Tenant Isolation in Microsoft Office 365

Published: November 9, 2016

This document describes how Microsoft implements logical isolation of customer data in a tenant within the Office 365 multi-tenant environment

© 2016 Microsoft Corporation. All rights reserved. This document is provided "as-is." Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it. Some examples are for illustration only and are fictitious. No real association is intended or inferred. This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

Introduction

One of the primary benefits of cloud computing is concept of a shared, common infrastructure across numerous customers simultaneously, leading to economies of scale. This concept is called *multi-tenancy*. Microsoft works continuously to ensure that the multi-tenant architecture of Microsoft Office 365 supports enterprise-level security, confidentiality, privacy, integrity, and availability standards.

Based upon the significant investments and experience gathered from Trustworthy Computing and the Security Development Lifecycle, Microsoft cloud services, including Office 365, were designed with the assumption that all tenants are potentially hostile to all other tenants, and we have implemented security measures to prevent the actions of one tenant from affecting the security or service of another tenant, or accessing the content of another tenant.

The two primary goals of maintaining tenant isolation in a multi-tenant environment are:

1. Preventing leakage of, or unauthorized access to, customer content across tenants; and
2. Preventing the actions of one tenant from adversely affecting the service for another tenant

Multiple forms of protection have been implemented throughout Office 365 to prevent customers from compromising Office 365 services or applications or gaining unauthorized access to the information of other tenants or the Office 365 system itself, including:

- Logical isolation of customer content within each tenant for Office 365 services is achieved through Azure Active Directory authorization and role-based access control.
- SharePoint Online provides data isolation mechanisms at the storage level.
- Microsoft uses rigorous physical security, background screening, and a multi-layered encryption strategy to protect the confidentiality and integrity of customer content. All Office 365 datacenters have biometric access controls, with most requiring palm prints to gain physical access. In addition, all U.S.-based Microsoft employees are required to successfully complete a standard background check as part of the hiring process. For more information on the controls used for administrative access in Office 365, see Office 365 Administrative Access Controls.
- Office 365 uses service-side technologies that encrypt customer content at rest and in transit, including BitLocker, per-file encryption, Transport Layer Security (TLS) and Internet Protocol Security (IPsec). For specific details about encryption in Office 365, see Data Encryption Technologies in Office 365.

Together, the above-listed protections provide robust logical isolation controls that provide threat protection and mitigation equivalent to that provided by physical isolation alone.

Isolation and Access Control in Azure Active Directory

Azure Active Directory was designed to host multiple tenants in a highly secure way through logical data isolation. Access to Azure Active Directory is gated by an authorization layer. Azure Active Directory isolates customers using tenant containers as security boundaries to safeguard a customer's

content so that the content cannot be accessed or compromised by co-tenants. Three checks are performed by Azure Active Directory's authorization layer:

1. Is the principal enabled for access to Azure Active Directory tenant?
2. Is the principal enabled for access to data in this tenant?
3. Is the principal's role in this tenant authorized for the type of data access requested?

No application, user, server, or service can access Azure Active Directory without the proper authentication and token or certificate. Requests are rejected if they are not accompanied by proper credentials.

Effectively, Azure Active Directory hosts each tenant in its own protected container, with policies and permissions to and within the container solely owned and managed by the tenant.

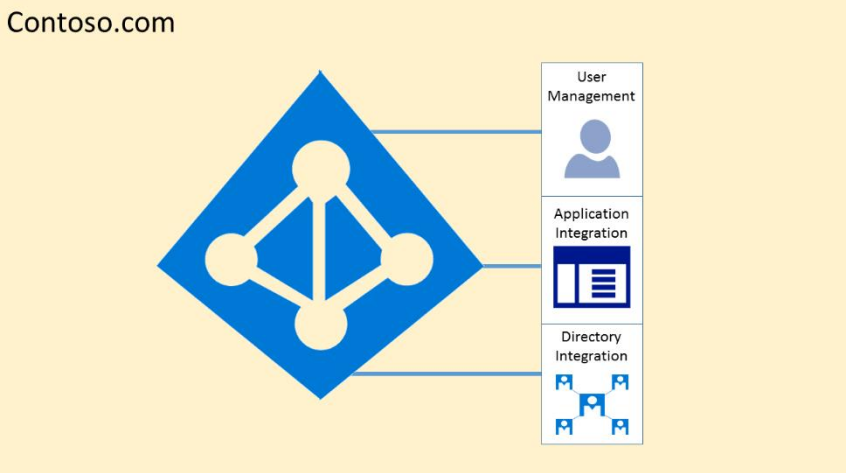

Figure 1 - Instance of Azure Active Directory for Contoso.com

The concept of tenant containers is deeply ingrained in the directory service at all layers, from portals all the way to persistent storage. Even when multiple Azure Active Directory tenant metadata is stored on the same physical disk, there is no relationship between the containers other than what is defined by the directory service, which in turn is dictated by the tenant administrator. There can be no direct connections to Azure Active Directory storage from any requesting application or service without first going through the authorization layer.

In the example below, Contoso and Fabrikam both have separate, dedicated containers, and even though those containers may share some of the same underlying infrastructure, such as servers and storage, they remain separate and isolated from each other, and gated by layers of authorization and access control.

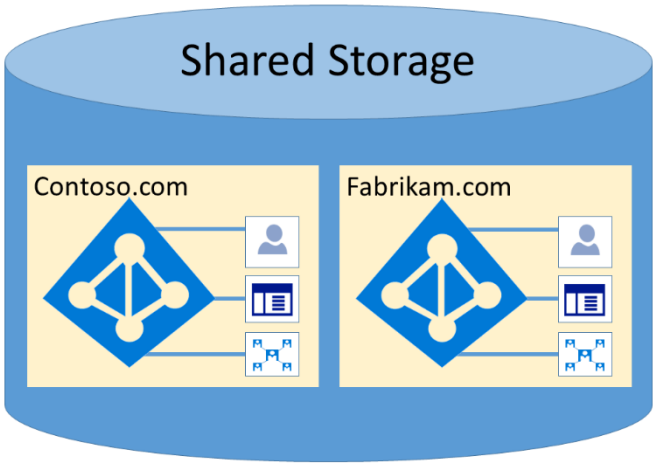

Figure 2 – Separate isolated instances of Azure Active Directory for multiple tenants on shared storage

In addition, there are no application components that can execute from within Azure Active Directory, and it is not possible for one tenant to forcibly breach the integrity of another tenant, access encryption keys of another tenant, or read raw data from the server.

By default, Azure Active Directory disallows all operations issued by identities in other tenants. Each tenant is logically isolated within Azure Active Directory through claims-based access controls. Reads and writes of directory data are scoped to tenant containers, and gated by an internal abstraction layer and a role-based access control (RBAC) layer, which together enforce the tenant as the security boundary. Every directory data access request is processed by these layers and every access request in Office 365 is policed by the logic above.

Azure Active Directory has North America, U.S. Government, European Union, Germany, and World Wide partitions. A tenant exists in a single partition, and partitions can contain multiple tenants. Partition information is abstracted away from users. A given partition (including all the tenants within it) is replicated to multiple datacenters. The partition for a tenant is chosen based on properties of the tenant (e.g., the country code). Secrets and other sensitive information in each partition is encrypted with a dedicated key. The keys are generated automatically when a new partition is created.

Azure Active Directory system functionalities are a unique instance to each user session. In addition, Azure Active Directory uses encryption technologies to provide isolation of shared system resources at the network level to prevent unauthorized and unintended transfer of information.

Isolation and Access Control in Office 365

Azure Active Directory and Office 365 use a highly complex data model that includes tens of services, hundreds of entities, thousands of relationships, and tens of thousands of attributes (entities, relationships and attributes are often application-specific). At a high level, Azure Active Directory and the service directories are the containers of tenants and recipients, and they are kept in sync using state-based replication protocols. In addition to the directory information held within Azure Active

Directory, each of the services also have their own directory services infrastructure (e.g., Exchange Online Directory Services, SharePoint Online Directory Services, etc.).

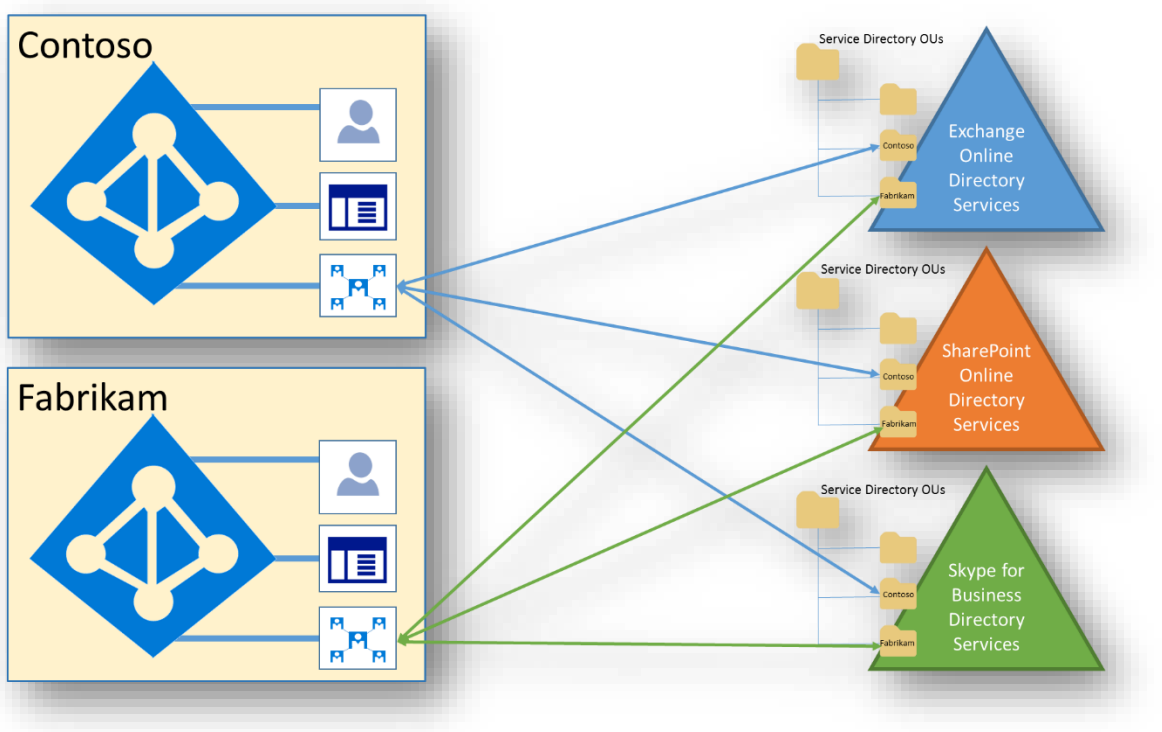

Figure 3 - Tenant data is synchronized from Azure Active Directory instance to separate Organizational Units within back-end service directories

Within this model, there is no single source of directory data. Every individual piece of data is owned by a specific system, but no single system holds all the data. Office 365 services cooperate with Azure Active Directory to realize the data model. Azure Active Directory is the "system of truth" for shared data, which is typically small and static data used often by every service. The federated model used within Office 365 and Azure Active Directory provides the shared view of the data.

Office 365 uses both physical storage and Azure cloud storage. Exchange Online (including Exchange Online Protection) and Skype for Business use their own storage for customer data. SharePoint Online leverages both its SQL Server storage and Azure storage, which necessitates the need for additional isolation of customer data at the storage level.

Exchange Online

Exchange Online stores customer data within mailboxes that are hosted within Extensible Storage Engine (ESE) databases called mailbox databases. This includes user mailboxes, linked mailboxes, shared mailboxes and public folder mailboxes. User mailboxes may also include saved Skype for Business content, such as conversation histories. User mailbox content includes emails and email attachments, calendaring and free/busy information, contacts, tasks, notes, Groups, and inference data.

Each mailbox database within Exchange Online contains mailboxes from multiple tenants. All mailboxes are secured by authorization code, including within a tenancy. As with an on-premises deployment of Exchange, by default only the assigned user has access to a mailbox. The access control list (ACL) that secures a mailbox contains an identity that is authenticated by Azure Active Directory at the tenant level. The mailboxes for a given tenant are limited to identities authenticated against that tenant's authentication provider, which include only users from that tenant. Content belonging to TenantA cannot in any way be obtained by users in TenantB, unless explicitly approved by TenantA.

Skype for Business

Skype for Business stores data in a variety of places:

- User and account information, which includes connection endpoints, tenant IDs, dial plans, roaming settings, presence state, contact lists, etc., is stored in the Skype for Business Active Directory servers, as well as in various Skype for Business database servers. Contact lists are stored in the user's Exchange Online mailbox if the user is enabled for both products, or on Skype for Business servers if the user is not. Skype for Business database servers are not partitioned per-tenant, but multi-tenancy isolation of data is enforced through RBAC.
- Meeting content, such as content users upload during Skype for Business meetings, is stored on Distributed File System shares. This content can also be archived in Exchange Online provided archiving is enabled. The DFS shares are not partitioned per-tenant but the content is secured with ACLs and multi-tenancy is enforced through RBAC.
- Call detail records, which is the activity history, such as call history, IM sessions, application sharing, IM history, etc., can also be stored in Exchange Online, but most call detail records are temporarily stored on call detail record (CDR) servers. Content is not partitioned per tenant, but multi-tenancy is enforced through RBAC.

SharePoint Online

There are several independent mechanisms unique to SharePoint Online that provide data isolation. SharePoint Online stores objects as abstracted code within application databases. For example, when a user uploads a file to SharePoint Online, that file is disassembled and translated into application code and stored in multiple tables across multiple databases.

If a user could gain direct access to the storage containing the data, the content would not be interpretable to a human or any system other than SharePoint Online. These mechanisms include security access control and properties. As described above, all SharePoint Online resources are secured by the authorization code and RBAC policy, including within a tenancy. The access control list (ACL) that secures a resource contains an identity that is authenticated at the tenant level. As with Exchange Online, in SharePoint Online, data for a given tenant are limited to identities authenticated against that tenant's authentication provider, which include only users from that tenant.

In addition to the ACLs, a tenant level property that specifies the authentication provider (which is the tenant-specific Azure Active Directory), is written once and cannot be changed once set. Once the

authentication provider tenant property has been set for a tenant, it cannot be changed using any APIs exposed to a tenant.

A unique *SubscriptionId* is also used for each tenant. All customer sites are owned by a tenant and are assigned a *SubscriptionId* unique to the tenant. The *SubscriptionId* property on a site is written once and cannot be changed. Once a site is assigned to a tenant, it cannot be moved to a different tenant later using the content store API. The *SubscriptionId* is also the key that is used to create the security scope for the authentication provider and is tied to the tenant.

SharePoint Online uses SQL Server and Azure storage for the storing of content. At the SQL level, the partition key for the content store is *SiteId*. When running a SQL query, SharePoint Online uses a *SiteId* that has been verified as part of a tenant-level *SubscriptionId* check.

SharePoint Online stores file binary blobs (e.g., the file streams) in Microsoft Azure. Each SharePoint Online farm has its own Microsoft Azure account and all the blobs saved in Azure are encrypted individually using a key that is stored in the SQL content store. The encryption key is not exposed directly to the end user, and is protected in code by the authorization layer. Finally, SharePoint Online has real-time monitoring in place to detect when an HTTP request reads or writes data for more than one tenant. It does this by tracking the *SubscriptionId* of the request identity against the *SubscriptionId* of the resource being accessed. A request accessing resources of more than one tenant should never happen by end-user. It can happen for service requests in a multi-tenant environment, though. For example, the search crawler pulls content changes for an entire database all at once. This usually involves querying sites of more than one tenant in a single service request, which is done for efficiency reasons.

Tenant Isolation in the Office Graph

The Office Graph models activity in Office 365 services, including Exchange Online, SharePoint Online, Yammer, Skype for Business, Azure Active Directory, and more, and in external services, such as other Microsoft services or third-party services. Office Graph components are used throughout Office 365. The Office Graph represents a collection of content and activity, and the relationships between them that happen across the entire Office suite. It uses sophisticated machine learning techniques to connect people to the relevant content, conversations and people around them. For example, the tenant index in SharePoint Online has an Office Graph index that is used to serve Delve queries, the Analytics Processing Engine in SharePoint Online is used to store signals and calculate insights, and Exchange Online calculates each user's recipient cache as input into tenant analytics.

The Office Graph contains information about enterprise objects, such as people and documents, as well as the relationships and interactions among these objects. The relationships and interactions are represented as *edges*. The Office Graph is segmented by tenant such that edges can only exist between *nodes* in the same tenancy. A *node* is an entity with a Uniform Resource Identifier (URI), node type, access control list, and a set of facets containing *metadata* and edges. Each node has associated metadata and edges, arranged into *facets* as in the Common Knowledge Model. *Metadata* are named

properties stored on a node which can be used for searching, filtering, or analysis within the office graph. A *facet* is a logical collection of metadata and edges on a node. Each facet describes one aspect of a node.

The Office Graph does not bring all the data into a single repository; rather, it stores metadata and relationships about data that lives elsewhere. The Office Graph consists of several data stores and processing components:

- The Tenant Graph Store provides bulk storage optimized for efficient analytics.
- The Active Content Cache provides quick random access to active node and edges to drive user experiences.
- The input router notifies components internal and external of changes to the tenant graph.

Analytics within each workload deduce insights relevant to the tenant-wide calculations and push them to the tenant graph. Tenant analytics reasons over all activity in a tenancy to produce insights into patterns of behavior. For example, Exchange Online calculates the recipient cache for each user with analytics that efficiently reason over each user's mailbox. These per-user analytics produce a set of *RecipientCache Edges* on each person, which are in turn pushed to the tenant graph. This keeps the as much of the analytics processing as close to the source data as possible.

Tenant Isolation in Delve

As mentioned previously, the Office Graph powers experiences that help people discover and collaborate on current activities in their enterprise, and provides an entity-centric platform for analytics to reason over content and activity across workloads and beyond Office 365. Delve is the first such experience powered by the Office Graph.

Delve is an Office 365 web experience that surfaces content from Office 365 and Yammer Enterprise to Office 365 users via the Office Graph. The web experience displays data as different boards, each with a certain topic, such as *Trending around me* or *Modified by me*. Each board consists of several document cards that display summary text and a picture from the document. The card lets users do things like open the document or a Yammer page for the document. There is a page for each person in an Office 365 tenant that displays the most relevant documents for this person, and icons that can invoke Exchange Online or Skype for Business for interacting with that person. Because Delve is based on the Office Graph API, it is bound by the tenant-based isolation of that API.

Tenant Isolation in Office 365 Search

SharePoint Online search uses a tenant separation model that balances the efficiency of shared data structures with protection against information leaking between tenants. With this model, we prevent the Search features from:

- Returning query results that contain documents from other tenants
- Exposing sufficient information in query results that a skilled user could infer information about other tenants

- Showing schema or settings from another tenant
- Mixing analytics processing information between tenants or store results in the wrong tenant
- Using dictionary entries from another tenant

For each type of tenant data, we use one or more layers of protection in the code to prevent accidental leaking of information. The most critical data has the most layers of protection to make sure that a single defect doesn't result in actual or perceived information leakage.

Tenant separation for the search index

The search index is stored on disk in the servers hosting the index components and the tenants share the index files. A tenant's indexed documents are visible only for queries for that tenant. Three independent mechanisms prevent information leakage:

- Tenant ID filtering
- Tenant ID term prefixing
- ACL checks

All three mechanisms would have to fail for Search to return documents to the wrong tenant.

Tenant ID filtering and Tenant ID term prefixing

Search prefixes every term that is indexed in the full-text index with the tenant ID. For example, when the term *"foo"* is indexed for a tenant with an ID of *"123"*, the entry in the full-text index is *"123foo."*

Every query is converted to include the tenant ID using a process called tenant ID filtering. For example, the query *"foo"* is converted to *"<guid>.foo AND tenantID:<guid>"*, where *<guid>* represents the tenant performing the query. This query conversion occurs within each index node and neither query nor content processing can influence it. Because the tenant ID is added to every query, the frequency of a term in other tenants can't be inferred by looking at best match ranking in one tenant.

Tenant ID term prefixing occurs only in the full-text index. Fielded searches, such as *"title:foo"*, go to a synthetic search index where terms aren't prefixed by tenant ID. Instead, fielded searches are prefixed with the field name. For example, the query *"title:foo"* is converted to *"fields.title:foo AND fields.tenantID:<guid>."* Because the frequency of a term doesn't influence ranking of hits in the synthetic search index, there's no need for tenant separation by term prefixing. For a fielded search like *"title:foo"*, tenant content separation depends on tenant ID filtering by query conversion.

Document Access Control List checks

Search controls access to documents through ACLs that are saved in the search index. Every item is indexed with a set of terms in a special ACL field. The ACL field contains one term per group or user that can view the document. Every query is augmented with a list of access control entry (ACE) terms, one for each group to which the authenticated user belongs.

For example, a query like *"<guid>.foo AND tenantID:<guid>"* becomes:

"<guid>.foo AND tenantID:<guid> AND (docACL:<ace1> OR docACL:<ace2> OR docACL:<ace3> …)"

Because user and group identifiers and hence ACEs are unique, this provides an extra level of security between tenants for documents that are only visible to some users. The same is the case for the special "Everyone except external users" ACE that grants access to regular users in the tenant. But since ACEs for "Everyone" are the same for all tenants, tenant separation for public documents depends on tenant ID filtering. Deny ACEs are also supported. The query augmentation adds a clause that removes a document from the result when there is a match with a deny ACE.

In Exchange Online search, the index is partitioned on mailbox ID for individual user's mailboxes instead of tenant ID (subscription ID) as in SharePoint Online. The partitioning mechanism is the same as SharePoint Online, but there is no ACL filtering.

Tenant Isolation in Azure Storage

Azure Storage is used to store data for multiple Office 365 services, including Office 365 Video and Sway. Azure Storage includes Blob storage, which is a highly-scalable, REST-based, cloud object store that is used for storing unstructured data. Azure Storage uses a simple access control model; each Azure subscription can create one or more Storage Accounts. Each Storage Account has a single secret key that is used to control access to all data in that Storage Account. This supports the typical scenario where storage is associated with applications and those applications have full control over their associated data; for example, Sway storing content in Azure Storage. All customer content for Sway is stored in shared Azure storage accounts. Each user's content is in a separate directory tree of blobs in Azure storage.

The systems managing access to customer environments (e.g., the Azure Portal, SMAPI, etc.) are isolated within an Azure application operated by Microsoft. This logically separates the customer access infrastructure from the customer applications and storage layer.

Tenant Isolation in Office 365 Video

Office 365 Video is an enterprise portal that provides organizations with a highly secure, organization-wide destination for posting, sharing, and discovering video content. In Office 365 Video, each tenant's videos are kept isolated and encrypted in all locations, and are only available to authenticated users that have access and permissions to the organization's videos. Office 365 Video uses a combination of technologies to accomplish this:

- SharePoint Online is used for storing the video file and metadata (video title, description, etc.). It also provides the security and compliance layer (including authentication), and search features.
- Azure Media Services provides transcoding, adaptive streaming, secure delivery (using AES encryption), and thumbnail features.

Azure Media Services is a platform-as-as-service offering for enabling end-to-end media workflows in the cloud. The platform provides a REST API for uploading, encoding, encrypting (with PlayReady), and delivery of media through Azure datacenters around the world.

All uploads for Office 365 Video occur via HTTPS. When a video file is uploaded, it is stored in SharePoint Online, and a copy of the file is sent via an encrypted channel to Azure Media Services. Azure Media Services transcodes the video into multiple formats that are optimized for viewing experience (e.g., mobile, low-bandwidth, high-bandwidth, etc.). These files, along with the original file acquired during upload, are stored in Azure Blob storage. The files are encrypted using AES 128 per the MPEG Common Encryption packaging algorithm (or an earlier PlayReady version) for playback, and encrypted using AES 256 storage encryption before being stored in Azure Blob storage.[1] Azure Media Services also generates a thumbnail for the video, which it transmits along with thumbnail metadata to SharePoint Online via an encrypted channel.

Resource Limits

Resource limits are enforced using quotas (limits) and throttling. Azure Active Directory and the individual Office 365 services use both. Limits are service-specific and change over time as new capabilities are added. For details on the current limits for the various services, see the following topics:

- Azure Active Directory service limits and restrictions
- Exchange Online Limits
- Exchange Online Protection Limits
- SharePoint Online software boundaries and limits
- Skype for Business Limits
- Yammer REST API and Rate Limits
- File Size Limits in Sway

In addition to these limits, several throttling mechanisms are used throughout Azure Active Directory and Office 365. Throttling within the service is especially important, given that network resources in Microsoft's datacenters are optimized for the broad set of customers that use the services. Throttling mechanisms include:

- Azure Active Directory and Office 365 feature user-level throttling, which limit the number of transactions or concurrent calls (by script or code) that can be performed by a single user.
- A default PowerShell throttling policy is assigned to each tenant at tenant creation. These settings affect other items, such as the maximum number of simultaneous PowerShell sessions that can be opened by a single administrator.

[1] Using the Azure Media Services Client SDK, customers can control which encryption is used. For example, a customer could apply Azure Media Services storage encryption (AES 256) to a high-value media asset before uploading it Azure Blob storage.

- Each Exchange Online customer has a default Exchange Web Services (EWS) policy that is tuned for EWS client operations, and throttling that applies to all Outlook clients.

Monitoring and Testing

Microsoft continuously monitors and explicitly tests for weaknesses and vulnerabilities in tenant boundaries, including monitoring for intrusion, permission violation attempts, and resource starvation. We also use multiple internal systems to continuously monitor for inappropriate resource utilization, which if detected, triggers built-in throttling.

Office 365 has internal monitoring systems that continuously monitor for any failure and drive automated recovery when failure is detected. Office 365 systems analyze deviations in service behavior and initiate self-healing processes that are built into the system. Office 365 also uses outside-in monitoring in which monitoring is performed from multiple locations both from trusted third-party services (for independent SLA verification) and our own datacenters to raise alerts. For diagnostics, we have extensive logging, auditing, and tracing. Granular tracing and monitoring helps us isolate issues and perform fast and effective root cause analysis.

While Office 365 has automated recovery actions where possible, Microsoft on-call engineers are available 24x7 to investigate all Severity 1 security escalations, and post-mortem reviews of every service incident contributes to continuous learning and improvement. This team includes support engineers, product developers, program managers, product managers, and senior leadership. Our on-call professionals provide timely backup and often can automate recovery actions, so that next time an event occurs, it can be self-healed.

Microsoft performs a thorough post-incident review each time an Office 365 security incident occurs regardless of the magnitude of impact. A post-incident review consists of an analysis of what happened, how we responded and how we prevent similar incidents in the future. In the interest of transparency and accountability, we share post-incident review for any major service incidents with affected customers. For specific details, see Office 365 Security Incident Management.

Assume Breach Methodology

Based on detailed analysis of security trends, Microsoft advocates and highlights the need for additional investments in reactive security processes and technologies that focus on detection and response to emerging threats, rather than solely the prevention of those threats. Because of changes in the threat landscape and in-depth analysis, Microsoft refined its security strategy beyond just preventing security breaches to one better equipped to deal with breaches when they do occur – a strategy which considers major security events not as a matter of *if*, but *when*.

While Microsoft's Assume Breach practices have been in place for many years, many customers are unaware of the work being done behind the scenes to harden the Microsoft cloud. Assume Breach is a mindset that guides security investments, design decisions and operational security practices. Assume Breach limits the trust placed in applications, services, identities, and networks by treating them all—internal and external—as insecure and already compromised. Although the Assume Breach

strategy was not borne from an actual breach of any Microsoft enterprise or cloud services, it was a recognition that many organizations, across the industry, were being breached despite all attempts to prevent it. While preventing breaches is a critical part of any organization's operations, those practices must be continuously tested and augmented to effectively address modern adversaries and advanced persistent threats. For any organization to prepare for a breach, they must first build and maintain robust, repeatable, and thoroughly-tested security response procedures.

While Prevent Breach security processes, such as threat modeling, code reviews, and security testing are very useful as part of the Security Development Lifecycle, Assume Breach provides numerous advantages that help account for overall security by exercising and measuring reactive capabilities in the event of a breach.

At Microsoft, we set out to accomplish this through ongoing war-games exercises and live site penetration testing of our security response plans with the goal of improving our detection and response capability. Microsoft regularly simulates real-world breaches, conducts continuous security monitoring, and practices security incident management to validate and improve the security of Office 365, Azure, and other Microsoft cloud services.

Microsoft executes its Assume Breach security strategy using two core groups:

- Red Teams (attackers)
- Blue Teams (defenders)

Both Microsoft Azure and Office 365 staff separate full-time red teams and blue teams.

Referred to as "Red Teaming", the approach is to test Azure and Office 365 systems and operations using the same tactics, techniques and procedures as real adversaries, against the live production infrastructure, without the foreknowledge of the Engineering or Operations teams. This tests Microsoft's security detection and response capabilities, and helps identify production vulnerabilities, configuration errors, invalid assumptions, and other security issues in a controlled manner. Every red team breach is followed by full disclosure between both teams to identify gaps, address findings, and improve breach response.

> **Note** No customer data is deliberately targeted during Red Teaming or live site penetration testing. The tests are against Office 365 and Azure infrastructure and platforms, as well as Microsoft's own tenants, applications and data. Customer tenants, applications, and content hosted in Office 365 or Azure are never targeted.

Red Teams

The red team is a group of full-time staff within Microsoft that focuses on breaching Microsoft's infrastructure, platform and Microsoft's own tenants and applications. They are the dedicated adversary (a group of ethical hackers) performing targeted and persistent attacks against Online Services (Microsoft infrastructure, platforms and applications but not end-customers' applications or content).

The role of the red team is to attack and penetrate environments using the same steps as an adversary:

Figure 4 - Breach Stages

Among other functions, red teams specifically attempt to breach tenant isolation boundaries to find bugs or gaps in our isolation design.

Blue Teams

The blue team is comprised of either a dedicated set of security responders or members from across the security incident response, Engineering, and Operations organizations. Regardless of their make-up, they are independent and operate separately from the red team. The blue team follows established security processes and uses the latest tools and technologies to detect and respond to attacks and penetration. Just like real-world attacks, the blue team does not know when or how the red team's attacks will occur or what methods may be used. Their job, whether it is a red team attack or an actual assault, is to detect and respond to all security incidents. For this reason, the blue team is continuously on-call and must react to red team breaches the same way they would for any other breach.

When an adversary, such as a red team, has breached an environment, the blue team must:

- Gather evidence left by the adversary
- Detect the evidence as an indication of compromise
- Alert the appropriate Engineering and Operation team(s)
- Triage the alerts to determine whether they warrant further investigation
- Gather context from the environment to scope the breach
- Form a remediation plan to contain or evict the adversary
- Execute the remediation plan and recover from breach

These steps form the security incident response that runs parallel to the adversary's, as shown below:

Figure 5 - Breach Response Stages

Red team breaches allow for exercising the blue team's ability to detect and respond to real-world attacks end-to-end. Most importantly, it allows for practiced security incident response prior to a genuine breach. Additionally, because of red team breaches, the blue team enhances their situational awareness which can be valuable when dealing with future breaches (whether from the red team or another adversary). Throughout the detection and response process, the blue team produces actionable intelligence and gains visibility into the actual conditions of the environment(s) they are trying to defend. Frequently this is accomplished via data analysis and forensics, performed by the

blue team, when responding to red team attacks and by establishing threat indicators, such as indicators of compromise. Much like how the red team identifies gaps in the security story, blue teams identify gaps in their ability to detect and respond. Furthermore, since the red teams model real-world attacks, the blue team can be accurately assessed on their ability, or inability, to deal with determined and persistent adversaries. Finally, red team breaches measure both readiness and impact of our breach response.

Summary

Microsoft continuously works to ensure that the multi-tenant architecture of Office 365 supports enterprise-level security, confidentiality, privacy, integrity, and availability standards. Multiple forms of protection have been implemented throughout Office 365 to prevent customers from compromising Office 365 services or applications or gaining unauthorized access to the information of other tenants or the Office 365 system itself. Together, these protections provide robust logical isolation controls that provides equivalent threat protection and mitigation to that provided by physical isolation alone.

All Office 365 services are built on top of Azure Active Directory and use the same authorization and RBAC model. All Office 365 requests are mediated through the authorization and access control features in Azure Active Directory. All Office 365 data sessions are either user-scoped or tenant-scoped, and users can't see outside of their tenant scope. Access to Office 365 objects is controlled with user account permissions that are enforced by Azure Active Directory and operating system ACLs. The authorization stack prevents a user from accessing data for which they don't have the appropriate credentials. Finally, there is no service code that allows a user from one tenant to execute commands against another tenant.

Microsoft also uses Red Teaming to test Azure and Office 365 systems and operations using the same tactics, techniques and procedures as real adversaries, against the live production infrastructure, without the foreknowledge of the Engineering or Operations teams. This tests Microsoft's security detection and response capabilities, and helps identify production vulnerabilities, configuration errors, invalid assumptions, and other security issues in a controlled manner. In addition, it allows us to measure both the readiness and impact of our breach response.

Materials in this Library

Microsoft publishes a variety of content for customers, partners, auditors, and regulators around security, compliance, privacy, and related areas. Below are links to other content in the Office 365 CXP Risk Assurance Documentation library.

Name	Abstract
Auditing and Reporting in Office 365	Describes the auditing and reporting features in Office 365 and Azure Active Directory and the various audit data that is available to customers via the Office 365 Security & Compliance Center, remote PowerShell, and the Management Activity API.
Controlling Access to Office 365 and Protecting Content on Devices	Describes the Conditional Access features in Office 365 and Microsoft Enterprise Mobility + Security, and how they are designed with built-in data security and protection to keep company data safe, while empowering users to be productive on the devices they love.
Data Encryption Technologies in Office 365	Provides an overview of the various encryption technologies that are used throughout Office 365, including features deployed and managed by Microsoft and features managed by customers.
Data Resiliency in Office 365	Describes how Microsoft prevents customer data from becoming lost or corrupt in Exchange Online, SharePoint Online, and Skype for Business, and how Office 365 protects customer data from malware and ransomware.
Defending Office 365 Against Denial of Service Attacks	Discusses different types of Denial of Service attacks and how Microsoft defends Office 365, Azure, and their networks against attacks.
Financial Services Compliance in Microsoft's Cloud Services	Describes how the core contract amendments and the Microsoft Regulatory Compliance Program work together to support financial services customers in meeting their regulatory obligations as they relate to the use of cloud services.
Microsoft Response to New FISC Guidelines in Japan (English) (Japanese)	Explains how Microsoft addresses the risks and requirements described in the FISC Revised Guidelines, and it describes features, controls, and contractual commitments that customers can use to meet the requirements in the Revised Guidelines.
Microsoft Threat, Vulnerability, and Risk Assessment of Datacenter Physical Security	Provides an overview regarding the risk assessment of Microsoft datacenters, including potential threats, controls and processes to mitigate threats, and indicated residual risks.
Office 365 Administrative Access Controls	Provides details on Microsoft's approach to administrative access and the controls that are in place to safeguard the services and processes in Office 365. For purposes of this document, Office 365 services include Exchange Online, Exchange Online Protection, SharePoint Online, and Skype for Business. Additional information about some Yammer Enterprise access controls is also included in this document.
Office 365 Customer Security Considerations	Provides organizations with quick access to the security and compliance features in Office 365 and considerations for using them.
Office 365 End of Year Security Report 2014	Covers security and legal enhancements made to Office 365 in calendar year 2014 than enables customers and partners to meet legal requirements surrounding independent verification and audits of Office 365.
Office 365 End of Year Security Report and Pen Test Summary 2015	Office 365 End of Year Security Report and Pen Test Summary for CY 2015.
Office 365 Mapping of CSA Cloud Control Matrix 3.0.1	Provides a detailed overview of how Office 365 maps to the security, privacy, compliance, and risk management controls defined in version 3.0.1-11-24-2015 of the Cloud Security Alliance's Cloud Control Matrix.
Office 365 Risk Management Lifecycle	Provides an overview of how Office 365 identifies, evaluates, and manages identified risks.
Office 365 Security Incident Management	Describes how Microsoft handles security incidents in Microsoft Office 365.
Privacy in Office 365	Describes Microsoft's privacy principles and internal privacy standards that guide the collection and use of customer and partner information at Microsoft and give employees a clear framework to help ensure that we manage data responsibly.
Self-Service Handling of Data Spills in Office 365 (restricted to Federal customers)	Reviews the spillage support provided by Office 365, the tools available to customers, and the configuration settings that should be reviewed in environments that are prone to data spills.