehouse as the fast path to compliance with HIPAA Security Requirements and Technical Safeguards.

REDUCE RISK AND ACCELERATE YOUR ORGANIZATION'S ACQUISITION OF BEST-IN-CLASS, HIPAA-READY, CLOUD DATA WAREHOUSE TECHNOLOGY BUILT FOR THE CLOUD.

Visit snowflake.net

CLOUD DATA WAREHOUSE HIPAA-READY CHECKLIST	
✓	Encryption everywhere
✓	Multi-factor authentication
✓	Time travel
✓	Federated authentication and SSO
✓	IP whitelisting
✓	OCSP
✓	Logging and reporting

About Snowflake

Snowflake started with a clear vision: Make modern data warehousing effective, affordable and accessible to all data users. Snowflake delivers the performance, concurrency and simplicity needed to store and analyze all of an organization's data in one location. Because traditional on-premises and cloud solutions struggle with this, Snowflake developed a new product with a new built-for-the-cloud architecture that combines the power of data warehousing, the flexibility of big data platforms and the elasticity of the cloud at a fraction of the cost of traditional solutions. Snowflake: Your data, no limits.

Visit [snowflake.net](https://www.snowflake.net)

DISCLAIMER

This ebook is an information resource developed by Snowflake Computing. It is neither a complete compendium of technology requirements for HIPAA compliance, nor a legal document. For additional information on HIPAA security requirements and compliance, business associates should consult their Chief Privacy Officer and General Counsel.