LANDMARK
ASSOCIATES INC

Data Security

Landmark Associates Inc.

Overview

Data privacy and security are of the utmost importance here at Landmark Associates, hence we utilize the resources of Salesforce and Amazon Web Services, in tandem, to support the data security needs of our stakeholders.

Our system was designed based on the best practices of other cloud-based technologies and offers our clients a secure environment for their data.

The goal of this document is to provide you with an in-depth look of how Landmark Associates manages, stores and secures your data.

Multi-Layer Data Encryption from In Motion to At Rest

All files transmitted or stored in our system are protected through multiple layers of encryption. For data that is at rest, our system employs a strong multi-factor encryption through Amazon's S3-managed encryption keys (SSE-S3). With SSE, every each file is encrypted with a unique key, which is also encrypted by a separate master key which rotates regularly.The encrypted data, encryption keys and master keys are stored and secured on separate hosts for additional layers of protection.

- When a file is stored on our system, the file is uploaded to our AWS S3 bucket via HTTPS from our Salesforce-managed client portal. Then, by using Amazon's Lambda function we encrypt the data using a 256-bit Advanced Encryption Standard (AES-256). This is one of the strongest block ciphers available to enable the server side encryption setting as described above.
- For data that is in motion our site uses HyperText Transfer Protocol with Secure SocketsLayer or HTTPS, the protocol which supports secure transmission of encrypted data over the web.

The Map of Landmark's System

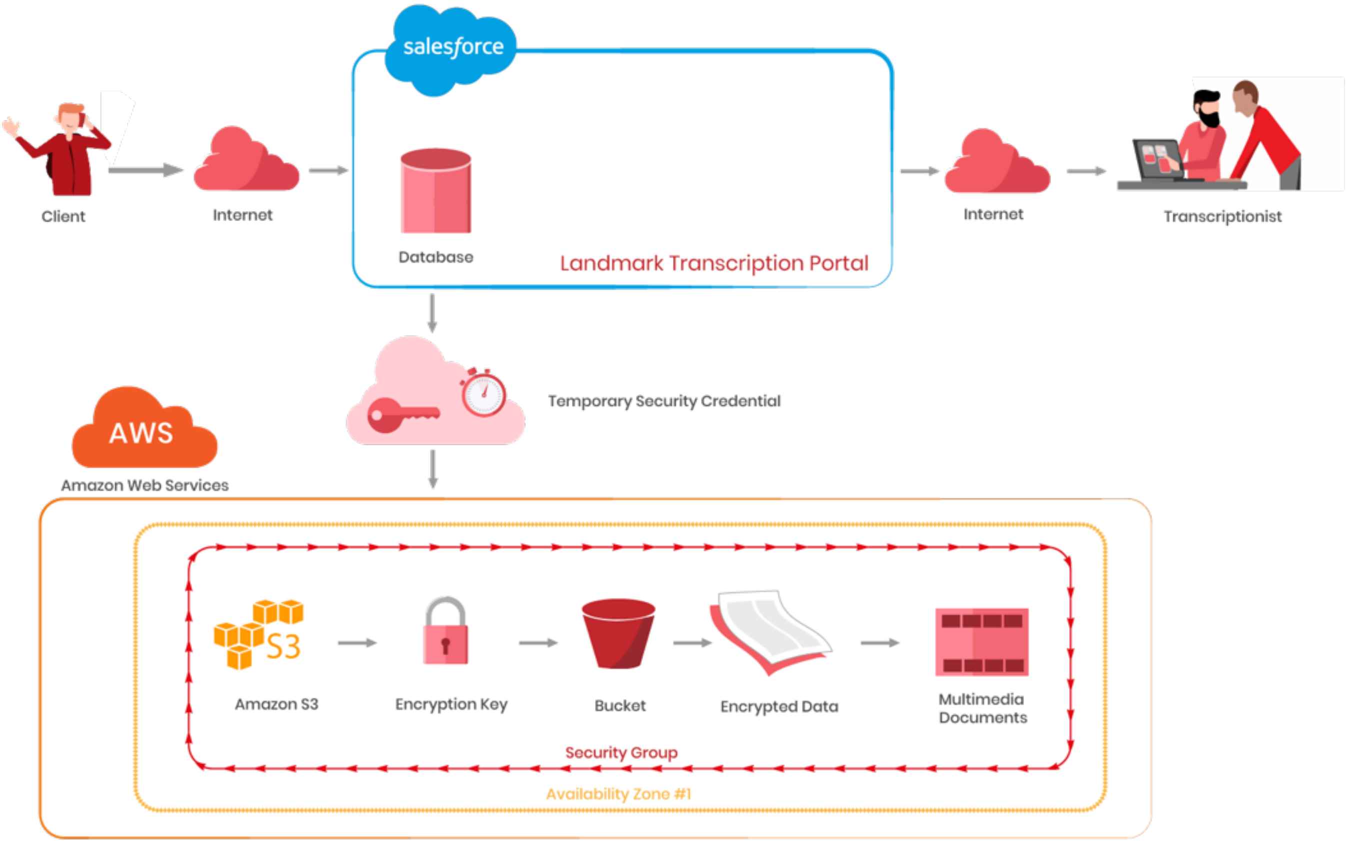

Security Policies and Best Practices

In response to the evolving threat landscape of cloud-based technology, Landmark stays up-to-date on the best practices and the latest security features to further support the existing system. For instance, Landmark internally implements two-factor authentication for our administrative users, which adds an additional safeguard in the login process. Login attempts that do not have the valid credentials from both sources are not granted access to our system. In addition, Landmark implements IP restrictions, preventing those who try to login to our system outside the designated IP address will not be granted access.

Each employee, executive staff member and contractor must sign a non-disclosure and confidentiality agreement, and follow Landmark's security policies. We also accommodate additional confidentiality requirements for our clients, including project-specific non-disclosure agreements or confidentiality pledges.

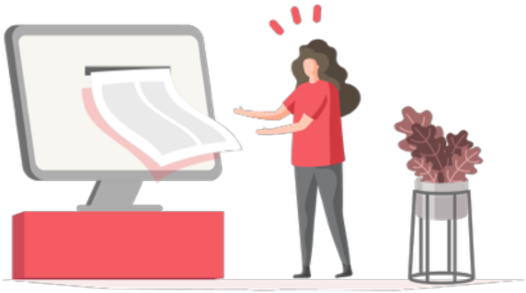

Salesforce CRM and Platform Security

Introduction

This paper is an overview of Salesforce CRM and Force.com platform security, including infrastructure, operations, application security features and controls, and the Trust organization at salesforce.com.

Principles of Trust

Our vision is to be the trusted cloud CRM and platform provider, based on the values of maintaining the confidentiality, integrity, and availability of our customers' data. Our methods to fulfill this vision are built upon an executive commitment to ensure and improve the security of our services, and include:

- Defense-in-depth: whenever possible, multiple controls and technologies are applied to limit the possibility of any single point of failure
- Investment: in personnel, tools, and technologies to manage, analyze, and improve our security effectiveness.
- Transparency: trust cannot be maintained without open communications regarding our service performance, reliability, and security, and to that end we strive to be industry leaders in transparency. See trust.salesforce.com and the Developer Security pages.

Application Security

Salesforce.com addresses application security by combining the strengths of multi-tenancy with modern development and management processes to minimize security vulnerabilities and maximize performance and usability. This section describes the software development lifecycle methodology and some of the core features that secure Salesforce CRM and the force.com platform.

Multi-Tenant Security Highlights

To achieve high scalability and performance, the database behind the salesforce.com CRM products is a single instance shared by thousands of customers. Our application ensures that users see only the data to which they have assigned privileges:

- Every record of the database contains the customer's orgID
- During login, the authenticated user is mapped to their org and access privileges according to the sharing model
- Every request to the database is formed by the application and is limited to the user's orgID and privileges
- Every row returned from the database is then validated against the orgID
- An error in the query process does not return any data to the client

For a more in-depth discussion of the multi-tenant design, please refer to the salesforce.com Multi-tenancy whitepaper (see Resources, page 9)

Software Development Lifecycle (SDLC)

The SDLC used by salesforce.com incorporates security as a core consideration. Before a product can be considered "done," it must meet security requirements as well as functional requirements. To ensure high-quality code, security is part of each of the design, code, test, and release phases of the SDLC.

- Training: all developers must take our internal secure coding training promptly after employment at salesforce.com.

Secure coding and testing standards are also provided to supplement the live training.

- Design: security principles are an integral part of the design phase. The salesforce.com Product Security team assists with new feature design and code review, establishes coding standards, and provides secure coding training to our developers.
- Coding: a secure development framework to prevent introduction of vulnerabilities. All code is peer-reviewed before check in, and high-risk features go through a threat-modeling exercise to further anticipate and reduce risks.
- Test: salesforce.com uses both automated and manual techniques to validate code. All features are tested with automated tools including static and dynamic analysis tools, automated web application scanners, and fuzzers to test for input validation flaws. Critical features receive additional investigation and manual penetration testing from the product security team.
- Release: When all the phases of development have been completed and passed the security requirements, the feature is “done” from a security standpoint

Infrastructure

Security by design starts with a secure physical infrastructure. Salesforce is hosted from dedicated spaces in top-tier data center facilities. The appearance of our data centers are low-profile and designed as anonymous buildings without any company signage. The exterior walls of the facilities are bullet resistant. Concrete bollards are positioned around the facility perimeter to aid in providing further security protection. All facilities maintain multiple transit access routes and are within close proximity to local law enforcement and fire/emergency services. All data centers selected are at core internet hubs with diverse physically protected routes into the facility.

All data center facilities maintain the following physical security controls:

- Continuous, 24x7x365 physical security presence
- All facility visitors are required to sign-in at the security desk to validate their identity to ensure only authorized persons are granted physical access.
- Enforcement of mantrap to reduce tailgating
- Emergency exit doors alarmed and integrated into the security access control system Closed circuit television (CCTV) cameras capture & record all motion throughout the facility (internally/externally)
- Video retention available for up to 92 days for the salesforce.com computer room
- All physical access activity is electronically logged into security access management system available to salesforce.com for further review.
- Multiple layers of physical security including multi-factor biometric challenge in order to reach the salesforce.com computer room

In addition to securing the data center locations it is critical that all facilities maintain robust critical infrastructure to support salesforce.com through the following services:

Power

- Next generation UPS systems (N+1)
- No single point of failure from cabinet circuit to UPS
- Minimum 48 hours fuel supply onsite
- Standing contracts with fuel suppliers

Cooling

- N+1 Cooling Infrastructure
- Industry best practices for efficient cooling with hot aisle/cold aisle configuration
- Overhead cabling allows better under floor cooling to our equipment
- Multiple sources of water: Industrial, onsite backup water supply
- Leak detection underneath raised floor

Fire Detection and Suppression

- Minimum (1) certified functional fire extinguisher per computer room
- VESDA (very early smoke detection apparatus) deployed to detect combustion prior to visible smoke
- Pre-action dual interlocking dry pipe separated by zones

Network

- Multi-gigabit IP transit for external customer service
- Access to thousands of global Internet peering points
- Private peering with key carriers
- Diverse physically protected secure paths into facilities for redundancy

All salesforce.com computing infrastructure in our data center facilities is managed by our full time employees. All infrastructure is redundant and fault tolerant across components, including network, application servers, and database servers.

Host Security

The Salesforce CRM suite of applications is powered entirely b

y Linux and Solaris systems, built with an automated process that ensures compliance to standardized build specifications, including removal of unnecessary processes, accounts, and protocols and use of non-root accounts to run services.

Monitoring and validation of host security includes:

- File integrity monitoring for unexpected changes to system configuration
- Malicious software detection on application servers
- Vulnerability detection and remediation, including internal and external scanning and patching
- All host logs are forwarded to a centralized log aggregation and event correlation server for review and alerting

Network Security

Access to salesforce.com is via the public Internet, and connections are secured via SSL/TLS. We contract with multiple carriers to provide the connectivity and bandwidth to host business critical data.

Figure 1 shows a high-level diagram of the Salesforce infrastructure including a “pod” of Salesforce servers. (A pod is a physical instance of the Salesforce CRM application, comprised of a pool of application servers and a clustered database host, indicated by the dashed line.) Protecting the perimeter of the environment are edge routers and stateful firewalls. The allowed network traffic passes through the perimeter firewalls and reaches redundant pairs of load balancers that also terminate SSL connections.

The load balancers make connections through core switches to reach the pool of application servers. Finally, the application servers run the Salesforce CRM application and connect to the database through another tier of firewalls and to other resources.

The application environment is completely secured from the Internet and only required services are allowed. Internal traffic is routed on a private RFC 1918 network, with network address translation (NAT) to public IP addresses.

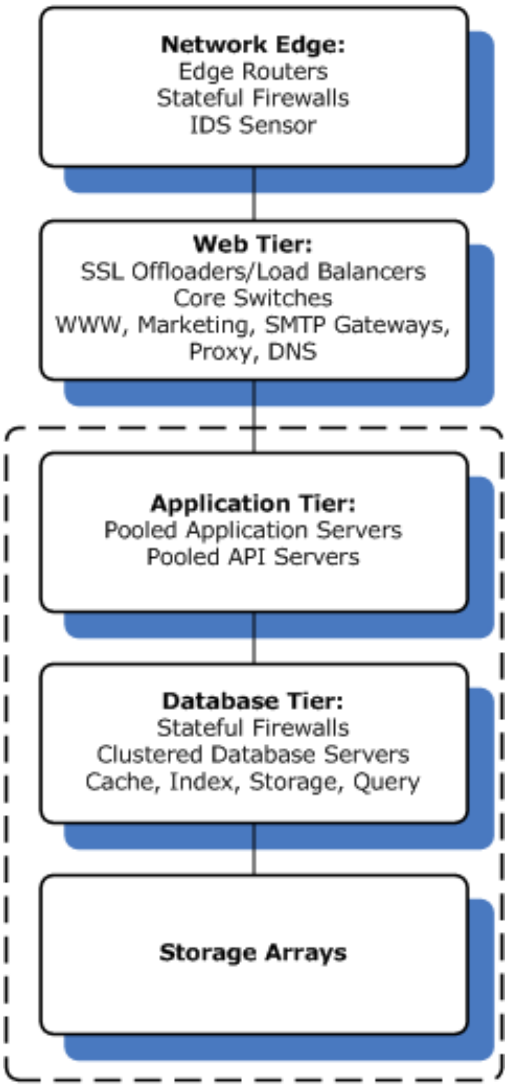

Figure 1. Pod Infrastructure

Transmission Layer Security

Connections to Salesforce are served over SSLv3 or TLS 1.0 (HTTPS), using 128-bit Global Step-up Certificates from reputable Certificate Authorities (CA) such as VeriSign, CyberTrust, and Verizon. This allows clients to force preferred ciphers if necessary. Email relay connections can utilize the TLS protocol, allowing our customers to create secure connections to their own servers.

Database Security

The database is hardened according to industry and vendor guidelines, and is accessible only by a limited number of salesforce.com employees with DBA access. Customers do not have direct database or OS level access to the Salesforce environment. Customer passwords for Salesforce CRM are hashed via SHA 256 before being stored in the database.

Customers can specify that certain field types use encryption. These custom fields are encrypted by the application before being saved to the database and can be configured to mask the display of their contents according to user access.

Disaster Recovery and Backups

Customer data are mirrored, backed up locally, and also mirrored over an encrypted network (AES 128) to a 100% full-scale replica disaster recovery data center. This is made possible through the following services:

- Software multipathing to ensure availability of reaching enterprise class storage through redundant storage switching architecture without a single point of failure
- Enterprise storage arrays configured with hardware RAID to protect against disk failure
- Salesforce.com completes a minimum of 1 annual DR exercise for compliance purposes.
- Validated disaster recovery plan with objectives as follows:

  o RPO: 4hrs

  o RTO: 12hrs
- All new "pod" infrastructure undergoes a DR exercise as part of the validation process prior to go live.

Backup Media Security

All backup tape media are maintained by salesforce.com employees at our secure data center facilities to ensure proper chain of custody. Electronic logs record and retain all physical access attempts to reach backup media. Salesforce.com employees follow industry guidelines for the secure destruction of backup media that has become corrupted or reached end of life.

Operations and Staffing

Our operational practices at salesforce.com ensure data security from policies, procedures, and monitoring through audits and continuous improvement.

ISO 27001 Certified Information Security Management System

Our information security management system follows ISO 27002 practices and is certified to the ISO 27001 standard. We maintain a broad registration, including all our customer data centers and large office sites. This means that our security practices are required by policy, managed and tracked in implementation, and improved as necessary. These practices include:

- Background checks for all employees (the salesforce.com environment is managed by full-time employees), including criminal history, job history and educational credentials.
- Security awareness education for all personnel, including first-day orientation, annual certification, and regular internal communications to train and remind our staff on such topics as notification requirements and threat recognition.
- Requirements for access management, code development and release, encryption, change management, and disciplinary action.

Administration Procedures

Administrator access to the production environment is highly restricted. All administrators must use a secure desktop client and authenticate first to a secure server using two-factor authentication. The secure server hosts all management applications; the desktop client is a 'sandbox' environment that presents only pixel data and prohibits copy/paste and writing to local storage. Management connections are authenticated using two factors again at a bastion host, and then connections can be made via SSH with Kerberos authentication to the production systems. Access to each server is restricted to the least number of administrators possible. Outbound Internet traffic from the secure server is limited via proxy to appropriate support sites.

Site Operations

The operations team consists of full-time mid and senior-level analysts who manage the Salesforce production environment through a 24x7x365 follow-the-sun business model. Operations centers are located in the United States, Europe, and Asia.

The team manages tier 1 and 2 network, systems, database and storage related alerts and issues. Their management responsibilities include but are not limited to monitoring, incident response, and change management of the Salesforce production environment.

Computer Security Incident Response Team

The Computer Security Incident Response Team (CSIRT) runs in parallel with site operations to provide monitoring and incident response. The CSIRT consists of senior level security analysts and manages a variety of tools and third-party resources that include:

- Intrusion Detection Systems (IDS): monitor every network in the production environment for potentially malicious network traffic.
- Security Event Management (SEM): Activity logs from all production devices and server are sent in real-time to a SEM that correlates, reports, and alerts on events such as successful/failed logins, SU changes, and system messages.
- Threat Monitoring: the salesforce.com information security team receives and reviews threat alerts from a variety of sources including SANS, CERT, OWASP, and manufacturers of our equipment and software. Threats that are deemed critical are escalated for the appropriate response.
- Perimeter monitoring: third-party security firms provide periodic vulnerability scanning and continuous perimeter monitoring to detect changes in IP address or ports opened, service versions, and SSL certificate expirations.
- External Certificate Authority monitors certificate validity and renewal

Product Security Features

In addition to all of the steps salesforce.com takes to ensure the security and privacy of our customers’ data, the Salesforce CRM product also incorporates many features to help customers secure their data. This section provides a high-level overview of some of the more important features related to security. For a more detailed description, please see the salesforce.com Security Implementation Guide.

Sharing

Salesforce CRM provides a sophisticated model known as “sharing” to control the users’ access to data. For each object type, defaults can apply across the entire company or for individual profiles to allow or restrict access to records that users do not own. In addition, rules can be created to allow various access levels (read/write/modify) based on organizational charts, sales territories, or other custom-defined groups of users. Sharing rules should be carefully monitored and reviewed to ensure sensitive items are restricted appropriately.

Password Policies

Password complexity and expiration settings within Salesforce CRM should be configured to comply with your internal policies. Note that the default settings are not assumed to fit all use cases and may not be appropriate for companies with stronger security policies (see Single Sign-on Options)

The available password settings include:

- Password expiration timers
- Prevent re-use of previous passwords;
- Password complexity restrictions
- Invalid lockout attempts
- Lockout timers

Session Settings

Several settings can be used to place restrictions on active user sessions. These include configuring the idle session timeout, locking sessions to the IP address used at login, and requiring secure (HTTPS) connections. As a best practice, the customer should review the default settings and adjust them appropriately; for example, the default idle session timeout value is two hours, and while it may be appropriate for customers with lower privacy and security requirements, larger customers or those with regulatory requirements will probably need to set a shorter duration.

Login and Authentication Settings

By default, all users can log in to salesforce.com CRM from any IP address at any time of day, subject to the restrictions of the Identity Confirmation feature described below. You can restrict user login access to specific work hours and/or defined ranges of IP addresses. These restrictions are defined based on User Profiles.

Time-of-Day Restrictions

User logins can be restricted to specific times of the day. Different time-of-day restrictions can be defined for different types of users.

IP Address Restrictions

User logins can be restricted to specific IP addresses or ranges of IP addresses. IP range restrictions can be configured for the entire organization or for each particular class of user.

Single Sign-On Options

In addition to the standard username and password authentication, Salesforce CRM supports two types of single sign-on methods. To use your own account management system, salesforce.com recommends enabling one of the following options:

- Delegated Authentication – When delegated authentication is enabled, Salesforce CRM makes a Web services call to your organization to authenticate your users, rather than using the native salesforce.com CRM passwords.
- Federated Authentication – Federated authentication directs Salesforce CRM to use the Security Assertion Markup Language (SAML) for user authentication.

Identity Confirmation

The Identity Confirmation feature was developed in part to provide a defense against phishing attacks and stolen user credentials. This feature is enabled by default for all Customers.

If IP address restrictions are not used, Salesforce CRM checks for the presence of a cookie and queries an IP address history to determine whether the user's browser or current IP address were previously used to log in to Salesforce CRM.

If the browser has the cookie or is using a previously known IP address, the login proceeds. If the cookie is not present and the connection is coming from a new IP address, the user is directed to a special screen and prompted to click a Send Activation Link button, which sends an activation email to the address on record for the user's account; the requestor must be able to login to that account to click the activation link on the email.

Auditing Salesforce.com Security Features

Salesforce.com provides a free AppExchange utility called the "Security Health Check." The Security Health Check reviews and scores your org's salesforce.com security settings, then provides a dashboard output to show relative strength of settings. These scores can be saved to track improvements and provide internal auditing information.

Logging

Salesforce.com provides built-in logging features available to the org admin:

- Login history: a six-month history of all login attempts to the org, including username, IP address, success/failure, and time and date is available upon demand.
- Setup audit trail: a 180-day history of setup changes made by your org's administrators is also available upon demand, and can be used to troubleshoot and audit administrative activities.

The Trust Organization

Salesforce.com was a pioneer in the concept of customer and public trust, and developed our Trust organization to communicate our policies, practices, and performance; manage and improve our security; and to maintain leadership in the industry for awareness and transparency. The Trust Organization includes:

- Chief Trust Officer: leads the Trust Organization from the executive level and represents the Trust mission to both our customers and company.
- Product Security: builds and improves our application and platform security controls including code reviews and automated testing, developer secure coding practices and training, AppExchange partner reviews, internal assessment and testing of our infrastructure.
- Enterprise Security: salesforce.com internal corporate and end user security
- Awareness: salesforce.com personnel security awareness
- Technology Risk, Audit and Compliance: audits and reports on security control implementation and effectiveness and facilitates external audits; manages customer queries and on-site audits; performs risk assessments; and manages security policies.
- Privacy Counsel: participates in policy formation and reviews privacy laws and regulations to provide guidance to the Trust organization

Additional Security and Privacy Audits and Certifications

In addition to ISO 27001 certification, the Salesforce platform is audited against several standards:

- SSAE 16 SOC-1 (formerly SAS 70)
- AT-101 SOC-3 (SysTrust)
- PCI-DSS
- TRUSTe (including verification of US-EU Safe Harbor compliance)
- US Federal Information Security Management Act (FISMA) - Moderate Authority to Operate
- Tüv Geprüfter Datenschutz
- Japan Privacy Mark

Conclusion

Of course, one whitepaper can't address all security concerns, especially as they apply to the feature-rich and flexible Salesforce CRM application and Force.com platform. We invite you to review the additional resources and contact us if you have any questions regarding the security of our facilities, platform, and applications. Please note that salesforce.com is committed to continuous improvement, so this is information is subject to change.

Resources

trust.salesforce.com – our public pages on availability, performance, and security:

https://trust.salesforce.com/index.html

Developer Wiki Security Page – administration guides, developer tools, and much more:

http://wiki.developerforce.com/index.php/Security

Security Implementation Guide – administrator's setup guide for salesforce.com security features:

https://na1.salesforce.com/help/doc/en/salesforce_security_impl_guide.pdf

Security Health Check – free tool to rank your salesforce.com org's security feature settings:

http://appexchange.salesforce.com/listingDetail?listingId=a0N300000018mjUEAQ

Salesforce.com Multitenancy Whitepaper:

http://www.developerforce.com/media/ForcedotcomBookLibrary/
Force.com_Multitenancy_WP_101 508.pdf

Secure, private, and trustworthy: enterprise cloud computing with Force.com:

http://www.salesforce.com/assets/pdf/misc/WP_Forcedotcom-Security.pdf

Test your ability to identify potential fraudulent emails and websites:

http://security.force.com/phishingquiz

trust.salesforce.com

Amazon Web Services (AWS) and S3 Security

Overview of AWS Security

Storage Services

(Please consult http://aws.amazon.com/security/ for the latest version of this paper)

Notices

This document is provided for informational purposes only. It represents AWS' current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS' products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Storage Services

Amazon Web Services provides low-cost data storage with high durability and availability. AWS offers storage choices for backup, archiving, and disaster recovery, as well as block and object storage.

Amazon Simple Storage Service (Amazon S3) Security

Amazon Simple Storage Service (S3) allows you to upload and retrieve data at any time, from anywhere on the web. Amazon S3 stores data as objects within buckets. An object can be any kind of file: a text file, a photo, a video, etc. When you add a file to Amazon S3, you have the option of including metadata with the file and setting permissions to control access to the file. For each bucket, you can control access to the bucket (who can create, delete, and list objects in the bucket), view access logs for the bucket and its objects, and choose the geographical region where Amazon S3 will store the bucket and its contents.

Data Access

Access to data stored in Amazon S3 is restricted by default; only bucket and object owners have access to the Amazon S3 resources they create (note that a bucket/ object owner is the AWS Account owner, not the user who created the bucket/ object). There are multiple ways to control access to buckets and objects:

Identity and Access Management (IAM) Policies.

AWS IAM enables organizations with many employees to create and manage multiple users under a single AWS Account. IAM policies are attached to the users, enabling centralized control of permissions for users under your AWS Account to access buckets or objects. With IAM policies, you can only grant users within your own AWS account permission to access your Amazon S3 resources.

Access Control Lists (ACLs).

Within Amazon S3, you can use ACLs to give read or write access on buckets or objects to groups of users. With ACLs, you can only grant other AWS accounts (not specific users) access to your Amazon S3 resources.

Bucket Policies.

Bucket policies in Amazon S3 can be used to add or deny permissions across some or all of the objects within a single bucket. Policies can be attached to users, groups, or Amazon S3 buckets, enabling centralized management of permissions. With bucket policies, you can grant users within your AWS Account or other AWS Accounts access to your Amazon S3 resources.

Type of Access Control	AWS Account-Level Control?	User-Level Control?
IAM Policies	No	Yes
ACLs	Yes	No
Bucket Policies	Yes	Yes

You can further restrict access to specific resources based on certain conditions. For example, you can restrict access based on request time (Date Condition), whether the request was sent using SSL (Boolean Conditions), a requester’s IP address (IP Address Condition), or based on the requester's client application (String

Conditions). To identify these conditions, you use policy keys. For more information about action-specific policy keys available within Amazon S3, refer to the Amazon Simple Storage Service Developer Guide.

Amazon S3 also gives developers the option to use query string authentication, which allows them to share Amazon S3 objects through URLs that are valid for a predefined period of time. Query string authentication is useful for giving HTTP or browser access to resources that would normally require authentication. The signature in the query string secures the request.

Data Transfer

For maximum security, you can securely upload/download data to Amazon S3 via the SSL encrypted endpoints. The encrypted endpoints are accessible from both the Internet and from within Amazon EC2, so that data is transferred securely both within AWS and to and from sources outside of AWS.

Data Storage

Amazon S3 provides multiple options for protecting data at rest. For customers who prefer to manage their own encryption, they can use a client encryption library like the Amazon S3 Encryption Client to encrypt data before uploading to Amazon S3. Alternatively, you can use Amazon S3 Server Side Encryption (SSE) if you prefer to have Amazon S3 manage the encryption process for you. Data is encrypted with a key generated by AWS or with a key you supply, depending on your requirements. With Amazon S3 SSE, you can encrypt data on upload simply by adding an additional request header when writing the object. Decryption happens automatically when data is retrieved.

Note that metadata, which you can include with your object, is not encrypted. Therefore, AWS recommends that customers not place sensitive information in Amazon S3 metadata.

Amazon S3 SSE uses one of the strongest block ciphers available – 256-bit Advanced Encryption Standard (AES-256). With Amazon S3 SSE, every protected object is encrypted with a unique encryption key. This object key itself is then encrypted with a regularly rotated master key. Amazon S3 SSE provides additional security by storing the encrypted data and encryption keys in different hosts. Amazon S3 SSE also makes it possible for you to enforce encryption requirements. For example, you can create and apply bucket policies that require that only encrypted data can be uploaded to your buckets.

For long-term storage, you can automatically archive the contents of your Amazon S3 buckets to AWS' archival service called Amazon Glacier. You can have data transferred at specific intervals to Glacier by creating lifecycle rules in Amazon S3 that describe which objects you want to be archived to Glacier and when. As part of your data management strategy, you can also specify how long Amazon S3 should wait after the objects are put into Amazon S3 to delete them.

When an object is deleted from Amazon S3, removal of the mapping from the public name to the object starts immediately, and is generally processed across the distributed system within several seconds. Once the mapping is removed, there is no remote access to the deleted object. The underlying storage area is then reclaimed for use by the system.

Data Durability and Reliability

Amazon S3 is designed to provide 99.999999999% durability and 99.99% availability of objects over a given year. Objects are redundantly stored on multiple devices across multiple facilities in an Amazon S3 region. To help provide durability, Amazon S3 PUT and COPY operations synchronously store customer data across multiple facilities before returning SUCCESS. Once stored, Amazon S3 helps maintain the durability of the objects by quickly detecting and repairing any lost redundancy. Amazon S3 also regularly verifies the integrity of data stored using checksums. If corruption is detected, it is repaired using redundant data. In addition, Amazon S3 calculates checksums on all network traffic to detect corruption of data packets when storing or retrieving data.

Amazon S3 provides further protection via Versioning. You can use Versioning to preserve, retrieve, and restore every version of every object stored in an Amazon S3 bucket. With Versioning, you can easily recover from both unintended user actions and application failures. By default, requests will retrieve the most recently written version. Older versions of an object can be retrieved by specifying a version in the request. You can further protect versions using Amazon S3 Versioning's MFA Delete feature. Once enabled for an Amazon S3 bucket, each version deletion request must include the six-digit code and serial number from your multi-factor authentication device.

Access Logs

An Amazon S3 bucket can be configured to log access to the bucket and objects within it. The access log contains details about each access request including request type, the requested resource, the requestor's IP, and the time and date of the request. When logging is enabled for a bucket, log records are periodically aggregated into log files and delivered to the specified Amazon S3 bucket.

Cross-Origin Resource Sharing (CORS)

AWS customers who use Amazon S3 to host static web pages or store objects used by other web pages can load content securely by configuring an Amazon S3 bucket to explicitly enable cross-origin requests. Modern browsers use the Same Origin policy to block JavaScript or HTML5 from allowing requests to load content from another site or domain as a way to help ensure that malicious content is not loaded from a less reputable source (such as during cross-site scripting attacks). With the Cross-Origin Resource Sharing (CORS) policy enabled, assets such as web fonts and images stored in an Amazon S3 bucket can be safely referenced by external web pages, style sheets, and HTML5 applications.

Amazon Glacier Security

Like Amazon S3, the Amazon Glacier service provides low-cost, secure, and durable storage. But where Amazon S3 is designed for rapid retrieval, Amazon Glacier is meant to be used as an archival service for data that is not accessed often and for which retrieval times of several hours are suitable.

Data Upload

To transfer data into Amazon Glacier vaults, you can upload an archive in a single upload operation or a multipart operation. In a single upload operation, you can upload archives up to 4 GB in size. However, customers can achieve better results using the Multipart Upload API to upload archives greater than 100 MB. Using the Multipart Upload API allows you to upload large archives, up to about 40 TB. The Multipart Upload API call is designed to improve the upload experience for larger archives; it enables the parts to be uploaded independently, in any order, and in parallel. If a multipart upload fails, you only need to upload the failed part again and not the entire archive.

When you upload data to Amazon Glacier, you must compute and supply a tree hash. Amazon Glacier checks the hash against the data to help ensure that it has not been altered en route. A tree hash is generated by computing a hash for each megabyte-sized segment of the data, and then combining the hashes in tree fashion to represent ever-growing adjacent segments of the data.

As an alternate to using the Multipart Upload feature, customers with very large uploads to Amazon Glacier may consider using the AWS Import/Export service instead to transfer the data. AWS Import/Export facilitates moving large amounts of data into AWS using portable storage devices for transport. AWS transfers your data directly off of storage devices using Amazon's high-speed internal network, bypassing the Internet.

You can also set up Amazon S3 to transfer data at specific intervals to Amazon Glacier. You can create lifecycle rules in Amazon S3 that describe which objects you want to be archived to Amazon Glacier and when. You can also specify how long Amazon S3 should wait after the objects are put into Amazon S3 to delete them.

To achieve even greater security, you can securely upload/download data to Amazon Glacier via the SSL-encrypted endpoints. The encrypted endpoints are accessible from both the Internet and from within Amazon EC2, so that data is transferred securely both within AWS and to and from sources outside of AWS.

Data Retrieval

Retrieving archives from Amazon Glacier requires the initiation of a retrieval job, which is generally completed in 3 to 5 hours. You can then access the data via HTTP GET requests. The data will remain available to you for 24 hours.

You can retrieve an entire archive or several files from an archive. If you want to retrieve only a subset of an archive, you can use one retrieval request to specify the range of the archive that contains the files you are interested or you can initiate multiple retrieval requests, each with a range for one or more files. You can also limit the number of vault inventory items retrieved by filtering on an archive creation date range or by setting a maximum items limit. Whichever method you choose, when you retrieve portions of your archive, you can use the supplied checksum to help ensure the integrity of the files provided that the range that is retrieved is aligned with the tree hash of the overall archive.

Data Storage

Amazon Glacier automatically encrypts the data using AES-256 and stores it durably in an immutable form. Amazon Glacier is designed to provide average annual durability of 99.999999999% for an archive. It stores each archive in multiple facilities and multiple devices. Unlike traditional systems which can require laborious data verification and manual repair, Amazon Glacier performs regular, systematic data integrity checks and is built to be automatically self-healing.

Data Access

Only your account can access your data in Amazon Glacier. To control access to your data in Amazon Glacier, you can use AWS IAM to specify which users within your account have rights to operations on a given vault.

AWS Storage Gateway Security

The AWS Storage Gateway service connects your on-premises software appliance with cloud-based storage to provide seamless and secure integration between your IT environment and AWS' storage infrastructure. The service enables you to securely upload data to AWS' scalable, reliable, and secure Amazon S3 storage service for cost-effective backup and rapid disaster recovery.

AWS Storage Gateway transparently backs up data off-site to Amazon S3 in the form of Amazon EBS snapshots. Amazon S3 redundantly stores these snapshots on multiple devices across multiple facilities, detecting and repairing any lost redundancy. The Amazon EBS snapshot provides a point-in-time backup that can be restored on-premises or used to instantiate new Amazon EBS volumes. Data is stored within a single region that you specify.

AWS Storage Gateway offers three options:

- Gateway-Stored Volumes (where the cloud is backup). In this option, your volume data is stored locally and then pushed to Amazon S3, where it is stored in redundant, encrypted form, and made available in the form of Elastic Block Storage (EBS) snapshots. When you use this model, the on-premises storage is primary, delivering low-latency access to your entire dataset, and the cloud storage is the backup.
- Gateway-Cached Volumes (where the cloud is primary). In this option, your volume data is stored encrypted in Amazon S3, visible within your enterprise's network via an iSCSI interface. Recently accessed data is cached on- premises for low-latency local access. When you use this model, the cloud storage is primary, but you get low- latency access to your active working set in the cached volumes on premises.
- Gateway-Virtual Tape Library (VTL). In this option, you can configure a Gateway-VTL with up to 10 virtual tape drives per gateway, 1 media changer and up to 1500 virtual tape cartridges. Each virtual tape drive responds to the SCSI command set, so your existing on-premises backup applications (either disk-to-tape or disk-to-disk-to-tape) will work without modification.

No matter which option you choose, data is asynchronously transferred from your on-premises storage hardware to AWS over SSL. The data is stored encrypted in Amazon S3 using Advanced Encryption Standard (AES) 256, a symmetric- key encryption standard using 256-bit encryption keys. The AWS Storage Gateway only uploads data that has changed, minimizing the amount of data sent over the Internet.

The AWS Storage Gateway runs as a virtual machine (VM) that you deploy on a host in your data center running VMware ESXi Hypervisor v 4.1 or v 5 or Microsoft Hyper-V (you download the VMware software during the setup process). You can also run within EC2 using a gateway AMI. During the installation and configuration process, you can create up to 12 stored volumes, 20 Cached volumes, or 1500 virtual tape cartridges per gateway. Once installed, each gateway will automatically download, install, and deploy updates and patches. This activity takes place during a maintenance window that you can set on a per-gateway basis.

The iSCSI protocol supports authentication between targets and initiators via CHAP (Challenge- Handshake Authentication Protocol). CHAP provides protection against man-in-the-middle and playback attacks by periodically verifying the identity of an iSCSI initiator as authenticated to access a storage volume target. To set up CHAP, you must configure it in both the AWS Storage Gateway console and in the iSCSI initiator software you use to connect to the target.

After you deploy the AWS Storage Gateway VM, you must activate the gateway using the AWS Storage Gateway console. The activation process associates your gateway with your AWS Account. Once you establish this connection, you can manage almost all aspects of your gateway from the console. In the activation process, you specify the IP address of your gateway, name your gateway, identify the AWS region in which you want your snapshot backups stored, and specify the gateway time zone.

AWS Import/Export Security

AWS Import/Export is a simple, secure method for physically transferring large amounts of data to Amazon S3, EBS, or Amazon Glacier storage. This service is typically used by customers who have over 100 GB of data and/or slow connection speeds that would result in very slow transfer rates over the Internet. With AWS Import/Export, you prepare a portable storage device that you ship to a secure AWS facility. AWS transfers the data directly off of the storage device using Amazon's high-speed internal network, thus bypassing the Internet. Conversely, data can also be exported from AWS to a portable storage device.

Like all other AWS services, the AWS Import/Export service requires that you securely identify and authenticate your storage device. In this case, you will submit a job request to AWS that includes your Amazon S3 bucket, Amazon EBS region, AWS Access Key ID, and return shipping address. You then receive a unique identifier for the job, a digital signature for authenticating your device, and an AWS address to ship the storage device to. For Amazon S3, you place the signature file on the root directory of your device. For Amazon EBS, you tape the signature barcode to the exterior of the device. The signature file is used only for authentication and is not uploaded to Amazon S3 or EBS.

For transfers to Amazon S3, you specify the specific buckets to which the data should be loaded and ensure that the account doing the loading has write permission for the buckets. You should also specify the access control list to be applied to each object loaded to Amazon S3.

For transfers to EBS, you specify the target region for the EBS import operation. If the storage device is less than or equal to the maximum volume size of 1 TB, its contents are loaded directly into an Amazon EBS snapshot. If the storage device's capacity exceeds 1 TB, a device image is stored within the specified S3 log bucket. You can then create a RAID of Amazon EBS volumes using software such as Logical Volume Manager, and copy the image from S3 to this new volume.

For added protection, you can encrypt the data on your device before you ship it to AWS. For Amazon S3 data, you can use a PIN-code device with hardware encryption or TrueCrypt software to encrypt your data before sending it to AWS. For EBS and Amazon Glacier data, you can use any encryption method you choose, including a PIN-code device. AWS will decrypt your Amazon S3 data before importing using the PIN code and/or TrueCrypt password you supply in your import manifest. AWS uses your PIN to access a PIN-code device, but does not decrypt software- encrypted data for import to Amazon EBS or Amazon Glacier.

AWS Import/Export Snowball uses appliances designed for security and the Snowball client to accelerate petabyte-scale data transfers into and out of AWS. You start by using the AWS Management Console to create one or more jobs to request one or multiple Snowball appliances (depending on how much data you need to transfer), and download and install the Snowball client. Once the appliance arrives, connect it to your local network, set the IP address either manually or with DHCP, and use the client to identify the directories you want to copy. The client will automatically encrypt and copy the data to the appliance and notify you when the transfer job is complete.

After the import is complete, AWS Import/Export will erase the contents of your storage device to safeguard the data during return shipment. AWS overwrites all writable blocks on the storage device with zeroes. If AWS is unable to erase the data on the device, it will be scheduled for destruction and our support team will contact you using the email address specified in the manifest file you ship with the device.

When shipping a device internationally, the customs option and certain required subfields are required in the manifest file sent to AWS. AWS Import/Export uses these values to validate the inbound shipment and prepare the outbound customs paperwork. Two of these options are whether the data on the device is encrypted or not and the encryption software's classification. When shipping encrypted data to or from the United States, the encryption software must be classified as 5D992 under the United States Export Administration Regulations.

Further Reading

https://aws.amazon.com/security/security-resources/

Introduction to AWS Security Processes
Overview of AWS Security - Storage Services
Overview of AWS Security - Database Services
Overview of AWS Security - Compute Services
Overview of AWS Security - Application Services
Overview of AWS Security - Analytics, Mobile and Application Services Overview of AWS Security – Network Services