Gain confidence about data security in the cloud

dashDB: A use case

Walid Rjaibi

August 07, 2014

This tutorial demystifies cloud security and arms you with the know-how to adopt the cloud with confidence. Learn how cloud security is a shared responsibility between the cloud service provider and the client. The responsibilities of each party are explored. Also walk through a data security use case involving dashDB. When certain criteria are met, clients can achieve data security equal to or better than what they can achieve onsite.

Introduction

The public cloud offers many benefits to enterprises seeking a competitive advantage. Cost savings, elasticity, and scalability are driving more enterprises toward the public cloud. A recent Gartner report predicts that the public cloud services market will exceed $244 billion by 2017.

Until recently, public clouds were used mostly for functions that are not mission-critical, such as testing and development. Despite some successful cases of using the public cloud for actual production workloads (for example, Netflix's use of Amazon Web Services), most enterprises are still wary of entrusting their sensitive data and mission-critical functions to a third party due to security and compliance concerns.

Security and compliance concerns are legitimate issues. But, we've found that these concerns are partly due to lack of clarity and best practices regarding cloud security. This is not surprising, as there are different cloud computing models, and different expectations and requirements for each model. And not all cloud service providers are created equal. Some providers initially focused primarily on scalability, ease of use, and accessibility, then bolted on security. On the other hand, other providers started out with built-in security and a strategic differentiator.

In this article, learn about data security in public clouds — specifically for data stored within the dashDB data warehousing and analytics environment. The article also discusses:

- Cloud computing models offered by today's cloud service providers
- Cloud security
- A dashDB use case

Cloud computing models

A discussion of cloud security requires an understanding of the cloud computing models because there are different expectations and requirements of each. In its special publication SP 800-145, the U.S. National Institute of Standards and Technology (NIST) defines three cloud computing models, as follows:

Software as a Service (SaaS)
: The consumer uses the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface, such as a web browser (for example, web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure that includes network, servers, operating systems, storage, or even individual application capabilities — with the possible exception of limited user-specific application configuration settings. Examples of SaaS providers include Salesforce.com and Google.

Platform as a Service (PaaS)
: The consumer can deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment. Examples of PaaS providers include IBM Bluemix™ and Microsoft® Windows® Azure.

Infrastructure as a Service (IaaS)
: The consumer has the capability of provision processing, storage, networks, and other fundamental computing resources where they can deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure, but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (for example, host firewalls). Examples of IaaS providers include IBM SoftLayer and Amazon Web Services (AWS).

Figure 1 shows what is managed by the cloud service provider and what is managed by the client in the three cloud computing models.

Figure 1. Cloud computing models

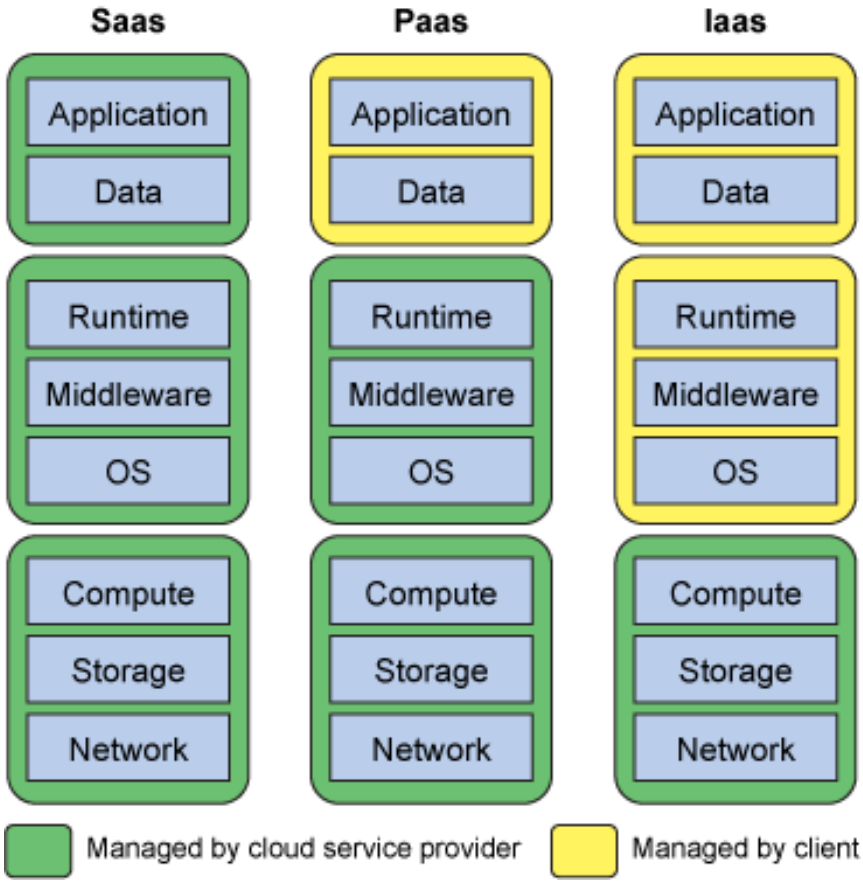

Demystifying cloud security

Adopting any of the cloud computing models involves concerns that span three security domains:

- *Identity,* which encompasses user identification, access management, and entitlements.
- *Protection,* which involves the protection and availability of infrastructure, data, and applications.
- *Insight,* which refers to gaining insight into user activities, threat intelligence, and compliance.

Client security objectives

The primary security objectives of clients typically reflect their responsibilities when adopting the cloud. Table 1 summarizes these objectives for each cloud computing model and provides a sample of security capabilities required to meet the objectives.

Table 1. Client security objectives

Cloud computing model	Organization or buyer	Primary security objectives	Sample security capabilities required
SaaS	CxO (CIO, CMO, CHRO)	• Governance of user access to SaaS applications • Sensitive data protection	• Identity federation and single sign-on • Data tokenization

		• Complete visibility into enterprise SaaS usage	• Monitoring and compliance reporting
PaaS	Application teams and LOBs	• Enable developers to compose secure cloud applications and services • Protection against application and data threats • Complete visibility into application and data services usage	• Authentication and authorization APIs • Database encryption • Monitoring and compliance reporting
IaaS	CIO and IT teams	• Protect the cloud infrastructure to securely deploy workloads and meet compliance objectives • Have full operational visibility across cloud deployments, and govern usage	• Privileged user management • Network intrusion prevention systems and firewalls • Monitoring and compliance reporting

Sharing responsibilities for security

The most important thing to recognize with cloud security is that the overall responsibility for any client workload deployed on a public cloud is *shared* between the client and the cloud service provider. It is critical for clients to understand the division of responsibility for security before moving their workloads to a public cloud. It is equally critical for clients to ask the cloud service provider to show certifications and audit reports that verify their posture regarding their share of the security responsibility.

Responsibilities of the cloud service provider

The responsibilities of the cloud service provider start with physical and environmental security. After all, the cloud service provider is operating a set of data centers. The key things to look for are physical employee access, fire detection and suppression, electrical power continuity, climate and temperature control for servers and other hardware devices, and sanitization for decommissioned storage devices. Figure 2 shows typical cloud service provider responsibilities for security.

Figure 2. Cloud service provider responsibilities for security

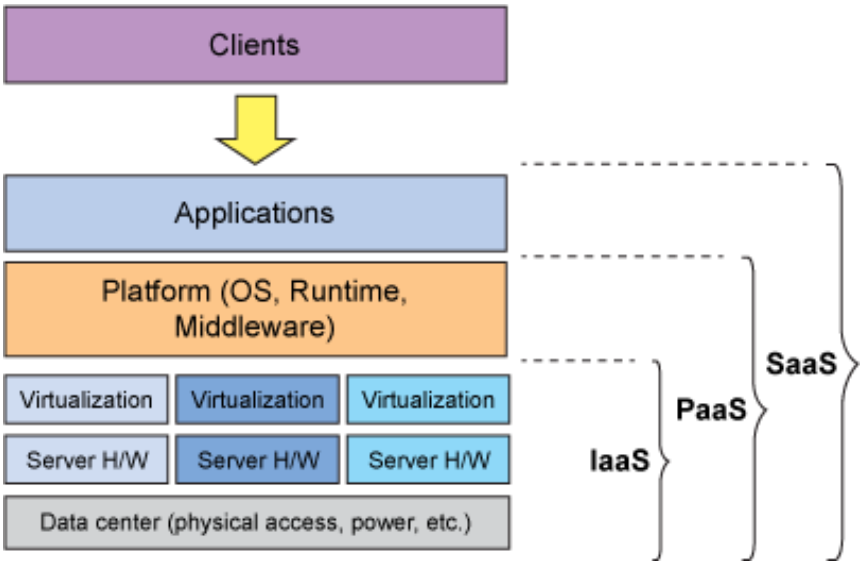

The next level of responsibility is network security, which includes firewalls and other network security devices to monitor and control communications at the network's external boundaries and at strategic internal boundaries within the network. This is meant to protect against traditional network security issues such as Distributed Denial of Service (DDoS) attacks, unauthorized port scanning, packet sniffing, and IP spoofing. Network security also includes the protection of data transmission between the client organization and the cloud service provider. Examples include uploading data to an object store such as SWIFT on IBM SoftLayer or AWS Simple Storage Service (S3). Secure Socket Layer (SSL) is typically used in this situation to protect against eavesdropping and tampering.

The next levels of responsibilities depend on the cloud computing model adopted by the client. For example, if the client adopted an IaaS model, the cloud service provider's responsibilities would stop at the hypervisor level. This is the model adopted by two leading public cloud providers: SoftLayer and AWS. If the customer adopted a PaaS model, the cloud service provider's responsibilities would extend to additional security domains, such as identity and access management, data security, and vulnerability management. If the customer adopted a SaaS model, the cloud service provider's responsibilities would further extend to include web application vulnerability testing and remediation.

To convince clients to adopt cloud services with confidence, the cloud service provider must get security certifications and share audit reports that attest to the security of their cloud services. Essentially, the cloud service provider must demonstrate that the cloud IT infrastructure is designed and managed with security best practices and industry standards. The best way to prove this is through certifications such as ISO 27001, Federal Information Security Management Act (FISMA), Federal Risk and Authorization Management Program (FedRamp), Payment Card Industry Data Security Standard (PCI DSS), SOC, Cloud Security Alliance (CSA), Service Organization Control reports (SOC2, SOC3), and Safe Harbor. SoftLayer and AWS publish their security statements at SoftLayer Cloud Security and AWS Security Center, respectively.

Responsibilities of the client

The client's first responsibility is to understand and qualify the risk profile of the workloads they intend to move to the cloud. Due diligence is expected when selecting a cloud service provider who will demonstrate that their cloud services are designed and managed in alignment with security best practices and industry standards. This includes inquiring about high availability of the IT infrastructure and the location where the client's data will actually be stored. For example, if the client's data must not leave a geographical location, or must not reside in a particular geographical location, the client must get assurance that their data will be stored according to their geographical requirements.

The client should consider the responsibilities for the technical controls to meet the security, privacy, and compliance requirements. These responsibilities start where the cloud service provider's responsibilities end. They also depend on the cloud service model being used. For example, if the client adopted an IaaS model, they are fully responsible for anything they deploy on the infrastructure they provisioned. This includes managing user access to the workloads and system administration; the protection of the workloads such as OS hardening, database hardening, storage encryption, and host firewall; and monitoring and compliance reporting. On the other hand, if the client adopted a PaaS model, he needs to ensure that the applications he's building are secure. This includes managing user access to the applications, application security testing and data encryption, and monitoring and compliance reporting. If the customer adopted a SaaS model, responsibilities would include managing user access to the SaaS application, data tokenization, and monitoring and compliance.

To meet security, privacy, and compliance requirements, clients could employ technical controls that they bring onto the cloud or they could consume those services from the cloud. For example, a client using an IaaS model could choose to harden a Linux® OS deployed on the provisioned infrastructure by using IP tables or by leveraging a solution offered by the cloud service provider, such as the Security Groups concept on AWS. Similarly, the client could choose to bring and deploy a virtual appliance such as Guardium® Database Activity Monitoring to protect the database deployed on the provisioned infrastructure. Another example is a PaaS client who takes advantage of application security scanning services from the cloud to scan the application for security vulnerabilities.

dashDB: A use case

dashDB is a data warehousing and analytic environment offered exclusively on IBM BlueMix

Learn more about dashDB: Data Warehousing and Analytics for Everyone.

System architecture

Figure 3 shows the system architecture for dashDB, which is comprised of the following components:

- DB2® BLU — The IBM DB2 database system with the BLU Acceleration feature.
- Web Console — The system's overall administrative console.
- LDAP Server — An embedded OpenLDAP system for managing users and groups.
- Tools — A set of downloadable tools including drivers, InfoSphere® Data Architect (IDA), IBM Data Studio, and Cognos Framework Manager (FM).
- Warehouse Packs — Physical data models and sample reports for accelerating data warehousing projects. The included packs are Customer Insight, Supply Chain, and Market and Campaign.

Figure 3. dashDB system architecture

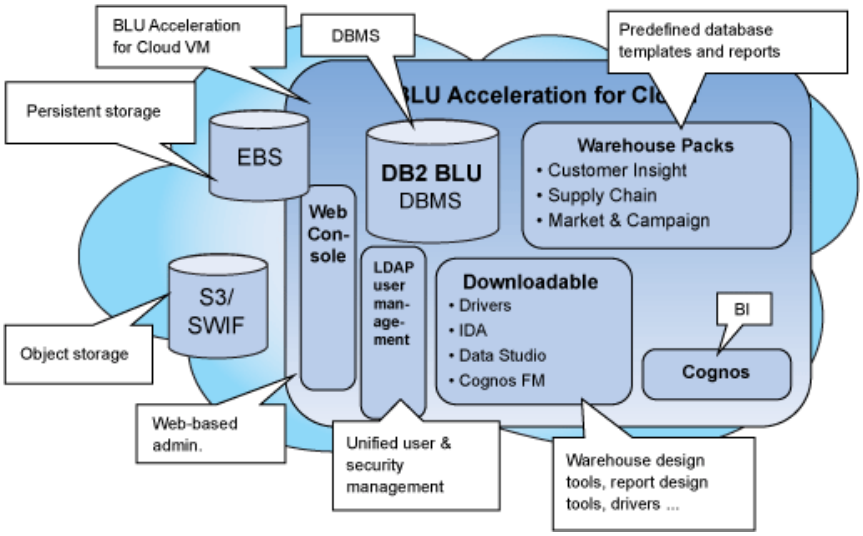

Using dashDB

A client uses dashDB as follows:

1. The client connects to the dashDB database and creates their warehouse schemas.

2. The client uploads data to his dashDB database.
3. The client connects his applications to the dashDB database and starts exploring just as he would with an onsite database. The client can also log on to the administrative console to add users, assign roles, or perform other administrative tasks.

Built-in security capabilities

dashDB provides a rich set of built-in security capabilities to help clients meet security, privacy, and compliance needs. The built-in security capabilities include:

User management
dashDB users are managed in the embedded LDAP server. The internal system components, such as DB2, Cognos, and the Web Console, are configured to perform user authentication through the embedded LDAP server. SSO across these components is also supported. User management is performed through the Web Console interfaces.

Role-based access control
When a user is created, they are assigned a specific role determining the level of access to the system: Administrator, Developer, or User.

Row-level access control
This allows clients to enforce stronger security policies by limiting the set of rows to which a user has access in a given table. For example, if a table contains employee data, a client can easily set up a rule that limits an employee's access to their own data or to employees who report to them.

Dynamic data masking
This allows clients to enforce stronger security policies by limiting access to sensitive columns in a given table. For example, if a table contains a Social Security number (SSN) column, a client can easily set up a rule whereby when that column is accessed by an unauthorized user, a masked value is returned instead of the actual SSN value.

Trusted contexts
This allows clients to further restrict when a user can exercise a particular privilege. For example, a client can easily implement a rule that permits connecting to the database only from a given IP address. For three-tier applications, trusted contexts allows the mid-tier application to assert the end-user identity to the database for access control and auditing purposes.

Encryption for data at rest
When provisioning a dashDB system, the client has the option to indicate whether to encrypt the database. The default is an encrypted database. The encryption uses Advanced Encryption Standard (AES) in Cipher Block Chaining (CBC) mode with a 256-bit key. Encryption and key management are totally transparent to applications and schemas. Upon provisioning, the client has the option to indicate the master key rotation period. The default is 90 days, but the client may choose a different value. The master key rotation is automatic and transparent. Database and tablespace backup images are automatically compressed and encrypted. As for online data, backup images are also encrypted using AES in CBC mode with 256-bit keys. Data is compressed first and then encrypted. This is particularly important because if the order is reversed, the compression ratio will not be interesting since encryption, by definition, removes any patterns.

Encryption for data in transit
SSL is supported for safeguarding both the database traffic, as well as the Web Console traffic.

Auditing
This allows clients to implement audit policies to hold users accountable for their actions and to track any malicious activities.

Strong security and privacy require protection at every level of the stack. Built-in capabilities include:

Linux hardening
dashDB employs a host firewall to protect listening services against port scans and other network security threats. As such, only the required TCP ports are open — those for the DB2 instance, the Web Console, the Cognos BI Web Console, and Secure Shell (SSH). dashDB takes advantage of the IP tables concept to implement the host firewall.

DB2 hardening
The DB2 database is automatically hardened upon provisioning. CONNECT authority to the database is revoked from PUBLIC, and SELECT privilege on the catalog tables and views is also revoked from PUBLIC. The AUTHENTICATION database manager configuration parameter is set to SERVER_ENCRYPT, which means that user authentication credentials are never flown in clear text between a user application and the database server. These credentials are automatically encrypted with AES 256 when flown over the network regardless of whether SSL is used or not.

Web Console restrictive interfaces
Non-administrative users have no way of invoking administrative functions because, by design, the Web Console does not show administrative features to users who are not members of the administrative role.

Guardium Database Activity Monitoring

Though the built-in security capabilities discussed above are sufficient for many clients, you can also add your own data security solution. For example, Guardium Database Activity Monitoring allows you to take security to the next level. With Guardium, clients can meet even the most stringent of security requirements as demonstrated by the success of Guardium in protecting mission-critical databases for many large enterprises around the world. Some of the key Guardium capabilities include:

Separation of duties
The Guardium server can be deployed on a separate host with a separate administrator, enabling physical separation of duties between security administration and database administration. The Guardium server can be deployed on the cloud or even onsite if so desired.

Real-time monitoring and alerting
Guardium continuously monitors all traffic in and out of the database and takes action in real time per the security policy. This is critical to limiting security exposures by immediately detecting intrusions and misuse. For instance, an excessive number of logins might be an indication that someone is trying to brute force the database. Guardium security policies support three types of rules:

- Access rules, which apply to the client's database request. Depending on the request, the action could be to let the request through, audit, send an alert, or even block that request.
- Extrusion rules, which apply to the database's response to a client request. Depending on the response, the action could be to let the response through, audit, alert, or even dynamically mask sensitive fields within that request.
- Exception rules, which apply when a database exception occurs, such as when a failed login threshold is exceeded. Again, Guardium allows great flexibility in choosing what action to take when such an event occurs. For example, an alert could be sent to an administrator via e-mail to let them know of this event.

Sensitive data discovery

Unknown sensitive data is sometimes the reason behind a data security breach. Guardium is able to detect sensitive data that may reside in the database so policies can be put around the data to limit security exposures.

Vulnerability assessment

Even an initially hardened database may go out of compliance if administrators are not careful with the changes they make. Guardium is able to assess the configuration of the database and produce a report describing the security posture of that database, giving the administrator an opportunity to remedy any out of compliance issues.

Optim Data Masking user-defined functions

The Optim™ Data Privacy solution has helped many customers move data from production to test environments without sacrificing data security or the applications' legitimate need for meaningful test data. This is accomplished through context-aware masking algorithms. When masking a VISA credit card number, the masked value, although fake, still looks like a VISA number. Thus, an application that validates credit card numbers does not break in the test environment.

The Optim Data Masking algorithms are now available as user-defined functions (UDFs). They can be used, for example, when extracting data for testing purposes on another database. Because these algorithms are SQL UDF, you can use them in conjunction with built-in dynamic data masking capability. This allows the client to combine the benefits of both worlds by achieving stronger security, application transparency, and state-of-the-art masking technology.

Secure design principles

The development of dashDB followed secure development best practices as outlined in the IBM Secure Engineering Framework. This includes the completion of a risk assessment and a threat modeling document. The IBM Security AppScan tools are regularly used to conduct static and dynamic code analysis during the development process. In addition, the Guardium Vulnerability Assessment tool is used to validate that the dashDB database is properly hardened.

Conclusion

This tutorial has discussed what you need to know about cloud security so you can adopt the cloud with confidence. Clients need to understand and qualify the risk profile of the workloads they intend to move to the cloud. It is also critical to understand that cloud security is a shared

responsibility between the cloud service provider and the client. Clients must understand the division of responsibility for security before moving their workloads to a public cloud.

Not all cloud service providers are created equal. It is critical for a client to select a cloud service provider that can demonstrate that their cloud services are designed and managed in alignment with security best practices and industry standards.

By adhering to the recommendations here, along with negotiating a satisfactory Service-Level Agreement (SLA), clients can achieve data security equal to or better than what they can achieve onsite. For example, leading cloud service providers manage their infrastructures with a high degree of automation compared to what most clients do onsite. Automation minimizes the risks for a misconfiguration due to manual intervention. Statistics show that misconfiguration is responsible for more than 60 percent of security breaches. Availability, a basic tenet of security, is another example. A small business might not be able to offer multiple geographically dispersed disaster recovery (DR) sites. With the right cloud service provider, that small business has better chances for business continuity.

Leading cloud service providers implement continuous security monitoring, which many clients do not implement onsite. Continuous security monitoring provides visibility into threats and prevents situations where vulnerabilities get so far out of line over time that they pose significant risks. With the right cloud service provider, clients of all sizes can benefit from continuous security monitoring.

Related topics

- For information about cloud computing and computer security, read:
  - The NIST Definition of Cloud Computing, Special Publication SP 800-145
  - The NIST Guidelines for Media Sanitization, Special Publication SP 800-88
- Visit IBM Secure Engineering: Developing products and services with security in mind to explore product security vulnerabilities and security information.
- Download and read Data Security Best Practices: A practical guide to implementing row and column access control (R. Walid, 2012).
- Learn more about:
  - InfoSphere Guardium Data Security
  - InfoSphere Optim Data Privacy
  - IBM SoftLayer Security
  - Amazon Web Services Security
- Read Forecast: Public Cloud Services, Worldwide, 2011-2017, 3Q13 Update by Gartner Inc. in 2013, to learn more about strong growth in public cloud services.