Data Resiliency in Microsoft Office 365

Published: January 9, 2017

This document describes how Microsoft prevents customer data from becoming lost or corrupt in Office 365, and how Office 365 protects customer data against malware and ransomware

Introduction

Given the complex nature of cloud computing, Microsoft is mindful that it's not a case of *if* things will go wrong, but rather *when*. We design our cloud services to maximize reliability and minimize the negative effects on customers when things do go wrong. We have moved beyond the traditional strategy of relying on complex physical infrastructure, and we have built redundancy directly into our cloud services. We use a combination of less complex physical infrastructure and more intelligent software that builds data resiliency into our services and delivers high availability to our customers.

This document describes data resiliency in Microsoft Office 365 from two perspectives:

1. How Microsoft prevents customer data from becoming lost or corrupt in Exchange Online, SharePoint Online, and Skype for Business; and
2. How Exchange Online, SharePoint Online, and Skype for Business protect customer data against malware and ransomware

Resiliency and Recoverability Are Built-in

Building in resiliency and recovery starts with the assumption that the underlying infrastructure and processes will fail at some point: hardware (infrastructure) will fail, humans will make mistakes, and software will have bugs. While it would be incorrect to say that software developers were not thinking about these things before the cloud, how these issues were handled in a typical IT implementation was very different before the cloud:

- First, hardware and infrastructure protections were significant. This meant having datacenters with 99.99% reliability required significant power and network redundancy, and servers were implemented with hardware-based clustering, dual power supplies, dual network interfaces, and the like.
- Second, process was paramount. Operations teams maintained rigorous procedures, change windows were employed, and there was often significant project management overhead.
- Third, deployment took place at a glacial pace. Deploying code without owning the source meant waiting for patch releases, and major version releases involved hardware replacement and significant capital outlay. Moreover, the only way to correct a problem was to rollback. Thus, most IT organizations would deploy only major releases to avoid the work to keep up-to-date.
- Finally, the scale of deployed systems, as well as the level of their interconnectedness was historically much smaller than it is now.

Today, customers expect continuous innovation from Microsoft without compromising quality, and this is one of the reasons why Microsoft's services and software are built with resiliency and recoverability in mind.

Office 365 Data Resiliency Principles

Resiliency refers to the ability of a cloud-based service to withstand certain types of failures and yet remain fully-functional from the customers' perspective. Data resiliency means that no matter what failures occur within Office 365, critical customer data remains intact and unaffected. To that end, Office 365 services have been designed around five specific resiliency principles:

1. There is critical and non-critical data. Non-critical data (for example, whether a message was read) can be dropped in rare failure scenarios. Critical data (for example, customer data such as email messages) should be protected at extreme cost. As a design goal, delivered mail messages are always critical, and things like whether a message has been read is non-critical.
2. Copies of customer data must be separated into different fault zones or as many fault domains as possible (e.g., datacenters, accessible by single credentials (process, server, or operator)) to provide failure isolation.
3. Critical customer data must be monitored for failing any part of Atomicity, Consistency, Isolation, Durability (ACID).
4. Customer data must be protected from corruption. It must be actively scanned or monitored, repairable, and recoverable.
5. Most data loss results from customer actions, so allow customers to recover on their own using a GUI that enables them to restore accidentally deleted items.

Through the building of our cloud services to these principles, coupled with robust testing and validation, Office 365 is able meet and exceed the requirements of customers while ensuring a platform for continuous innovation and improvement.

Dealing with Data Corruption

One of the challenging aspects of running a large-scale cloud service is how to handle data corruption, given the large volume of data and independent systems. Data corruption can be caused by:

- Application or infrastructure bugs, corrupting some or all of the application state
- Hardware issues that result in lost data or an inability to read data
- Human operational errors
- Malicious hackers and disgruntled employees
- Incidents in external services that result in some loss of data

Because greater resiliency in data integrity means fewer data corruption incidents, Microsoft has built into Office 365 protection mechanisms to prevent corruption from happening, as well as systems and processes that enable us to recover data if it does. Checks and processes exist within the various stages of the engineering release process to increase resiliency against data corruption, including:

- System Design
- Code organization and structure

- Code review
- Unit tests, integration tests, and system tests
- Trip wires tests/gates

Within Office 365 production environments, peer replication between datacenters ensures that there are always multiple live copies of any data. Standard images and scripts are used to recover lost servers, and replicated data is used to restore customer data. Because of the built-in data resiliency checks and processes, Microsoft maintains backups only of Office 365 information system documentation (including security-related documentation), using built-in replication in SharePoint Online and our internal code repository tool, Source Depot. System documentation is stored in SharePoint Online, and Source Depot contains system and application images. Both SharePoint Online and Source Depot use versioning and are replicated in near real-time.

Exchange Online Data Resilience

There are two types of corruption that can affect an Exchange database: physical corruption, which is typically caused by hardware (in particular, storage hardware) problems, and logical corruption, which occurs due to other factors. Generally, there are two types of logical corruption that can occur within an Exchange database:

- Database logical corruption The database page checksum matches, but the data on the page is wrong logically. This can occur when the database engine (the Extensible Storage Engine (ESE)) attempts to write a database page and even though the operating system returns a success message, the data is either never written to the disk or it's written to the wrong place. This is referred to as a *lost flush*. ESE includes numerous features and safeguards that are designed to prevent physical corruption of a database and other data loss scenarios. To prevent lost flushes from losing data, ESE includes a lost flush detection mechanism in the database along with a feature (single page restore) to correct it.
- Store logical corruption Data is added, deleted, or manipulated in a way that the user doesn't expect. These cases are generally caused by third-party applications. It's generally only corruption in the sense that the user views it as corruption. The Exchange store considers the transaction that produced the logical corruption to be a series of valid MAPI operations. The In-Place Hold features in Exchange Online provides protection from store logical corruption (because it prevents content from being permanently deleted by a user or an application).

Exchange Online performs several consistency checks on replicated log files during both log inspection and log replay. These consistency checks prevent physical corruption from being replicated by the system. For example, during log inspection, there is a physical integrity check which verifies the log file and validates that the checksum recorded in the log file matches the checksum generated in memory. In addition, the log file header is examined to make sure the log file signature recorded in the log header matches that of the log file. During log replay, the log file undergoes further scrutiny.

For example, the database header also contains the log signature which is compared with the log file's signature to ensure they match.

Protection against corruption of mailbox data in Exchange Online is achieved by using Exchange Native Data Protection, a resiliency strategy that leverages application-level replication across multiple servers and multiple datacenters along with other features that help protect data from being lost due to corruption or other reasons. These features include native features that are managed by Microsoft or the Exchange Online application itself, such as:

- Database availability groups and multiple database copies
- Lagged database copies
- Transport resiliency features, such as Safety Net and Shadow Redundancy
- Single Bit Correction
- Online Database Scanning
- Lost Flush Detection
- Single Page Restore
- Mailbox Replication Service
- Log File Checks
- Deployment on Resilient File System

For more information on the native features listed above, click on the above hyperlinks, and see below for additional information and for details on items without hyperlinks. In addition to these native features, Exchange Online also includes data resiliency features that customers can manage, such as:

- Single Item Recovery (enabled by default)
- In-Place Hold and Litigation Hold
- Deleted Item Retention and Soft-Deleted Mailboxes (both enabled by default)

Database Availability Groups

Every mailbox database in Office 365 is hosted in a database availability group (DAG) and replicated to geographically-separate datacenters within the same region. The most common configuration is four database copies in four datacenters; however, some regions have fewer datacenters (databases are replicated to three datacenters in India, and two datacenters in Australia and Japan). But in all cases, every mailbox database has four copies that are distributed across multiple datacenters, thereby ensuring that mailbox data is protected from software, hardware, and even datacenter failures.

Out of these four copies, three of them are configured as highly available. The fourth copy is configured as a lagged database copy. The lagged database copy is not intended for individual mailbox recovery or mailbox item recovery. Its purpose is to provide a recovery mechanism for the rare event of system-wide, catastrophic logical corruption.

Lagged database copies in Exchange Online are configured with a seven-day log file replay lag time. In addition, the Exchange Replay Lag Manager is enabled to provide dynamic log file play down for lagged copies to allow lagged database copies to self-repair and manage log file growth. Although lagged database copies are used in Exchange Online, it is important to understand that they are not a guaranteed point-in-time backup. Lagged database copies in Exchange Online have an availability threshold, typically around 90%, due to periods where the disk containing a lagged copy is lost due to disk failure, the lagged copy becoming a highly-available copy (due to automatic play down), as well as the periods where the lagged database copy is re-building the log replay queue.

Transport Resilience

Exchange Online includes two primary transport resilience features: Shadow Redundancy and Safety Net. Shadow Redundancy keeps a redundant copy of a message while it is in transit. Safety Net keeps a redundant copy of a message after the message is successfully delivered.

With Shadow Redundancy, each Exchange Online transport server makes a copy of each messages it receives before it acknowledges successfully receiving the message to the sending server. This makes all messages in the transport pipeline redundant while in transit. If Exchange Online determines the original message was lost in transit, a redundant copy of the message is redelivered.

Safety Net is a transport queue that is associated with the Transport service on a Mailbox server. This queue stores copies of messages that were successfully processed by the server. When a mailbox database or server failure requires activating an out-of-date copy of the mailbox database, messages in the Safety Net queue are automatically resubmitted to the new active copy of the mailbox database. Safety Net is also redundant, thereby eliminating transport as a single point of failure. It uses the concept of a Primary Safety Net and a Shadow Safety Net wherein if the Primary Safety Net is unavailable for more than 12 hours, resubmit requests become shadow resubmit requests, and messages are re-delivered from the Shadow Safety Net.

Message resubmissions from Safety Net are automatically initiated by the Active Manager component of the Microsoft Exchange Replication service that manages DAGs and mailbox database copies. No manual actions are required to resubmit messages from Safety Net.

Single Bit Correction

ESE includes a mechanism to detect and resolve single-bit CRC errors (aka single-bit flips) that are the result of hardware errors (and as such they represent physical corruption). When these errors occur, ESE automatically corrects them and logs an event in the event log.

Online Database Scanning

Online database scanning (also known as *database checksumming*) is the process where an ESE uses a database consistency checker to read each page and check for page corruption. The primary purpose is to detect physical corruption and lost flushes that may not be getting detected by transactional operations. Database scanning also performs post-store crash operations. Space can be

leaked due to crashes, and online database scanning finds and recovers lost space. The system is designed with the expectation that every database is fully scanned once every seven days.

Lost Flush Detection

A lost flush occurs when a database write operation that the disk subsystem/operating system returned as completed did not actually get written to disk, or was written in the wrong location. Lost flush incidents can result in database logical corruption, so to prevent lost flushes from resulting in lost data, ESE includes a lost flush detection mechanism. As database pages are written to passive copies, a check is performed for lost flushes on the active copy. If a lost flush is detected, ESE can repair the process using a page patching process.

Single Page Restore

Single page restore, aka *page patching*, is an automatic process where corrupt database pages are replaced by healthy copies from a healthy replica. The repair process for a corrupt page depends on whether the database copy is active or passive. When an active database copy encounters a corrupted page, it can copy a page from one of its replicas, provided the page it copies is completely up-to-date. This is accomplished by putting a request for the page into the log stream, which is the basis of mailbox database replication. As soon as a replica encounters the page request it responds by sending a copy of the page to the requesting database copy. Single page restore also provides an asynchronous communication mechanism for the active to request a page from replicas, even if the replicas are currently offline.

In case of corruption in a passive database copy, including a lagged database copy, because these copies are always behind their active copy, it is always safe to copy any page from the active copy to a passive copy. A passive database copy is by nature highly available, so during the page patching process, log replaying is suspended, but log copying continues. The passive database copy retrieves a copy of the corrupted page from the active copy, waits until the log file which meets the maximum required log generation requirement is copied and inspected, and then patches the corrupt page. Once the page has been patched, log replay resumes. The process is the same for the lagged database copy, except that the lagged database first replays all log files that are necessary to achieve a patchable state.

Mailbox Replication Service

Moving mailboxes is a key part of managing a large-scale email service. There are always updated technologies and hardware and version upgrades to deal with, so having a robust, throttled system that enables our engineers to accomplish this work while keeping the mailbox moves transparent to users (by making sure they stay online throughout the process) is key and making sure that the process scales up gracefully as mailboxes get larger and larger.

The Exchange Mailbox Replication Service (MRS) is responsible for moving mailboxes between databases. During the move, MRS performs a consistency check on all items within the mailbox. If a consistency issue is found, MRS will either correct the problem, or skip the corrupted items, thereby removing the corruption from the mailbox.

Because MRS is a component of Exchange Online, we can make changes in its code to address new forms of corruption that are detected in the future. For example, if we detect a consistency issue that MRS is not able to fix, we can analyze the corruption, change the MRS code and correct the inconsistency (if we understand how to).

Log File Checks

All transaction log files generated by an Exchange database undergo several forms of consistency checks. When a log file is created, the first thing done is a bit pattern is written and then a series of log writes is performed. This enables Exchange Online to execute a series of checks (lost flush, CRC and other checks) to validate each log file as it is written, and again as it is replicated.

Deployment on Resilient File System

To help prevent corruption from occurring at the file system level, Exchange Online is being deployed on Resilient File System (ReFS) partitions to provide improved recovery capabilities. ReFS is a file system in Windows Server 2012 and later that is designed to be more resilient against data corruption thereby maximizing data availability and integrity. Specifically, ReFS brings improvements in the way that metadata is updated which offers better protection for data and reduces data corruption cases. It also uses checksums to verify the integrity of file data and metadata ensuring that data corruption is easily found and repaired.

Exchange Online takes advantage of several ReFS benefits:

- More resiliency in data integrity means fewer data corruption incidents. Reducing the number of corruption incidents means fewer unnecessary database reseeds.
- Checksum running on metadata enabling detections of corruption cases sooner and more deterministically, allowing us to fix customer data corruption before grey failures occur on data volumes.
- Designed to work well with extremely large data sets—petabytes and larger—without performance impact
- Support for other features used by Exchange Online, such as BitLocker encryption.

Exchange Online also benefits from other ReFS features:

- Integrity (Integrity Streams). ReFS stores data in a way that protects it from many of the common errors that can normally cause data loss. Office 365 Search uses Integrity Streams to help with early disk corruption detection and checksums of file content. The feature also reduces corruption incidents caused by "Torn Writes" (when a write operation does not complete due to power outages, etc.).
- Availability (Salvage). ReFS prioritizes the availability of data. Historically, file systems were often susceptible to data corruption that would require the system to be taken offline for repair. Although rare, if corruption does occur, ReFS implements salvage, a feature that removes the corrupt data from the namespace on a live volume and ensures that good data is not adversely affected by non-repairable corrupt data. Applying the Salvage feature and

isolating data corruption to Exchange Online database volumes means that we can keep non-affected databases on a corrupted volume healthy between the time of corruption and repair action. This increases the availability of databases that would normally be affected by such disk corruption issues.

SharePoint Online Data Resilience

A key principle for SharePoint Online is to never have a single copy of any piece of data. SharePoint Online uses SQL Server replication, which is a set of technologies for copying and distributing data and database objects from one database to another, and then synchronizing between databases to maintain consistency.

For example, when a user saves a file in SharePoint Online, the file is chunked, encrypted, and stored within Azure Blob storage. Azure Blob service provides mechanisms to ensure data integrity both at the application and transport layers. This post will detail these mechanisms from the service and client perspective. MD5 checking is optional on both PUT and GET operations; however, it does provide a convenience facility to ensure data integrity across the network when using HTTP. Additionally, since HTTPS provides transport layer security additional MD5 checking is not needed while connecting over HTTPS as it would be redundant. Azure Blob service provides a durable storage medium, and uses its own integrity checking for stored data. The MD5's that are used when interacting with an application are provided for checking the integrity of the data when transferring that data between the application and service via HTTP.

To ensure data integrity the Azure Blob service uses MD5 hashes of the data in a couple different manners. It is important to understand how these values are calculated, transmitted, stored, and eventually enforced to appropriately design your application to utilize them to provide data integrity. For more information, see Windows Azure Blob MD5 Overview.

Metadata and pointers to the file are stored in a SQL Server database (the content database). All the chunks – files, pieces of files, and update deltas – are stored as blobs in Azure storage that are randomly distributed across multiple Azure storage accounts. The SQL database is hosted on a RAID 10 storage array which is synchronously mirrored to another RAID 10 storage array in a separate rack within the same datacenter. Asynchronous log shipping is then used to replicate the data to another RAID 10 storage array in a second datacenter. In addition to protecting data with RAID 10 and synchronous and asynchronous replication, scheduled data backups are taken which are also asynchronously replicated to the second datacenter.

In SharePoint Online, data backups are performed every 12 hours and retained for 14 days. SharePoint Online also uses a hot standby system that includes paired geographically-separate datacenters within the same customer data location region (for example, Chicago and San Antonio for customers who have provisioned their tenant in the United States) configured as active/active. For example, there are live users that have Chicago as their primary datacenter and San Antonio as a failover datacenter, and live users that have San Antonio as their primary datacenter and Chicago as their failover datacenter.

Skype for Business Data Resilience

The key principles for Skype for Business are to:

1. Never have a single copy of any piece of customer data; and
2. Keep the system loosely-coupled

Skype for Business uses a blob store, which contains persistent data that is used to provision new front-end servers, as well as user data. An internal Microsoft technology is then deployed as part of the Skype for Business service and used for creating highly reliable, distributable, and scalable applications. This technology defines votes, determines where users are homed, and replicates data. Persistent user data is synchronously replicated to two or more front-end servers and Skype for Business clients only see success when the data is written to all replicas correctly.

A backup service replicates data between blob stores, and an internal monitoring service automatically probes the system for latencies or malfunctions at all the times. Instances that fail can be replaced seamlessly while other instances continue to operate and run independent stacks in more than one pool, either in the same region or in another region. In case of an event affecting availability, the system performs failover from the primary to one of the backups while recovery steps (automatic and/or initiated by the engineering team monitoring the service 24/7) are performed in the background, usually without the users being aware of the incident.

Customer data created in Skype for Business resides in the user's Exchange Online mailbox, and that data is protected by the resiliency features described above in Exchange Online Data Resilience.

Protecting Customer Data from Malware

Malware consists of viruses, spyware and other malicious software. Office 365 includes protection mechanisms to prevent malware from being introduced into Office 365 by a client or by an Office 365 server. The use of anti-malware software is a principal mechanism for protection of Office 365 assets from malicious software. The anti-malware software detects and prevents computer viruses, malware, rootkits, worms, and other malicious software from being introduced into any service systems. Anti-malware software provides both preventive and detective control over malicious software.

Each anti-malware solution in place tracks the version of the software and what signatures are running. The automatic download and application of signature updates at least daily from the vendor's virus definition site is centrally managed by the appropriate anti-malware tool for each service team.

The following functions are centrally managed by the appropriate anti-malware tool on each endpoint for each service team:

- Automatic scans of the environment
- Periodic scans of the file system (at least weekly)
- Real-time scans of files as they are downloaded, opened, or executed

- Automatic download and application of signature updates at least daily from the vendor's virus definition site
- Alerting, cleaning, and mitigation of detected malware

When anti-malware tools detect malware, they block the malware and generate an alert to Office 365 service team personnel, Office 365 Security, and/or the security and compliance team of the Microsoft organization that operates our datacenters. The receiving personnel initiate the incident response process. Incidents are tracked and resolved, and post-mortem analysis is performed.

SharePoint Online and OneDrive for Business Protection Against Malware

To further protect the service against malicious files, SharePoint Online (which includes OneDrive for Business) prohibits certain file types from being uploaded and prevents content from being executed directly in the service. This prohibits the potential spread of malware from within the service. Anti-malware software is installed both as part of the initial build on all systems, and on all SharePoint Online servers, enabling further protection by actively scanning document repositories and code within SharePoint Online sites and libraries.

Exchange Online Protection Against Malware

All email messages for Exchange Online travel through Exchange Online Protection (EOP), which quarantines and scans in real time all email and email attachments both entering and leaving the system for viruses and other malware. Administrators do not need to set up or maintain the filtering technologies; they are enabled by default. However, administrators can make company-specific filtering customizations using the Exchange Admin Center.

Using multiple anti-malware engines, EOP offers multilayered protection that's designed to catch all known malware. Messages transported through the service are scanned for malware (including viruses and spyware). If malware is detected, the message is deleted. Notifications may also be sent to senders or administrators when an infected message is deleted and not delivered. You can also choose to replace infected attachments with either default or custom messages that notify the recipients of the malware detection.

The following helps provide anti-malware protection:

- Layered Defenses Against Malware Multiple anti-malware scan engines used in EOP help protect against both known and unknown threats. These engines include powerful heuristic detection to provide protection even during the early stages of a malware outbreak. This multi-engine approach has been shown to provide significantly more protection than using just one anti-malware engine.
- Real-time Threat Response During some outbreaks, the anti-malware team may have enough information about a virus or other form of malware to write sophisticated policy rules that detect the threat even before a definition is available from any of the engines used by the service. These rules are published to the global network every 2 hours to provide your organization with an extra layer of protection against attacks.

- Fast Anti-Malware Definition Deployment The anti-malware team maintains close relationships with partners who develop anti-malware engines. As a result, the service can receive and integrate malware definitions and patches before they are publicly released. Our connection with these partners often allows us to develop our own remedies as well. The service checks for updated definitions for all anti-malware engines every hour.

Advanced Threat Protection

Advanced Threat Protection (ATP) is an email filtering service that provides additional protection against specific types of advanced threats, including malware and viruses. Exchange Online Protection currently uses a robust and layered anti-virus protection powered by multiple engines against known malware and viruses. ATP extends this protection through a feature called Safe Attachments, which protects against unknown malware and viruses, and provides better zero-day protection to safeguard your messaging system. All messages and attachments that don't have a known virus/malware signature are routed to a special hypervisor environment, where a behavior analysis is performed using a variety of machine learning and analysis techniques to detect malicious intent. If no suspicious activity is detected, the message is released for delivery to the mailbox.

Exchange Online Protection also scans each message in transit in Office 365 and provides time of delivery protection, blocking any malicious hyperlinks in a message. Attackers sometimes try to hide malicious URLs with seemingly safe links that are redirected to unsafe sites by a forwarding service after the message has been received. Safe Links proactively protects your users if they click such a link. That protection remains every time they click the link, and malicious links are dynamically blocked while good links are accessible.

ATP also offers rich reporting and tracking capabilities, so you can gain critical insights into who is getting targeted in your organization and the category of attacks you are facing. Reporting and message tracing allows you to investigate messages that have been blocked due to an unknown virus or malware, while the URL trace capability allows you to track individual malicious links in the messages that have been clicked.

For more information about ATP, see Introducing Exchange Online Advanced Threat Protection and Office 365 Advanced Threat Protection.

SharePoint Online and OneDrive for Business Protection Against Ransomware

There are many forms of ransomware attacks, but one of the most common forms is where a malicious individual encrypts a user's important files and then demands something from the user, such as money or information, in exchange for the key to decrypt them. Ransomware attacks are on the rise, particularly those that encrypt files that are stored in the user's cloud storage. For example, Crowti (also known as Cryptowall) and Tescrypt (also known as Teslacrypt) are two ransomware families that have infected over half a million PCs in the first half of 2015. For more information about ransomware, see the Microsoft Malware Protection Center.

You can use versioning to protect SharePoint Online lists and SharePoint Online and OneDrive for Business libraries from some, but not all, of these types of ransomware attacks. Versioning is enabled by default in OneDrive for Business and disabled by default in SharePoint Online.[1] When versioning is enabled in SharePoint Online site lists, you can look at earlier versions and recover them, if necessary. That enables you to recover versions of items that pre-date their encryption by the ransomware. Some organizations also retain multiple versions of items in their lists for legal reasons or audit purposes.

By default, versioning in SharePoint is turned off. To turn it on and implement your versioning decisions, you must either have Full Control or Design permissions. For detailed steps to enable versioning for SharePoint libraries and lists, see Enable and configure versioning for a list or library. For detailed steps to do this, see Restore a previous version of a document in OneDrive for Business.

SharePoint Online and OneDrive for Business Recycle Bins

SharePoint Online administrators can restore a deleted site collection by using the SharePoint Online admin center. SharePoint Online users have a Recycle Bin where deleted content is stored. They can access the Recycle Bin to recover deleted documents and lists, if they need to. Items in the Recycle Bin are retained for 90 days. The following data types are captured by the Recycle Bin:

- Site collections
- Sites
- Lists
- Libraries
- Folders
- List items
- Documents
- Web Part pages

Site customizations made through SharePoint Designer are not captured by the Recycle Bin. For more information, see Manage the Recycle Bin of a SharePoint site collection. See also, Restore a deleted site collection.

As illustrated in the example below, versioning does not protect against ransomware attacks that copy files, encrypt them, and then delete the original files. However, end-users can leverage the Recycle Bin to recover OneDrive for Business files after a ransomware attack occurs.

Scenario: Recovering SharePoint Online and OneDrive for Business files after a ransomware attack

Andrew is the head of Human Resources at Blue Yonder Airlines. One day, Andrew receives an email from Tom, an employee that was recently terminated. Tom says he is applying for a new job elsewhere and would like Andrew to review and approve an attachment that contains Tom's

[1] If you don't see the Version history command, version history may be turned off. Depending on how your organization has set up personal sites, you may be able to turn on document versioning.

statements to the prospective employer about his termination from Blue Yonder Airlines. Andrew opens the attachment which, unbeknownst to him, begins a drive-by download process that silently installs crypto-based ransomware onto Andrew's computer. After the ransomware is running on Andrew's computer:

1. It connects to a remote server, where it uploads connection information, such as the public IP address, location, and system information for Andrew's computer, including what operating system is running.
2. The remote server then generates a random 2048-bit RSA key pair that's associated with Andrew's computer.
3. The ransomware copies the public key to Andrew's computer and starts copying each file using a pre-determined list of file extensions. As each copy is created, it is also encrypted using the public key, and the original file is then deleted from Andrew's hard drive. This typically continues until all files with the specified file extensions have been copied and encrypted, including files that are synchronized to the cloud.
4. Once the encryption process has completed, the ransomware typically executes some local commands to stop the Volume Shadow Copy (VSS) service that runs on all modern versions of Windows and controls versioning (history), backup, and restoration of data on a host computer. The command run by the ransomware stops the VSS service and deletes the VSS cache.
5. Every time Andrew tries to open a file with the specified file extensions, he receives a message similar to the following:

Figure 1 - Example ransomware message

Andrew isn't worried, though, because his files are located on SharePoint Online and OneDrive for Business. The files on Andrew's hard drive are simply copies of files that are also stored in the cloud. To remove the ransomware and recover his files (without paying the ransom), Andrew follows this process:

1. Because he doesn't need any local data, Andrew can reformat his disk, reinstall Windows and reinstall his applications.
2. After his computer has been reinstalled, Andrew uses the Web interface for OneDrive for Business to access the Recycle Bin.

3. Once in the Recycle Bin, Andrew recovers all his files. The version he recovers is the version that existed right before the ransomware attack.
4. After the original files have been restored, Andrew can delete the encrypted files from OneDrive and synchronize the recovered files to his computer.

Monitoring, Alerting and Self-Healing

Given the scale of Office 365, it would be impossible to keep customer data resilient and safe from malware without built-in monitoring that is comprehensive, alerting that is intelligent, and self-healing that is fast and reliable. Monitoring a set of services at the scale of Office 365 is very challenging. New mindsets and methodologies needed to be introduced, and whole new sets of technology needed to be created to operate and manage the service in a connected global environment. We have moved away from the traditional monitoring approach of data collection and filtering to create alerts to an approach that is based on data analysis; taking signals and building confidence in that data and then using automation to recover or resolve the issue. This approach helps take humans out of the recovery equation, which in turn makes operations less expensive, faster, and less error prone.

Fundamental to Office 365 monitoring is a collection of technologies that comprise our Data Insights Engine, which is built on Azure, SQL Azure, and open-source streaming database technology. It is designed to collect and aggregate data and reach conclusions. Currently, it processes more than 500 million events per hour from more than 100,000 servers (~15 TB per day) scattered across dozens of datacenters in many regions, and these numbers are growing.

Office 365 uses *outside-in monitoring*, which involves creating synthetic transactions to test everything that is important. For example, in Exchange Online each scenario is testing every database worldwide every five minutes in a scattered fashion, providing near continuous coverage of everything that lives in the system. From multiple locations, 250 million test transactions per day are performed to create a robust baseline or heartbeat for the service.

Office 365 also uses the concept of *Red Alert*, which shrinks down all the monitoring signals from all of the machines in our datacenters to something manageable by a human being. The concept is quite simple: If something is happening across multiple signals, there must be something going on. It is not about building confidence in one signal, it is about having reasonable fidelity for each signal so that you get greater accuracy. This monitoring system is so powerful that we do not have 24x7 staff watching our monitors; all we have is the machinery that wakes up if it detects a problem, in which case it will page the appropriate on-call personnel, or more often as is the case, it will just go ahead and solve the problem. Once we start collecting signals and building red alerts off them, we can start triangulating across all our service partitions. Below is an example of an Office Service Alert which illustrates this.

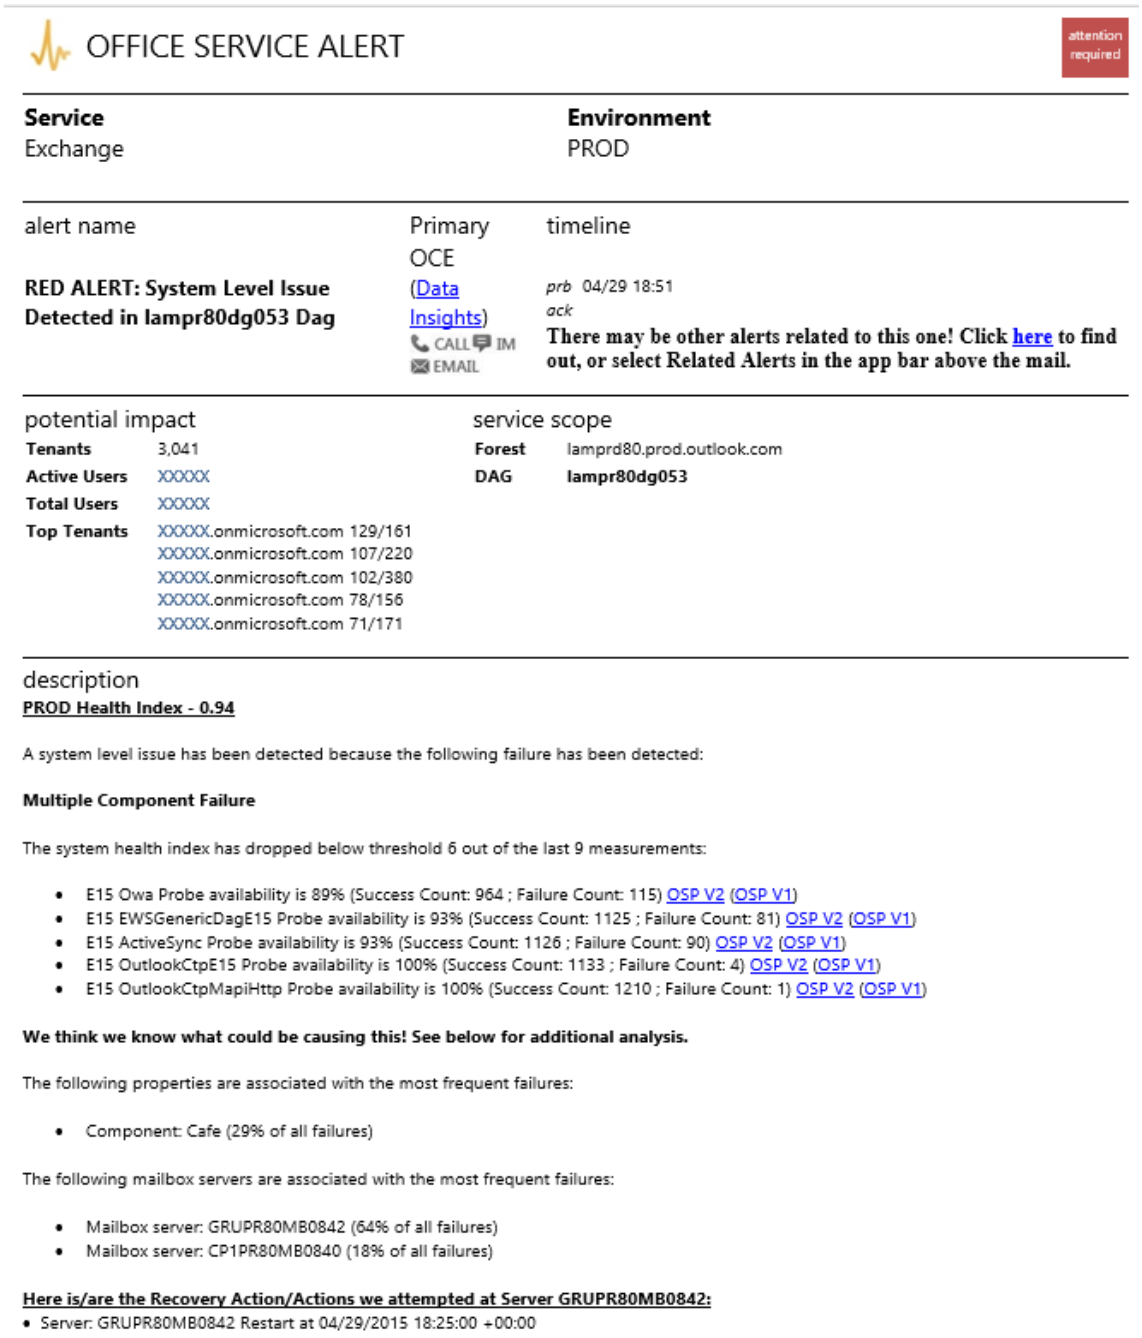

OFFICE SERVICE ALERT

attention required

Service
Exchange

Environment
PROD

alert name

RED ALERT: System Level Issue Detected in lampr80dg053 Dag

Primary OCE
(Data Insights)
CALL IM EMAIL

timeline
prb 04/29 18:51
ack
There may be other alerts related to this one! Click here to find out, or select Related Alerts in the app bar above the mail.

potential impact

Tenants	3,041
Active Users	XXXXX
Total Users	XXXXX
Top Tenants	XXXXX.onmicrosoft.com 129/161 XXXXX.onmicrosoft.com 107/220 XXXXX.onmicrosoft.com 102/380 XXXXX.onmicrosoft.com 78/156 XXXXX.onmicrosoft.com 71/171

service scope

Forest	lamprd80.prod.outlook.com
DAG	**lampr80dg053**

description

PROD Health Index - 0.94

A system level issue has been detected because the following failure has been detected:

Multiple Component Failure

The system health index has dropped below threshold 6 out of the last 9 measurements:

- E15 Owa Probe availability is 89% (Success Count: 964 ; Failure Count: 115) OSP V2 (OSP V1)
- E15 EWSGenericDagE15 Probe availability is 93% (Success Count: 1125 ; Failure Count: 81) OSP V2 (OSP V1)
- E15 ActiveSync Probe availability is 93% (Success Count: 1126 ; Failure Count: 90) OSP V2 (OSP V1)
- E15 OutlookCtpE15 Probe availability is 100% (Success Count: 1133 ; Failure Count: 4) OSP V2 (OSP V1)
- E15 OutlookCtpMapiHttp Probe availability is 100% (Success Count: 1210 ; Failure Count: 1) OSP V2 (OSP V1)

We think we know what could be causing this! See below for additional analysis.

The following properties are associated with the most frequent failures:

- Component: Cafe (29% of all failures)

The following mailbox servers are associated with the most frequent failures:

- Mailbox server: GRUPR80MB0842 (64% of all failures)
- Mailbox server: CP1PR80MB0840 (18% of all failures)

Here is/are the Recovery Action/Actions we attempted at Server GRUPR80MB0842:

- Server: GRUPR80MB0842 Restart at 04/29/2015 18:25:00 +00:00

Figure 2 - Example Office Service Alert

Based on the combination of the failure alert and the Red Alerts, this alert indicates exactly which components could be having a problem, and that the system is going to try to correct the problem by itself by restarting a mailbox server.

In addition to self-healing capabilities such as single page restore, Exchange Online includes several features that take an approach to monitoring and self-healing which focuses on preserving the end-user experience. These features include *Managed Availability*, which provides built-in monitoring and

recovery actions, and AutoReseed, which automatically restores database redundancy after a disk failure.

Managed Availability

Managed availability provides a native health checking and recovery solution that monitors and protects the end user's experience through recovery-oriented actions. Managed availability is the integration of built-in monitoring and recovery actions with the Exchange high availability platform. It's designed to detect and recover from problems as soon as they occur and are discovered by the system. Unlike previous external monitoring solutions and techniques for Exchange, managed availability doesn't try to identify or communicate the root cause of an issue. Instead, it's focused on recovery aspects that address three key areas of the end-user experience:

- **Availability** Can users access the service?
- **Latency** How is the experience for users?
- **Errors** Are users able to accomplish what they want?

Managed availability is an internal feature that runs on every Office 365 server running Exchange Online. It polls and analyzes hundreds of health metrics every second. If something is found to be wrong, most of the time it is fixed automatically. But there will always be issues that managed availability will not be able to fix on its own. In those cases, managed availability will escalate the issue to an Office 365 support team by means of event logging.

AutoReseed

Exchange Online servers are deployed in a configuration that stores multiple databases and their log streams on the same non-RAID disk. This configuration is often referred to as *just a bunch of disks* (JBOD) because no storage redundancy mechanisms, such as RAID, are being used to duplicate the data on the disk. When a disk fails in a JBOD environment, the data on that disk is lost.

Given the size of Exchange Online and the fact that deployed within it are millions of disk drives, disk drive failures are a regular occurrence in Exchange Online. In fact, more than 100 fail every day. When a disk fails in an on-premises enterprise deployment, an administrator must manually replace the failed disk and restore the affected data. In a cloud deployment the size of Office 365, having operators (cloud administrators) manually replacing disks is neither practical nor economically feasible.

Automatic Reseed, or *AutoReseed*, is a feature that is the replacement for what is normally operator-driven action in response to a disk failure, database corruption event, or other issue that necessitates a reseeding of a database copy. AutoReseed is designed to automatically restore database redundancy after a disk failure by using spare disks that have been provisioned on the system. If a disk fails, the database copies stored on that disk are automatically reseeded to a preconfigured spare disk on the server, thereby restoring redundancy.

Summary

Office 365 services have been designed around specific resiliency principles that are designed to protect data from corruption, to separate data into different fault zones, to monitor data for failing any part of the ACID test, and to allow customers to recover on their own. Office 365 is designed to maximize reliability and minimize the negative affects to customers when things do go wrong. We have moved beyond the traditional strategy of relying on complex physical infrastructure and we have built redundancy into our cloud services. We use a combination of less complex physical infrastructure and more intelligent software that builds customer data resiliency into our services and delivers high availability to our customers.

Because more resiliency in customer data integrity translates to less data corruption incidents, Microsoft has built into Office 365 several protection mechanisms to prevent corruption from happening, as well as systems and processes that enable us to recover when it happens. Protection against corruption of mailbox data in Exchange Online is achieved by using Exchange Native Data Protection, a resiliency strategy that leverages application-level replication across multiple servers and multiple datacenters along with other features that help protect data from being lost due to corruption or other reasons. SharePoint Online and OneDrive for Business files are stored within a SQL Server database that is hosted on a RAID 10 storage array and synchronously mirrored to another RAID 10 array in a separate rack within the same datacenter. Asynchronous log shipping is then used to replicate the data to another RAID 10 storage array in a second datacenter. In addition to protecting data with RAID 10 and synchronous and asynchronous replication, scheduled data backups are taken every few minutes which are also asynchronously replicated to the second datacenter. Skype for Business uses a blob store, which contains persistent data that is used to provision new front-end servers, as well as user data. Persistent user data is synchronously replicated to two or more front-end servers and Skype for Business clients only see success when the data is written to all replicas correctly.

Office 365 also includes protection mechanisms to prevent malware from being introduced into Office 365 by a client or by an Office 365 server. The use of anti-malware software is a principal mechanism for protection of Office 365 assets from malicious software. You can also use versioning and recycle bins to protect SharePoint Online lists and SharePoint Online and OneDrive for Business libraries from certain types of ransomware attacks.

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 our 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 available to customers. Also details the various audit data that is available to customers via the Office 365 Security & Compliance Center, remote PowerShell, and the Management Activity API. Also describes the internal logging data that is available to Microsoft Office 365 engineers for detection, analysis, and troubleshooting.
Controlling Access to Office 365 and Protecting Content on Devices	Describes the Conditional Access (CA) features in Microsoft 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. It also provides guidance on how to address common concerns around data access and data protection using Office 365 features.
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.

Name	Abstract
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.
Tenant Isolation in Office 365	Describes how Microsoft implements logical isolation of tenant data within Office 365 environment.