3E Cloud

Reliability and Availability Information Guide

Document Version 1.0
Date of issue: November 24, 2021

THOMSON REUTERS®

Contents

About this document

INTENDED READERSHIP

This guide is intended for current and potential customers who have technical roles such as CTO, IT Director, and the like.

IN THIS GUIDE

The purpose of this guide is to provide specific information regarding the availability and reliability of 3E Cloud. Among the topics covered on this topic are:

- Availability
- Reliability
- Backups
- Disaster recovery
- Service interruptions and incident management

WHERE TO FIND MORE INFORMATION

If you are an existing customer, you can find more information about 3E Cloud in our Knowledge Base. To access it, navigate to https://customerportal.elite.com/ and select the **Knowledge Base** option.

Availability and Reliability Overview

Thomson Reuters has decades of experience managing and operating enterprise-scale applications like Westlaw, Practical Law, ONESOURCE, and many more. We use availability and reliability measures to assess the accessibility of our applications.

Availability is a measure of the percentage of time a system or infrastructure is up and able to take requests.

Reliability is the ability of a system to perform its documented functions.

This information guide highlights the Thomson Reuters processes and structures designed to deliver a high standard of availability and reliability for 3E Cloud.

OUR APPROACH

We employ a diversified strategy designed to ensure a high standard of availability and reliability with 3E Cloud.

To start, we use Microsoft Azure native services and infrastructure. Azure data centers are designed and built to limit unauthorized access and withstand natural disasters.

Second, we employ scaling strategies that allow us to scale our services up or down based on the demand.

Third, we have an experienced and dedicated operations team focused on managing our cloud infrastructure. We have proactive monitoring tools that alert our team based on established triggers related to the infrastructure utilized by our applications or the applications themselves.

Lastly, we monitor the frequency of errors and other metrics, and our product teams improve our applications to address identified errors.

Availability

EXPECTATIONS

Please refer to the 3E Cloud and 3E Cloud CARE service level agreements (SLAs) at https://www.elite.com/terms/ for specific information regarding our operational standards pertaining to availability.

To summarize, subject to the exceptions and definitions detailed in the SLAs, Thomson Reuters uses commercially reasonable efforts to make 3E Cloud available for customer use at least 99.9 percent of the time.

Note that this availability requirement only applies to the live production environment; it does not apply to non-production environments. More details surrounding our scheduled maintenance can be found in the Service Interruptions & Incident Management section of this document.

HIGH AVAILABILITY FOR COMPUTE INFRASTRUCTURE

Most of the web applications in the 3E Cloud product suite are hosted on Azure App Service plans. These Azure App Service plans are configured to be zone redundant in case of zone failures. Along with zone redundancy, the App Service plans are configured to scale out to handle an increase in load.

The following link provides more details about the zone redundancy for Azure App Service: https://docs.microsoft.com/en-us/azure/app-service/how-to-zone-redundancy

HIGH AVAILABILITY FOR DATABASE INFRASTRUCTURE

Most of the 3E Cloud suite products use either Azure SQL Database or Azure Cosmos DB as database technologies. For example, 3E core uses Azure SQL Database, while 3E Templates uses Azure Cosmos DB.

For Azure SQL Database, Microsoft uses Azure Blob Storage to store data files. Blob storage provides built-in measures to address availability and redundancy. As for the compute for the Azure SQL Database, the compute is stateless and designed for redundancy and failover to provide high availability.

For products that use Cosmos DB, the Cosmos DB is configured for zone redundancy. Cosmos DB will ensure replicas are placed across multiple zones within a given region to provide high availability and resiliency to zonal failures. The following link has more information regarding the availability of Cosmos DB: https://docs.microsoft.com/en-us/azure/cosmos-db/high-availability

REDUNDANCY

To address availability in the 3E Cloud product suite, single points of failure are eliminated by providing redundancy using paired regions.

Azure operates in multiple geographies around the world. An Azure geography is a defined area of the world that contains at least one Azure region. An Azure region is an area within a geography containing one or more data centers.

Each Azure region is paired with another region within the same geography, together making a regional pair. Across the regional pairs, Azure serializes platform updates (planned maintenance) so that only one paired region is updated at a time. In the event of an outage affecting multiple regions, at least one region in each pair will be prioritized for recovery.

The benefits of paired regions are:

Physical isolation	Hundreds of miles of separation reducing risk during natural disasters.
Platform-provided replication	For example, geo-redundant storage.
Region recovery order	Especially useful when restoring service across multiple regions.
Sequential updates	Paired regions are never updated simultaneously.

Data residency	All but one geography (Brazil) offers paired regions within the same geography.

The below table indicates the paired regions for each geography.

Geography	Primary Azure Region	Secondary Azure Region
Canada	Canada Central (Toronto)	Canada East (Quebec City)
North America	East US 2 (Virginia)	Central US (Iowa)
Australia	Australia East (New South Wales)	Australia South East (Victoria)
Europe	West Europe (Netherlands)	North Europe (Ireland)
UK	UK South (London)	UK West (Cardiff)

For more information about paired regions, see https://docs.microsoft.com/en-us/azure/best-practices-availability-paired-regions.

MONITORING

We monitor the availability of our systems by running a series of scheduled availability tests.

All our distributed applications self-report availability on an established interval to the monitoring system.

A second set of tests and self-health checks, confirm availability for all dependent systems. If any of the tests fail, alerts are generated that confirm the part of the system that has experienced a failure.

Reliability

EXPECTATIONS

The 3E Cloud solution is architected to adjust resources elastically to allow you to operate without limiting the users connected to the system.

ADJUSTMENT / SCALABILITY

The scalability and elasticity have been architected into the 3E Cloud product suite through many measures:

- The 3E core offering, where most of the demand happens, is a single-tenant solution and can be scaled independently for each customer.
- Multi-tenant services such as 3E Templates and 3E Workspace are architected to provide scalability and elasticity to meet the needs of all customers using the services at the same time.
- Utilization metrics and monitoring are in place for both single-tenant and multi-tenant services that automatically scale resources based on resource needs and alert the operations team if manual intervention is required.

MONITORING

We monitor performance for all levels of the application down to the infrastructure our applications run on.

If performance thresholds are exceeded in either the infrastructure or applications, alerts are generated.

We are also alerted when exceptions occur in the backend applications. Application exceptions are unforeseen errors that the application is not able to handle. These may be benign or can be an indication of something more serious that should be addressed as quickly as possible.

UPDATES

We are continually evolving and improving our offering and products.

Environments are scaled periodically based on current usage and data processing requirements, as well as product platform changes and new expected features.

Our product releases include fixes and enhancements. For 3E Cloud, we have the following release types:

- **Feature Release** – Includes new feature enhancements, new defect fixes, and a roll-up of previous releases and relevant hotfixes. May also include noteworthy architecture and technology changes when required.
- **Hotfix Release** – Includes urgent fixes and updates that affect daily operations and need resolution without a Feature update. Usually includes a roll-up of relevant previous hotfixes.

For more information on our guidelines regarding 3E Cloud feature releases and hotfixes, please refer to the Upgrade Guidelines for 3E Cloud (Knowledge Base article TR-18382) or Upgrade Guidelines for 3E Cloud CARE customers (Knowledge Base article TR-18383).

Backups

Thomson Reuters leverages tools from Azure, other third parties, and custom-built solutions to automate, monitor, and manage data backups. This helps ensure data that is stored in databases, file systems, and other storage mechanisms are retrievable and restorable to help customers recover from an unplanned event.

PROTECTION

We have measures in place designed to protect backups against accidental deletion and compromise. We store multiple copies of each backup in the primary data center. Generally, there are three (3) copies stored across availability zones. Backups are also replicated to a paired failover data center in the same geopolitical region.

There are also measures in place designed to protect backups against unauthorized access. Backup data is stored in an Azure storage area that Thomson Reuters does not have direct access to, and backups are always encrypted. For SQL backups, the backup agent gets temporary, one-time, write-only access to the storage in an effort to ensure unauthorized access does not result in access, modification, or deletion of a backup.

PROCESS

The backup process varies depending on the type of backup being performed. The following outlines at a high level our backups for both databases and file/blob storage.

Database

- Full backups occur each week.
- Differential backups are taken every 12 to 24 hours.
- Transaction log backups happen every 5 to 10 minutes, based on database activity, to support point-in-time restores.

File and Blob Storage

- Files and blobs are backed up in multiple ways.
- Data is replicated to the failover region (geo-redundant storage) to protect against a region failure.
- Blob versions are continuously maintained to enable point-in-time restore.
- Point-in-time restore is designed to protect against logical errors such as accidental deletes or file corruption.

RETENTION

The backup retention timeframes vary depending on the type of backup. The following outlines at a high-level our backup retentions for both databases and file/blob storage.

Database

- Seven (7) days of point-in-time restore
- Four (4) weekly full backups
- Twelve (12) monthly backups
- Five (5) year-end backups

File and Blob Storage

- Point-in-time restore is maintained for thirty (30) days

CUSTOMER ACCESS

Customers have no visibility into the physical environment or the infrastructure of 3E Cloud. One of the key benefits of moving to the cloud is having Thomson Reuters monitor the environments, system health, and other behind-the-scenes operations of the 3E Cloud solution for you, such as maintaining and restoring backups as needed.

As such, customers should only require access to the applications themselves, and access to backups is not provided.

Disaster Recovery

In the past, Disaster Recovery (DR) has been tightly coupled to Availability. However, Disaster Recovery is different from Availability in that DR focuses on being able to recover from a natural or man-made disaster where geographical separation is important. These types of disasters happen very rarely, but when they do, we need to ensure we can recover from such events. In contrast, Availability is about being able to have redundancy and eliminate a single point of failure within a single region.

For 3E Cloud, Disaster Recovery means that a 3E product suite has a replicated, functional, usable, and operationally managed DR instance available in an alternate paired Azure region. This alternate paired DR Azure region can provide the appropriate product/service to our customers in the event of a disaster.

OUR APPROACH

The 3E Cloud product suite follows the combination of “Geo-Replication” and “Backup and Restore” approach for Disaster Recovery.

In this approach, Thomson Reuters geo-replicates storage to the alternate paired Azure region and regular backups of certain critical core databases are stored in the alternate paired Azure region. In the case of DR, Thomson Reuters would make storage in the alternate paired Azure region as primary and restore databases from their backups. After the storage and databases are up and running, Thomson Reuters adds and/or scales out the compute infrastructure in the alternate paired region, and the application will be available in the alternate paired region. All the products and services are recovered in the correct order with the most up-to-date data to ensure the smallest downtime.

We use several technologies from both Microsoft and Thomson Reuters that are designed to support smooth and reliable recovery from failures. Azure Backup and Azure Site Recovery are examples of those technologies.

RECOVERY PLAN

The 3E Cloud Disaster Recovery Plan includes steps for responding to the following types of issues:

Severity	Description
Severity 3	Data integrity-related, involving deletion or modification of data. Requires recovery of one or more databases or file systems to a point in time.
Severity 2	Application or service level integrity. One or more specific instances of an application, service, or virtual machine are no longer functional, and functionality cannot be restored through troubleshooting, system updates, so forth, but where the Azure infrastructure itself is not impacted.
Severity 1	Azure regional integrity or outage. An issue that results in widespread instability or total outage within an Azure region, such that the application needs to be recovered to another region in order to regain full operability.

Each severity level has a defined group of people who are responsible for deciding when to execute the disaster plan. The 3E Cloud operations team performs the disaster recovery procedures, with escalation to the product delivery teams if needed.

Thomson Reuters regularly runs disaster and failure simulations and does a full disaster recovery test at least once a year.

RPO / RTO

Two important concepts that need to be understood with any disaster recovery plan are the Recovery Point Objective (RPO) and the Recovery Time Objective (RTO).

The below table lists our time objectives for both of these. This table applies only to Severity 1 issues, which are described earlier in this section.

Term	Description	Time
Recovery Point Objective (RPO)	The maximum acceptable amount of data loss after a disaster, expressed as an amount of time.	1 hour
Recovery Time Objective (RTO)	The maximum amount of time the system can be down, before being restored to operational state.	24 hours

Service Interruptions & Incident Management

SERVICE INTERRUPTION NOTIFICATIONS

The following outlines the types of service interruptions, and what you can expect for notifications around such events. For more specific details, including the scheduled maintenance timeframes, see the 3E Cloud and 3E Cloud CARE Service Level Agreements (SLAs) at https://www.elite.com/terms/.

Direct notifications will often be issued via email. Customers are responsible for updating and maintaining support key contacts with Thomson Reuters via the customer support platform, which is at present the Thomson Reuters Customer Portal at https://customerportal.elite.com. These are the contacts who will be notified of 3E Cloud service interruptions. Customers are also responsible for informing end users of service interruptions, as appropriate.

Scheduled Maintenance

Scheduled maintenance can largely be classified as the planned rollout of Feature Releases and other necessary changes related to a Feature Release, such as architecture and technology updates, that could cause all or a portion of live/production 3E Cloud to become unavailable for a period of time.

These are fully planned, managed, and executed activities that are documented in advance in our 3E Cloud Release Schedule (Knowledge Base article TR-6446). All customers are strongly encouraged to subscribe to this article. This article will serve as the scheduled maintenance notification, enabling you to be informed in advance of such events to help with any necessary preparation. In addition to this schedule, you will be contacted (usually in the form of an email) at least two (2) weeks in advance of an upgrade to either Preview or Production instances.

For more information, please refer to our 3E Cloud Upgrade Guidelines. For 3E Cloud, see Knowledge Base article TR-18382. For 3E Cloud CARE, see Knowledge Base article TR-18383.

Unscheduled Maintenance

Unscheduled maintenance can largely be classified as the rollout to the live/production 3E Cloud environment of Hotfix Releases and architecture and technology updates not related to a Feature Release. Hotfix Releases normally are not scheduled but will be issued on an as-needed basis and thus usually are not listed on the 3E Cloud Release Schedule.

These updates may not cause any noticeable interruption of service, but at times they may. Where unscheduled maintenance is expected to have a material impact on your use of 3E Cloud, you will be directly notified in advance (usually in the form of an email). As this type of maintenance is often due to pressing situations that require a more urgent update, the advanced notice timeline will vary and cannot be predetermined.

Unplanned Interruptions

In the unfortunate event that there is an unplanned interruption in service, you will be notified directly (usually in the form of an email) in a timely manner that an outage is occurring and that it is being investigated. Additional details will be provided once determined. Regular updates will be sent to ensure you understand the status of the situation during the course of the service interruption.

APPLICATION INCIDENT MANAGEMENT

If problems arise that impact your use of the 3E Cloud applications, Thomson Reuters has a response process in place to address such incidents as they are identified.

Customers follow the standard procedure of opening a case with our Support team who will investigate the inquiry.

You can create a case through the Thomson Reuters Customer Portal at https://customerportal.elite.com or by contacting your regional Product Support team (see https://www.elite.com/support/ for contact information).

The Product Support team has 3E experts who handle inquiries for both 3E on-premises and 3E Cloud customers alike, as the application-level features are largely the same.

Initial response times depend on the severity of the case. The severity levels are:

- Urgent
- High
- Normal
- Low

As an example, a severity level of High means that the issue has a significant business impact on the customer. The initial response time for North America, ANZ, and EMEA is within 2 hours of initial notification. See the 3E Cloud and 3E Cloud CARE service level agreements (SLAs) at https://www.elite.com/terms/ for more information, including our standard support response times for each of the severity levels.

Each case is routed to a support analyst. If the case cannot be resolved by the initially assigned analyst, the case will be referred to a more senior support analyst. Depending on the nature of the issue, members of the development and 3E Cloud operations teams may also be involved.

3E Cloud CARE customers also have access to a Technical Account Manager (TAM), who helps coordinate the support resources.

For more specific details, please refer to our 3E Support Guidelines (Knowledge Base article TR-5472).

SECURITY INCIDENT MANAGEMENT

For incidents involving security-related issues, Thomson Reuters employs a tiered incident management and escalation model based on ITIL (Information Technology Infrastructure Library).

Incidents are triaged based on criticality and assigned through incident leads in each region. Incident command follows documented response practices, as well as established communications and escalation practices. Coordination of incidents also involves IT and product teams and the use of outside communications expertise and general counsel where necessary.