Security & Infrastructure Frequently Asked Questions

Introduction

The ThinkSafe platform provides robust and secure functionality for the rapid creation and deployment of connected, data-driven business applications. Our application architecture and failover design leverages the world class capabilities of Microsoft Azure to deliver a massively scalable, highly available and cost effective software offering.

Built on Microsoft Azure

ThinkSafe is hosted on Microsoft's Azure cloud infrastructure, which enables us to deliver highly scalable, available and fault tolerant services. Our application architecture has been designed to leverage Azure's powerful geo-redundancy, replication and recovery options, and is guided by Microsoft recommended best practices.

Frequently Asked Questions

Below is a set of system and security questions commonly asked of ThinkSafe. Please note that our infrastructure and system design is subject to change and thus may result in the answers below being revised from time to time. All answers apply to our cloud services unless otherwise indicated.

Privacy & Security

Is data "encrypted at rest", e.g. in static backups, databases, file storage?

Yes. All communications to/from storage is encrypted over HTTPS and access to storage is secured by passwords & secret keys.

Have any significant security breaches or incidents occurred in the past 5 years?

No.

How is physical security managed at Azure data centers?

All Microsoft data centers maintain state-of-the art physical security, including 24x7x365 surveillance, environmental protections and extensive secure access policies.
More information is available in the Windows Azure Security Overview document.

Are employees only provided with access to the network and network services that they have been specifically authorized to use based on their role? Clients?

Strictly employees only have network and infrastructure services access, access level is based on role. Client have no network or infrastructure services access.

Are privileged and generic account access tightly controlled and reviewed on a periodic basis, at least annually?

Yes.

Are shared user accounts prohibited for employees? Clients?

Some shared accounts are employed based on access role, otherwise employees have their own dedicated accounts. Clients have no access/accounts as mentioned above.

Does your password construction require multiple strength requirements, i.e. strong passwords and utilizes a random sequence of alpha, numeric and special characters?

We require a minimum 6 characters in passwords on our basic password management level. OWASP and NIST SP 800-63-3 password policy options will be available from May 2018.
Clients can implement their own choice of strength requirements by creating users & passwords through our APIs and turning off user password change functionality in the app.

Is the network boundary protected with a firewall with ingress and egress filtering?

Yes. All firewalls and load balancing facilities are provided by Microsoft's Azure platform. Refer to Microsoft's STAR self-assessment details found here: https://cloudsecurityalliance.org/star-registrant/microsoft-azure

Are public facing servers in a well-defined De-Militarized Zone (DMZ)?

Yes, this is inherited from Azure's default infrastructure zoning. Refer to Microsoft's STAR self-assessment details found here: https://cloudsecurityalliance.org/star-registrant/microsoft-azure

Is internal network segmentation used to further isolate sensitive production resources such as PCI data?

We do not store PCI data, but network segmentation is employed based on Azure's default configurations in this respect. Refer to Microsoft's STAR self-assessment details found here: https://cloudsecurityalliance.org/star-registrant/microsoft-azure

Is network intrusion Detection or Prevention implemented and monitored?

We run a broad spectrum of monitoring tools, supplemented by notifications and alerts provided by Azure. This includes intrusion detection and email confirmations of network access.

Are all desktops protected using regularly updated virus, worm, spyware and malicious code software?

Yes.

Are servers protected using industry hardening practices? Are the practices documented?

Yes, we utilise security services to provide regular system security audits.

Is there an ongoing program for network and vulnerability scanning, e.g. port scanning?

We subscribe to services that conduct internal penetration tests monthly using industry security standard tools and services.

Is there active vendor patch management for all operating systems, network devices and applications?

Yes. This is provided by Microsoft automatically via their Azure service.

Are all production system errors and security events recorded and preserved?

We preserve logs for a minimum of 1 month, with some remaining for up to 6 months, depending on severity and action required.

Are security events and log data regularly reviewed?

Yes. Logs are reviewed daily, weekly and monthly – depending on the nature of the log events.

Is there a documented privacy program in place with safeguards to ensure protection of client confidential information?

Yes.

Is there a process in place to notify clients if any privacy breach occurs?

Yes.

Do you store, process, transmit (i.e. "handle") Personally Identifiable Information (PII)?

Yes.

In what country or countries is PII stored?

This depends on where your account is hosted. We have 3 possible hosting locations – USA, EU and Australia.

Are system logs protected from alteration and destruction?

This is provided by Azure internally. Refer to Microsoft's STAR self-assessment details found here: https://cloudsecurityalliance.org/star-registrant/microsoft-azure

Are boundary and VLAN points of entry protected by intrusion protection and detection devices that provide alerts when under attack?

This is provided by Azure internally. Refer to Microsoft's STAR self-assessment details found here: https://cloudsecurityalliance.org/star-registrant/microsoft-azure

Are logs and events correlated with a tool providing warnings of an attack in progress?

Our monitoring tools provide access to the necessary logging events when seeking correlation to attacks.

Is system level security based on industry standard frameworks such as ISO27001, NIST800-53, or an equivalent framework as appropriate?

Microsoft Azure is audited annually by ISO27001 for compliance. ThinkSafe follows industry best practices for data and system security, including ISO 27001 recommendations. We are not currently audited or otherwise certified under such frameworks. We aim to formally gain a relevant certification in the future.

How is data is segregated from other clients within in the solution, including networking, front-ends, back-end storage and backups?

Every client account is logically separated from other clients, through the use of a required, persistent tenant identifier on all database records. Additionally all application code requires this tenant identifier for all operations - both read and write. An automated testing regime is also in place to protect code changes from regressions and possible cross-tenant data contamination. The tenant identifier is "hard linked" to every user account and logically enforced through fixed "WHERE" clauses on database queries and equivalent measures for file access. A platform user is not able to change or otherwise unlink their session or account from this tenant identifier. Thus there is no logical possibility of a user having login authorisation under a different tenant identifier. Even if they tried to access pages using a different tenant's id, the system would reject the request due to the user account not being registered to the requested tenant ID.

Do you have an Incident Response Plan?

Yes, we maintain a "living document" in our company cloud drive, which outlines disaster and incident response checklists, contact details and key system facilities for understanding and responding to incidents.

What level of network protection does ThinkSafe implement?

All network level security is managed by Microsoft Azure.
See: http://download.microsoft.com/download/C/A/3/CA3FC5C0-ECE0-4F87-BF4B-D74064A00846/AzureNetworkSecurity_v3_Feb2015.pdf

Does ThinkSafe install Microsoft Antimalware for Cloud Services and Virtual Machines or another antivirus solution on VMs, and can VMs be routinely reimaged to clean out intrusions that may have gone undetected?

We have the option to install Antimalware if needed, however our default configuration is the same as Microsoft's - which is antimalware is not installed. We don't remotely login or otherwise install software on our Cloud Services instances aside from our standard closed loop deployments through standard Azure management tools. Thus the risk of malware installation is minimal due to the lack of any direct login access to the instances.
Our servers are re-created using new, default Cloud Service instances every time we deploy a platform upgrade, which happens on average every 2 days or less.
This highly frequent re-creation of fresh instances also reduces any possible exposure time to malware in the highly unlikely event such was deployed to our servers.

Does the platform provide reports for Quality of Service (QOS) performance measurements (resource utilisation, throughput, availability etc)?

We don't provide such metrics to customers.

Is the disaster recovery program tested at least annually?

Yes, we perform recovery checks and tests annually.

What is the Recovery Time Objective (RTO) and Recovery Point Objective (RPO) of the system?

Our RTO is 4 hours, with RPO being 1 hour.

Do you provide backup and restore plans for individual customers?

All aspects are multi-tenanted, so backups are taken across entire customer base. We run complete file backups every 24 hours and benefit from Azure database point in time backups taken every 5 minutes.

What is the maximum time that back-ups are retained?

We retain database point-in-time backups for 30 days general file backups for 90 days minimum.

What is the expected turnaround time for a data restore?

Any customer restore in any non-disaster scenario must be requested and scheduled with our team. Turn around is between 1 and 2 business days.

Can a single entity (e.g. a Form) be restored without impacting the entire platform?

If restoration of a specific record or artefact is required by a customer, this can be performed online via a per request basis and is chargeable work. There is no impact on the platform or customer account.

Is High Availability provided – i.e. where one server instance becomes unavailable does another become available?

We run multiple server instances at all system tiers, including database (which is replicated). Failure of a server instance within the data centre is handled by Azure's load balancers, with the problem instance recycle and/or removed and replaced with a new instance.

Is data stored and available in another location (data centre) to meet disaster recovery requirements?

Yes. All data is replicated to a second data centre which differs by geographic location.

Is the ThinkSafe failover process an active/active, automated switchover process?

Failure of a server instance within the primary data centre is handled by Azure's load balancers, with the problem instance recycle and/or removed and replaced with a new instance.

In the event that the entire data centre were to have a critical failure, then switchover to our secondary centre is a manual process, as we need to perform a full assessment of the issue first to ensure there is no simple workarounds to keep the existing primary centre presence available. If we determine that a move to our secondary centre is required, then switchover will be initiated manually to meet our target recovery objectives.