Infrastructure,
Security,
Service Level Availability
& Support

TABLE OF CONTENTS

1 OVERVIEW

This document provides an overview of the Echo360 platform architecture, the security measures in place to protect the platform, the target service availability, and the support program to maintain Echo360-licensed solutions.

2 ECHO360 ARCHITECTURE

Echo360 is a cloud-based Software as a Service (SaaS) platform designed and built on Amazon Web Services (AWS). We chose AWS because it provides us with the flexibility to scale systems to meet traffic demands on a robust and secure platform. The Echo360 platform is deployed in four separate AWS Regions around the world. These regional deployments allow us to keep a specific region’s data isolated to that local geographical area and ensures low latency access for users to the platform. Each regional service deployment is isolated from others and each has their own unique URL for access. We leverage AWS CloudFront as our Content Delivery Network (CDN) for streaming media and serving assets to end users which ensures smooth reliable playback of video to the user.

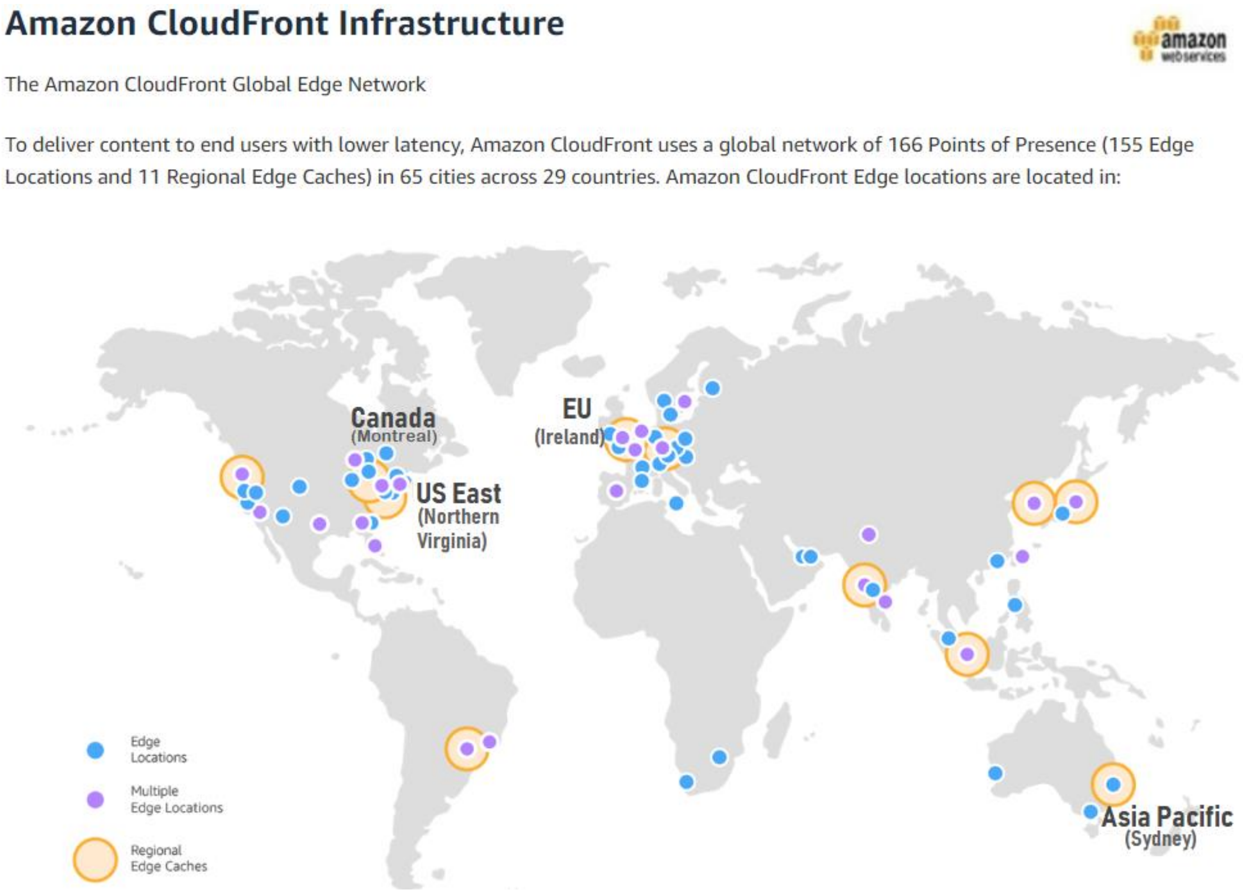

Amazon CloudFront Infrastructure

The Amazon CloudFront Global Edge Network

To deliver content to end users with lower latency, Amazon CloudFront uses a global network of 166 Points of Presence (155 Edge Locations and 11 Regional Edge Caches) in 65 cities across 29 countries. Amazon CloudFront Edge locations are located in:

2.1 Echo360 Regions

Our platform currently operates out of the following four regions:

- U.S. East – Northern Virginia
- Canada – Montreal, Quebec
- European Union (EU) – Dublin, Ireland
- Asia Pacific – Sydney, Australia

Each Region consists of two or more customer accessible Availability Zones. Availability Zones are distinct locations that are engineered to be insulated from failures in other Availability Zones with different and redundant sources of power, internet, and other required utilities. They also provide inexpensive, low-latency network connectivity to other Availability Zones in the same region. The Echo360 platform spans at least three Availability Zones in each of the AWS regions where the platform is provisioned.

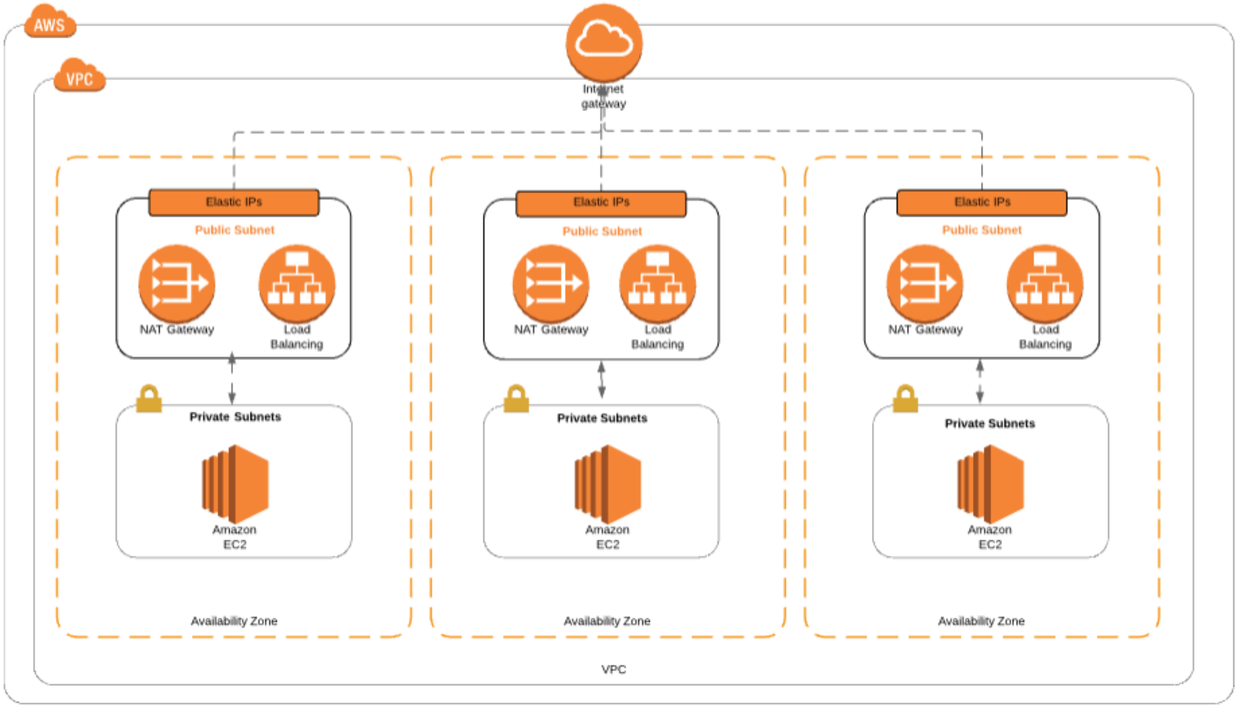

Our platform serves traffic continuously out of multiple Availability Zones within a region to allow us to maintain the user experience even if an Availability Zone is offline.

2.2 Scalability

The Echo360 platform was built from the beginning to be a multi-tenant application and our systems are designed to scale as traffic increases. Systems that receive variable amounts of traffic depending on the usage of the platform such as live streaming, media processing, and application servers scale up and down throughout the day to respond quickly to changes in traffic and resource demands. This is completely seamless to the end user and allows our platform to adapt to large spikes in traffic without any manual intervention.

2.3 Microservices

The Echo360 platform is built on top of Amazon Web Services which provides a highly secure foundation and allows us to deliver a highly secure and scalable solution. Our platform leverages numerous AWS services such as EC2, S3, DynamoDB, RDS, and CloudFront. On top of these AWS services, our platform is built around a microservices concept which is part of our high-availability strategy.

A microservice is a small independently deployable service that contributes a feature or subset of a feature to our overall application. These smaller services isolate the impact of a service outage to just that functionality that is covered by the specific service rather than the entire platform. This architecture also limits the scope and complexity of a given service which allows us to provide a more secure and performant platform. Our distributed microservices architecures provides a highly-scalable data processing system.

2.4 Architecture Diagram

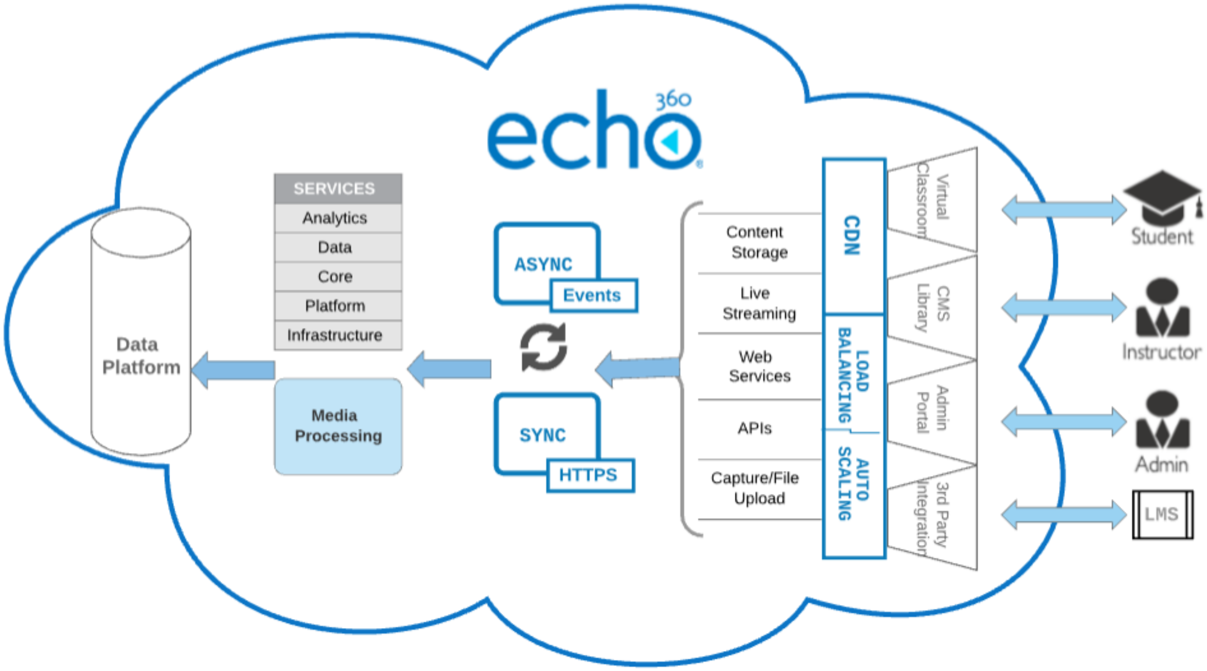

2.5 Site Reliability Engineering Team

Echo360 employs a Site Reliability Engineering team that operates with a mandate to ensure the platform is secure, performing efficiently, and available to customers at all times. This SRE team is involved in every aspect of the software development lifecycle and is directly responsible for all Echo360 infrastructure and services.

2.6 Infrastructure

Echo360 designed and built our infrastructure platform to maximize maintainability, scalability, and security for our customers. The infrastructure platform provides a standardized way for Engineering teams to design, release, and monitor applications and services within the Echo360 platform. These standardized patterns allow us to bake internal service monitoring, logging, and eventing into every service and feature added to the platform which allows the SRE and Engineering teams to receive real time feedback on how a new feature or service is performing.

Abnormal error rates, logs, or other metrics are flagged for review and bugs are inserted into the bug pipeline for evaluation and resolution. These industry leading patterns and our approach to automation and monitoring at an infrastructure and service level allows us to deliver an incredibly stable platform. Our product and engineering teams use the platform telemetary to guide feature development, bug fixes, and incident response. Ultimately, this creates a continuous cycle of engineering improvement, speed of iteration, and predictability of results. Detailed information about our infrastructure can be found in the sections that follow.

2.6.1 Virtual Private Cloud Network

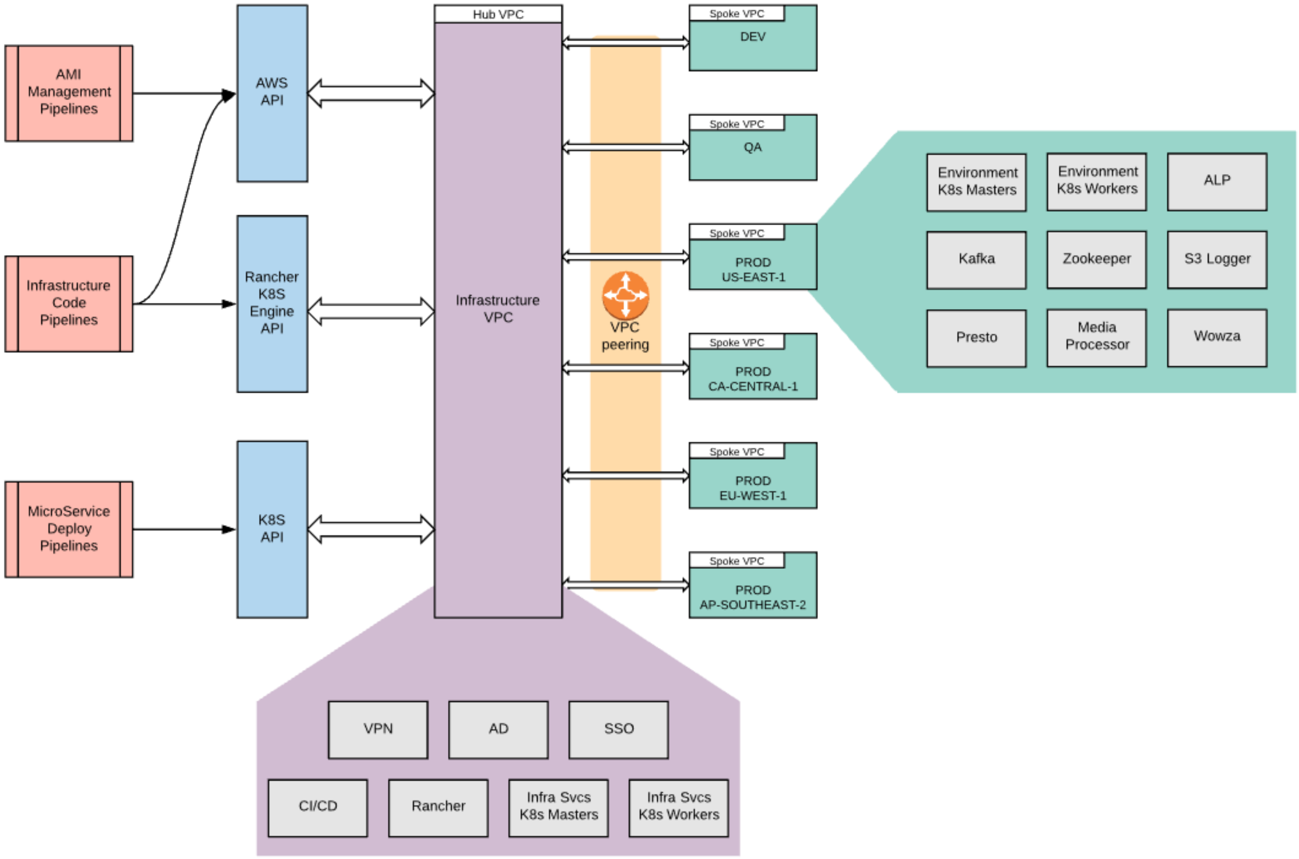

2.6.2 Compute

With over 60+ internal services powering the Echo360 platform, no single solution fits all compute scenarios. The Echo360 Engineering teams have developed infrastructure services that are designed to be highly available by scheduling services to run across all accessible availability zones in a region. This approach also provides excess capacity which we use to rapidly scale services in the event of traffic demand, an availability zone outage, or some other unexpected problem.

User-facing services, like our application servers, live streaming, and video transcoding infrastructure, include self-healing and auto-scaling capabilities which automatically detect and replace unhealthy nodes and provide on-demand upgrades and patching.

Our many back-end services are stateless Docker applications scheduled on Kubernetes clusters running on Amazon Elastic Compute Cloud (EC2). EC2 provides control over the instances allowing for latency optimization and high levels of redundancy. The containerized applications are self-healing and the clusters seamlessly rebalance these services as needed making it easy to deploy, manage, and scale.

2.6.3 Storage

The Echo360 platform uses three primary AWS services for persisting data:

- **Amazon Simple Storage Service (S3)** is utilized for object storage such as captures, uploaded files, and other similar content.
- **DynamoDB** is utilized for key/value storage and serves numerous services as part of the Echo360 Platform
- **MySQL Aurora** provides highly scalable and performant relational database functionality to various services that make up the platform including analytics.

These and many other AWS services share the same underlying data infrastructure within AWS that provide a high degree of security, durability, and availability of data. We encrypt data at rest in all of the above services, and when an item is written to one of these services it is actually written at least six times in three different physical datacenters before that write is acknowledged to our services.

2.6.4 Service Communication

Internal services communicate within the Echo360 platform using a combination of asynchronous messaging powered by Apache Kafka and standard REST-based HTTP calls. The connections between these services are encrypted in flight using SSL / TLS and any data at rest is encrypted.

2.6.5 Monitoring and Alerting

The platform is monitored and managed by the Echo360 Site Reliabiltiy Engineering team which has implemented off the shelf and in house services to monitor the health and security of the platform. Infrastructure monitoring services check the health of the entire compute platform several times a second ensuring critical, dependant services are functioning properly and are accessible. Every infrastructure component, system and service produces numerous health metrics which allow automated systems to take corrective action in the event of failure. These gathered metrics are reviewed and key metrics indentified as indicators of either an impending or occurring system degredation or outage. The Site Reliability Engineering team maintains a 24x7x365 oncall rotation with both a primary and secondary team member.

2.7 Echo360 Business Continuity and Disaster Recovery

Business Continuity and Disaster Recovery were driving requirements for our platform from the very beginning. Echo360 chose Amazon Web Services to provide our Echo360 SaaS solution because it has a proven track record of providing reliable services. Our platform technologies and patterns not only provide for incredibly high levels of data durability and availability, but also implement automated self-healing technologies that allow services to automatically recover in the event of a detected problem. These events are logged and investigated by our Site Reliability Engineering team who determines root cause and then prioritizes and coordinates with other engineering teams to fix the underlying problems.

Echo360 also maintains an **Enterprise Support Plan** with AWS which includes a 15 minute service level agreement response time and dedicated technical contacts available 24x7x365 who will guide us through any AWS outage or incident.

More information about our AWS Enterprise support plan can be found here:
https://aws.amazon.com/premiumsupport/plans/enterprise/

Echo360 is also an **AWS Education Partner** which requires internal certification and review by AWS to ensure compliance with the **AWS Well-Architected Framework** which confirms our platform follows best practices around data security, reliability, and performance.

More information about AWS Education Partner can be found here:
https://aws.amazon.com/education/partner-solutions/

More information about the AWS Well-Architected Framework can be found here:
https://aws.amazon.com/architecture/well-architected/

Echo360 has the readiness and capability to cope effectively with whatever major incidents and disasters occur, including those that were not, and perhaps could not have been, foreseen protecting both our owners and foremost our customers.

2.7.1 Business Continuity Design

BUSINESS CONTINUITY

The Echo360 solution and supporting infrastructure has been designed and engineered to provide resilience and is materially unaffected by most disruptions through the use of AWS as the backbone of our system. Automated self-healing technologies allow services to automatically recover and restore if they fail for an unexpected reason. Echo360 also has established a general business continuity plan as a last-resort response if resilience and recovery arrangements should prove inadequate.

2.7.1.1 Service Regions & Availability Zones

We deploy a specific customer in one of our offered Regions (United States, Canada, Europe, and Australia); these regions consist of two or more Availability Zones. Availability Zones are physically separated within a typical metropolitan region and are located in lower risk flood plains. In addition to discrete uninterruptable power supply (UPS) and onsite backup generation facilities, each are fed via different grids from independent utilities to further reduce single points of failure. Furthermore, Availability Zones are all redundantly connected to multiple tier-1 transit providers to ensure uninterrupted network connectivity. At all times, requests to our services are load balanced across all available Availability Zones in a Region. In the event of an Availability Zone failure, traffic is automatically routed to the remaining data centers.

2.7.1.2 Redundancy

All of our systems are tested for redundancy on an ongoing basis. As new internal services are built for the platform, the SRE team participates during the development process to provide insight and guidance to ensure the services being built are secure, scalable, and self-healing where appropriate. Following microservice design patterns, our platform consists of many smaller services meaning that if a single service were to go down, its impact would not be felt as severely as if the entire application went down. All of our infrastructure and services are redundant and self-healing supporting automatic recovery. In addition, all of our services are striped across AWS Availability Zones; therefore, in the event an Availability Zone went completely offline, service would continue uninterrupted. This ensures that customers are not affected by an Availability Zone outage.

2.7.1.3 Data Backup

Data is stored at rest in three primary datastores. S3 is used for object storage such as content and file uploads. All S3 data is versioned with non-current versions being maintained in Glacier (offline archival storage) for 90 days before automatic expiration. Data stored in DynamoDB or MySQL Aurora have two recovery mechanisms in place. The first is via a point-in-time recovery process that keeps track of changes in real time. This allows us to restore a specific table or database to any specific time in the last 12 days down to a minute resolution. We also create, at a minimum, once daily full snapshots of the data. We retain these full daily snapshots for one week, and then we keep a Sunday's snapshots for a year. These backups are tested periodically and are monitored by separate internally developed services to ensure that snapshots are being taken at the regular intervals. Large sudden changes in an individual tables backup size is flagged for review by the SRE team.

2.7.1.4 Infrastructure as Code

The Echo360 infrastructure and its configuration is created and modified following standards similar to software development with peer review and a combination of automated/manual tests. Infrastructure is managed by an open source tool called Terraform and changes are initiated by a pull request to a git repository which goes through automated policy review to ensure the change is inline with our standards. If inline with the standards, the system does a dry run of the change to show what would be achieved. This change and the associated output from the tool is then reviewed by at least one other person on the SRE team before it is merged into the system. These plans are manually reviewed before they are ran however as an extra precaution we have ensured these automated tools do not have the capability or even permissions to destroy data repositories like DynamoDB/RDS Databases, S3 buckets or any backups,

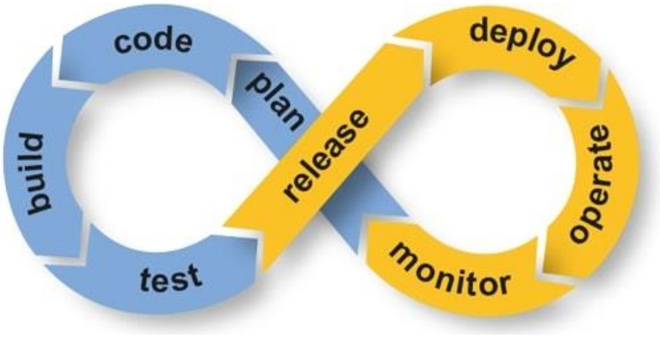

A change is then promoted through QA, Staging, and finally production environments. Each environment is automatically and manually monitored in accordance with our change control policies. This includes running extensive automated testing, manual testing, and real time feedback from our monitoring and alerting systems before the change is pushed forward.

Changes to the infrastructure are checked against policies that automatically prohibit certain insecure or suspicious actions and all of our AWS accounts are monitored in real time where potentially insecure or anomoulous changes are flagged and reviewed.

2.7.1.5 Immutable Infrastructure

We use an immutable infrastructure pattern to apply the same principals that we do with software development and the Infrastructure as Code approach. Using industry leading open source technology solutions such as Packer and Ansible allows us to define the “images” that are used to create an EC2 instance as code. Changes to these images are proposed with a pull request which triggers an AWS EC2 instance to be created from scratch, built with the configuration defined, and then automatically tested with a set of automated tests specific to the purpose that the image will serve to verify that the build process was successful. If this process is successful then an Amazon Machine Image or AMI is created from the EC2 Instance.

With the tests passing successfully and at least two members of the SRE Team signing off on the new AMI version the image is then promoted and deployed into our QA environments where extensive application level automated testing is performed to ensure the changes did not impact the environment or services.

This image is then promoted through our other environments following a process similar to the Infrastructure as Code process. We believe that this process is vital to ensuring the platform is performant, secure, and scalable that we use this process as our patch management strategy and we regularly replace all of our underlying infrastructure.

2.7.1.6 Deployment

Service deployments are promoted using the same method as Infrastructure as Code and Immutable infrastructure. Changes are triggered via an internally developed tool which monitors all systems for “state changes.” For example, a state change could be a new EC2 image version being deployed; a change to an AWS service configuration; or a newly deployed version of a service.

Any state change automatically triggers both application security scanning and automated testing to verify the functionality of the platform. If either of these tests fail, a rollback process is invoked, our engineering teams involved in the changes as well as our Site Reliability Engineering team is notified, and the changes will be blocked from being promoted to another environment.

2.7.1.7 Monitoring

All of our services and infrastructure are monitored in real time using a combination of both internally developed and third party tools. Our monitoring and alerting system Datadog ingests hundreds of thousands of AWS service, system, and service metrics with most coming in at a one second resolution. These metrics allow us to automatically perform actions such as looking for anomalous patterns that could indicate a potential bottleneck or impending outage, and are used during incident response and root cause analysis procedures.

For logging, we use a combination of open source, AWS, and internally developed services to record billions of logs a day which are ingested into an internal elasticsearch cluster. We utilize rule based alerting and anomaly detection to alert the SRE team of issues that require further investigation. Our various in-house services that monitor the platform in real time 24x7x365. For example, we have services that constantly monitor the ability for internal infrastructure hosts to communicate with dependant services and simulated user automated testing that exercises the system every 5 minutes testing core functionality. We evaluate this information, combined with high level metrics such as request latency, request error rates, and others to determine overall platform health.

We also utilize Datadog Application Performance Monitoring which provides end-to-end distributed tracing of requests in our microservices platform enabling us to quickly spot potential problems.

While a large segment of the metrics we gather have no alerts, hundreds do generate alerts. These metrics are monitored in real time with thresholds defined to indicate either Severity 1 or Severity 2 alerts. Severity 1 alerts are acted on and investigated immediately and will page the on call SRE 24x7x365. Examples of a Severity 1 alert might be an event that is about to impact the users ability to access the platform or any potential or suspected security related item.

2.7.1.8 Incident Response

Severity 1 alerts will trigger our Incident Response escalation policies and procedures. At any given point in time, three members of the SRE team are on call. A primary and secondary person as well as the team manager. We utilize a third party service to escalate incidents to the appropriate on call person in the schedule. This generally consists of a combination of a text message and a phone call to the primary on call person. If the primary person does not acknowledge the page within 15 minutes, the secondary on call person will be contacted in the same manner by both text and phone call. If this goes unacknowledged, then it escalates to the SRE team manager. Depending on the severity of the incident, the Primary may elect to respond and also contact additional resources as needed to resolve the incident.

2.7.1.9 Root Cause Analysis

Every Severity 1 and on call escalation triggers an automatic root cause analysis procedure. This extensive process will often involve multiple teams which work to identify the extent at which the incident impacted users, the origin of the failure, a detailed timeline of the incident, as well as documenting what went well, didn't not go well, and where we got lucky. Finally, we document corrective steps for the short and long term that are required to prevent this incident from impacting users again. These action items are then prioritized and assigned out to the appropriate engineering teams. This internal report is then reviewed by our support, product, and executive teams and a limited version is turned into a customer facing document.

2.7.2 Disaster Recovery

Echo360 is equipped with the business continuity plans and procedures to enable the recovery or continuation of our platform's infrastructure and systems following an unforeseen event, natural or human-induced disaster. While it is not possible to plan for every type of catastrophe, our disaster recovery and business continuity plans are integral parts of our overall risk management process. Our ability to recover our system includes automatic backup of critical systems and data, recovery with minimal user impact and flexible recovery options based on our platform design which includes, at a minimum, restoring individual services to reinstating the entire infrastructure.

3 SECURITY

Echo360 uses a risk management approach to develop and implement security controls, objectives, policies and procedures that address security and privacy objectives in conjunction with our business and operational considerations.

Echo360 is committed to ensuring the security and protection of the personal information that we process and providing a compliant and consistent approach to data protection. Our data protection program complies with existing law and abides by data protection principles. We are dedicated to safeguarding the personal information under our remit and in developing a data protection regime that is effective, fit for purpose and includes the development and implementation of data protection roles, policies, procedures, controls and measures to ensure maximum and ongoing security.

The Echo360 Information Security Program has drawn from relevant controls, objectives, standards and guidelines including fundamental guidance, procedures and commentary based on the following Information Security frameworks:

- **CSA** - **C**loud **S**ecurity **A**lliance
  A not-for-profit organization with the mission to promote the use of best practices for providing security assurance within Cloud Computing who provides guidance for critical areas of focus within the domain of cloud computing and has delivered a practical, actionable roadmap for organizations seeking to adopt the cloud paradigm. The CSA uses the NIST model for cloud computing as its standard and also endorses the ISO/IEC model within its domain framework.

- **COBIT** – **C**ontrol **OB**jectives for **I**nformation and Related **T**echnology
  An IT governance framework that allows managers to bridge the gap between control requirements, technical issues, and business risk. It was written by the Information Systems Audit and Control Association (ISACA) and is a supporting toolset that allows managers to bridge the gap between control requirements, technical issues and

business risks. COBIT enables clear policy development and good practice for IT control throughout organizations.

- **ISO/IEC 27000-series** – **I**nternational **O**rganization for **S**tandardization (ISO) and the **I**nternational **E**lectrotechnical **C**ommission (IEC)
  ISO/IEC 27000 series provides guidance for information technology security techniques and information security management systems.

- **(SSAE) No. 18** – **S**tatement on **S**tandards for **A**ttestation **E**ngagements
  A generally accepted auditing standard produced and published by the American Institute of Certified Public Accountants (AICPA) Auditing Standards Board. It prescribes three levels of service: examination, review, and agreed-upon procedures. It also prescribes two types of reports for reporting on an examination of controls at a service organization relevant to user entities' internal control over financial reporting: Type 1, which includes an assessment of internal control design, and Type 2, which additionally includes an assessment of the operating effectiveness of controls.

- **GDPR** – **G**eneral **D**ata **P**rotection **R**egulation
  A regulation in European Union law on data protection and privacy for all individuals within the EU. It also addresses the export of personal data outside the EU. The GDPR aims primarily to give control to citizens and residents over their personal data and to simplify the regulatory environment for international business by unifying the regulation within the EU.

3.1 Security Attestations, Certifications and Accreditations

Another advantage of building our platform on the AWS cloud is that it allows us to deliver a highly secure platform under the AWS shared responsibility model. We benefit greatly from the economies of scale that AWS data centers and networks operate at including Physical and Environmental controls as well 24x7x365 Network and Security operations centers. The AWS architecture is built to meet the requirements of the most security-sensitive organizations in the finance, government, health, and education industries. We recognize security and compliance is a shared responsibility and to that end, we strive to meet and exceed our responsibilities when it comes to the security of the platform and its data.

3.1.1 AWS Education Competency

Echo360 achieved AWS Education Competency status in July 2017. The designation recognizes Echo360 for technical proficiency and proven success building and scaling cloud-based solutions that support mission-critical workloads of educators across higher education. Achieving the AWS Education Competency differentiates Echo360 as an AWS Partner Network (APN) member that has delivered proven customer success providing specialized solutions aligned with AWS architectural best practices to support the academic experience of teachers and learners. The AWS Partner Competency Program has validated the partners have demonstrated success in providing specialized solutions aligning with AWS architectural best practices to help support teaching and learning, administration, and academic research efforts in education.

3.1.2 Echo360 Penetration Test Attestation

Nettitude Inc. conducts a yearly, in-depth Authenticated and Unauthenticated Web Application Penetration Test and cyber-security review of the Echo360 Platform. In their most recent assessment, Nettitude determined the security posture for the Echo360 platform to be STRONG indicating any identified vulnerability is not critical to the application. No Critical, High or Medium Vulnerabilities were reported. A copy of the Nettitude Attestation is available upon request.

3.1.3 Privacy Shield Program

Privacy Shield Framework

Echo360 is a member of the Privacy Shield Program. The Privacy Shield Principles comprise a set of seven commonly recognized privacy principles combined with 16 equally binding supplemental principles, which explain and augment the first seven. Collectively, these 23 Privacy Shield Principles lay out a set of requirements governing participating organizations' use and treatment of personal data received from the EU under the Framework as well as the access and recourse mechanisms that participants must provide to individuals in the EU. Once an organization publicly commits to comply with the Privacy Shield Principles, that commitment is enforceable under U.S. law.

3.1.4 AWS Security Certifications and Attestations

AWS is compliant with various certifications and third-party attestations. For more information on AWS compliance please see http://aws.amazon.com/compliance/.

3.2 Network Security

3.2.1 Secure Network Architecture

Network devices, including firewall and other boundary devices, are in place to monitor and control communications at the external boundary of the network and at key internal boundaries within the network. These boundary devices enforce the flow of information between network fabrics through rule sets, access control lists (ACLs), and device configurations.

3.2.2 Secure Access Points

To allow for more comprehensive monitoring of inbound and outbound communications and network traffic, AWS has strategically placed a limited number of access points to the cloud. In addition, AWS has implemented network devices that are dedicated to managing interfacing communications with Internet Service Providers. AWS employs a redundant connection to more than one communication service at each Internet-facing edge of the AWS network. These connections each have dedicated network devices.

3.2.3 Network Monitoring and Protection

AWS uses a variety of automated monitoring systems to provide a high level of service performance and availability. AWS monitoring tools are designed to detect unusual or unauthorized activities and conditions at ingress and egress communication points. These tools monitor server and network usage, port scanning activities, application usage, and unauthorized intrusion attempts. The tools can be used to set custom performance metric thresholds for unusual activity.

3.2.4 Distributed Denial of Service (DDoS) Attack

AWS Application Programming Interface (API) endpoints are hosted on large, Internet-scale, world-class infrastructure that benefits from the same engineering expertise that has built Amazon into the world’s largest online retailer. Proprietary DDoS mitigation techniques are utilized. Additionally, the AWS networks are multi-homed across a number of providers to achieve Internet access diversity.

3.2.5 Man in the Middle (MITM) Attacks

All data is encrypted in transit to prevent “man in the middle” attacks and eavesdropping. We also conduct server session verification to ensure browsers authenticated with the platform are verified users. Should this verification fail in response to some sort of session modification, the user would immediately be logged out.

3.2.6 IP Spoofing

Amazon EC2 instances cannot send spoofed network traffic. The AWS-controlled, host-based firewall infrastructure will not permit an instance to send traffic with a source IP or MAC address other than its own.

3.2.7 Intrusion Detection

Echo360 uses multiple cloud security platforms including AWS Guarduty, Lacework, and Datadog to provide real time monitoring and intrusion detection for our systems. These services monitor the platform infrastructure in real time looking for anomalous network, file, or API activities.

3.3 System Security

3.3.1 System Access

System access is granted only to the Site Reliabiliy Engineering team and is managed by our EC2 Instance build pipeline ensuring any access changes are highly visible and tracked. The team uses SSH keys to access EC2 Instances and keys are not allowed to be shared. All system access requires VPN access which is managed by Okta and secured with Multi-Factor Authentication. All commands and activities performed by the employee are logged and monitored in real time by various security products.

3.3.2 Vulnerability Assesment and Remediation

All systems are scanned regularly by Lacework for known vulnerabilities. Vulnerabilities that are identified are flagged and triaged to determine the potential attack vectors and the required steps to mitigate. System patches are pushed out as required following our immutable infrastructure and deployment patterns to ensure our ability to rapidly patch systems at scale.

3.3.3 System Security Monitoring

Echo360 runs Lacework and Datadog on each system. Each system is monitored in real time providing behavioral analysis and looking for signs of compromise such as user logins, TTY timelines, file access, priviledge escalations, login attemps and failures, suspicious commands, network connections, process monitoring, as well as scanning the system regularly for vulnerabilities.

3.3.4 AWS Security Monitoring

The AWS configuration of services and the overall accounts are monitored in real time using Lacework, AWS Config, and AWS Guarduty. This real time monitoring flags unexpected service or configuration changes that could potentially be exploited if discovered.

3.4 Platform Security

Echo360 enforces stringent application security standards. Not only do we perform extensive code reviews, static code analysis, and automated security scanning, but we engage a third party penetration testing company to ensure we have the most secure system possible. The Echo360 platform was built from the beginning to be a multi-tenant application and to that end, we have implemented the strictest safeguards around data privacy.

3.4.1 Data Security

The Echo360 platform architecture allows for segregation at the institution level throughout the platform. Every institution is assigned a randomly generated unique identifier. Our databases and services are partitioned using that identifier. We also limit the expanse of sensitive data attributes, such as PII, to our various internal services that power the system and instead rely on randomly generated unique identifiers to reference objects inside the platform. These are resolved “just in time” either when rendering a display to a user or by an authenticated, internal service asking the source service to resolve the information. This pattern reduces the amount of data that needs to be replicated and moved around while also improving the overall security as we limit the locations that sensitive information is stored.

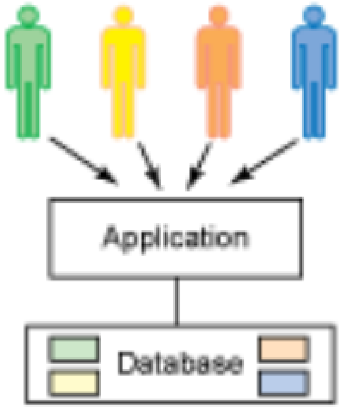

All data in transit, both end-user to Echo360 and within the Echo360 platform, is encrypted over SSL/TLS. Regardless of the data containing PII or an obscure identifier, all data is encrypted at rest in all of our systems and data stores. We leverage the AWS Key Management Service (KMS) which uses FIPS 140-2 validated hardware security modules to generate and protect keys. KMS keys are not shared or distributed outside of the region where they are generated. Access to these keys regardless of the source is logged and monitored in realtime using AWS CloudTrail and Lacework.

3.4.2 Identity Management

We optionally support integration with an institution’s identity services in order to provide single sign on to the Echo360 web application as well as all mobile and desktop applications. We support authentication with local Echo360 user accounts as well as SSO via LTI when integrated with an institution’s LMS, Active Directory, ADFS, Shibboleth, Google Identity, or any other SAML 2 compliant provider. We are listed as a service provider with inCommon, and have numerous customers using the examples stated.

3.4.3 Role Based Access

Roles are assigned to a user during the initial data load process. Roles can be inherited from the LMS or adjusted by an administrator via the Echo360 administrator web interface.

3.4.4 Automated Security Scanning

The Echo360 platform utilizes various third party security scanners that perform daily and on demand authenticated security scanning of the platform. These scans are automatically triggered upon changes to the environment. These scans are performed inside our own security tenant in each regional deployment of the platform and they test for vulnerabilities such as SQL Injections, Reflected XSS, Local File Inclusions, Remote File Inclusions, and more.

3.4.5 Platform Penetration Testing

Echo360 has consistently performed and intends to continue to perform third party penetration testing on an annual basis and will remediate any deficiencies found. These multi-week engagements utilize a penetration testing company who attempts to exploit all aspects of the platform with heavy focus on data security, tenant isolation, user priveledges, and authentication and authorization.

3.5 Corporate Security

3.5.1 Physical Security

Echo360 utilizes best practices to authorize, monitor, and control all methods of access to its facilities and all information systems. Access Control is divided into three categories: external perimeter facility access, interior access and computer security. To the maximum extent possible, authorization and facility access control practices address the following:

- **Exterior doors:** perimeter access is achieved through any door that opens to the exterior of the Echo360 office space through a traditional key/lock or electronic card reader. External perimeter access is maintained via building time schedules (7:00 AM – 7:00 PM).
- **Interior doors:** interior access control is determined by the needs of the individual on a case-by-case basis. In certain circumstances, such as in restricted areas where a higher level of security is necessary, i.e. spaces where equipment is stored, installation of traditional locks with key or electronic card reader restricted to required employees achieve an extra level of security.
- **Computer security:** automated and non-automated enforcement practices have been implemented to safeguard computers and to prevent company and customer data from becoming vulnerable and being compromised.

3.5.2 Device Security

Echo360 employee systems contain a Trusted Platform Module (TPM) which is a secure cryptographic processor that provides the system with a secure mechanism to protect encryption keys. Leveraging the TPM, all Echo360 employee systems have full disk encryption enabled as well as a sophisticated antivirus and threat detection product suite from Sophos installed. These systems are centrally monitored,and updated by our Corporate IT team.

3.5.3 Corporate System Access

Okta Identity Cloud provides one trusted platform to secure and manage user identities in our organization. We utilize Okta to manage access to our corporate systems with single sign-on service and multi-factor authentication for all employees and contractors throughout our organization.

3.5.4 Corporate Device Management

We utilize a corporate-wide mobile device management solution to provision and deploy our Apple equipment issued to employees. This includes setting and enforcing password requirements with mandatory, timed auto-locking; managing application configuration profiles and use policies; administering operating system and application updates to protect devices and secure user data; and remotely locking and/or wiping a lost, stolen, or otherwise compromised device.

3.5.5 Data Handling and Training

Access to customer data is highly restricted to a very limited number of individuals in the company such as our Customer Support and SRE teams. These employees who may come in contact with customer data undergo additional training to ensure they are familiar with our information handling and data protection policies and procedures. Access is logged, and we require a valid reason for accessing the data, for example, investigating a customer reported bug or issue. This access is reviewed by management.

3.5.6 Onboarding

All Echo360 employees regardless of position in the company undergo a criminal background check as a condition of employment.

4 DATA RETENTION

The standard retention period for output products is currently two (2) years; however, customers can purchase an extension to the retention period that will provide an additional two (2) years for a total of four (4) years of content retention.

The Echo360 platform automatically deletes Source Media Files as they reach 30 days and only if the source media was successfully transcoded. Source file retention does not affect the output media files (those used for user download and playback), the ability to edit and re-use media, playback of media, or any other functionality in the platform.

Source Media Files are the original media files uploaded from an Echo360 hardware capture device (PRO, POD, or SCHD) or Echo360 capture software (PCAP, CCAP, or Universal Capture.) Source Media Files are uploaded to the Echo360 platform and are transformed (transcoded) into the output products found in the Echo360 Classroom and embedded media player. This does not currently include media uploaded directly to the platform, such as video or presentation uploads to an instructor or student library.

5 SERVICE LEVEL AVAILABILITY

5.1 Service Availability

Echo360 shall use commercially reasonable efforts to ensure the Software Service is Available twenty four (24) hours a day, seven (7) days a week with a targeted uptime of 99.5% per calendar month. The Echo360 Service shall be considered “Available” when an Authorized End User request receives a response within thirty (30) seconds. “Unavailable” means an Authorized End User cannot access the Service within thirty (30) seconds due to hardware failure or sustained latency within the Amazon Web Services (AWS) facility Echo360 uses to deliver the Software Service. Notwithstanding the foregoing, the Availability of the Service shall be determined without regard to any (i) network, AWS solution or Echo360 software scheduled maintenance, (ii) the inability of a user to connect with the Service due to Internet or telecommunications problems outside the control of Echo360, or (iii) Force Majeure. For purposes of this section, “Force Majeure” means causes beyond Echo360’s reasonable control, including without limitation, acts or omissions of government or military authority, acts of God, materials shortages, transportation delays, fires, floods, labor disturbances, riots, wars, terrorist acts or inability to obtain any export or import license or other approval or authorization of any government authority.

5.2 Service Availability Measurement

“Service Availability” will mean, with respect to any particular calendar month, the ratio obtained by subtracting Unscheduled Downtime during such month from the total time during such month, and thereafter dividing the difference so obtained by the total time during such month. Represented algebraically, System Availability for any particular calendar month is determined as follows:

$$\textbf{Echo360 Service Availability (\%)} \ = \ \boxed{\frac{100 \ \text{X} \ (\text{TMH} – (\text{TMD} – \text{SD}))}{\text{TMH}}}$$

TMD = Total Monthly Downtime or total of all downtime in a month (measured in hours.) “Total Monthly Downtime” is deemed to include all hours in the relevant calendar month to the extent such hours are included within the Term of Agreement
SD = Schedule Downtime in the given month, not to exceed 3 hours in any calendar month
TMH = Total Monthly Hours in each calendar month (e.g., 30 days X 24 hours per day = 720 hours)

5.3 System Uptime

“System Uptime” is defined as the total amount of time during any calendar month, measured in hours or fractions thereof, during which Echo360’s Services are available to Customer. System Uptime is measured within Echo360’s operating environment and is not reduced by any periods in which Customer requests fail to reach the Echo360 operating environment.

5.4 Scheduled Downtime

“Scheduled Downtime” is defined as the total amount of time during any calendar month, measured in hours or fractions thereof, during which an Authorized End User is not able to access Echo360’s services due to planned system maintenance performed by Echo360 or AWS. Echo360 updates the Service on a two week cycle – Scheduled System Maintenance. These updates usually entail no disruption of Service Availability or “zero system downtime”. Regardless, Echo360 will exercise reasonable efforts to perform Echo360 service scheduled maintenance between the hours of 8:00 pm and 7:00 am Eastern Standard Time for the U.S. and Canada Regions; between the hours of 8:00 pm and 7:00 am GMT time for the EU – Ireland Region and between the hours of 8:00 pm and 7:00 am Australian Eastern Standard Time for the Asia Pacific – Sydney Region. Echo360 has sole discretion to perform updates when there is no required disruption of Service Availability. These updates, enhancements and bug fixes will be documented in release notes. Refer to section *Customer Communication* for more information.

If there is the need for Scheduled System Maintenance that requires a temporary disruption of the Service, Echo360 will provide Customers with five (5) business days advance notice. Echo360 will attempt to schedule these maintenance activities for evenings on the weekend not to exceed 3 hours per month. Echo360 in its sole discretion may disrupt the service for Unscheduled System Maintenance if it is the only way to address a persistent degradation of performance, and in that event will attempt to notify customer in advance if possible. Such Unscheduled Maintenance will be counted against the uptime guarantee.

5.5 Unscheduled Downtime

“Unscheduled Downtime” is defined as the total amount of time during any calendar month, measured in hours or fractions thereof, during which Echo360’s services are unavailable to an Authorized End User, other than Scheduled Downtime, as defined above, and excluding periods of unavailability resulting from reasons outside Echo360’s control (e.g., failure of customer’s internal systems or issues outside the managed environment itself, such as Internet or telecommunications failures, act of government, flood, fire, earthquake, civil unrest, act of terror, Third Party Application, or denial of service attack or “Force Majeure” events.

Echo360 in its sole discretion may take the service down for Unscheduled Maintenance and in that event will attempt to notify customer in advance in accordance with the Notice section set forth in Terms of Agreement. Such unscheduled maintenance will be considered Unscheduled Downtime for calculating Echo360 Service Availability.

Echo360 will monitor Unscheduled Downtime. Customer may also notify Echo360 in the event Unscheduled Downtime occurs by opening a ticket via the Customer Portal. Unscheduled Downtime will be deemed to begin when Echo360 receives accurate notification thereof from customer, or when Echo360 first becomes aware of such Unscheduled Downtime, whichever first occurs.

5.6 Media Turnaround Time

“Media Turnaround Time” is defined as the difference in time between the time that the raw media arrives at the AWS data center and the time the processed media is made available to an authorized end user at a specific customer.

Echo360 will undertake commercially reasonable measures to ensure that Media Turnaround Time is 12 hours per one hour of recorded media, provided that any failures occurring as a result of (i) Customer’s breach of any provision of this Agreement; (ii) non-compliance by Customer with any provision of this Attachment II; (iii) incompatibility of Customer’s equipment or software with the Echo360 Services; (iv) poor or inadequate performance of Customer’s systems; (v) Force Majeure Events or (vi) any matters outside Echo360’s reasonable control, shall not be considered toward any reduction in Media Turnaround Time measurements. In no event will Echo360 have responsibility for any inability of Customer to access the Echo360 Services due to Internet or telecommunications failures, failures of Customer’s internal systems, or due to any other issues occurring outside the managed environment itself.

5.7 Customer Requirements

5.7.1 Minimum System Requirements

The service standards set forth in the Terms of Agreement assume that the customer and/or its authorized end users, as applicable, meet Echo360’s published minimum system standards for compatibility, (e.g., compatible browser software, published LTI standard for LMS integrations, etc.) for access to Echo360 services.

5.7.2 Additional Customer Obligations

Except as otherwise agreed between the parties, the customer is responsible for (i) maintenance and management of its computer network(s), Internet connectivity, third party software systems, Web site(s), classroom technology, and any equipment or services related to maintenance and management of the foregoing; and (ii) correctly configuring customer’s systems with Echo360 service account(s).

5.7.3 Non-Performance by Customer

The obligations of Echo360 set forth the Terms of Agreement will be excused to the extent any loss of Service Availability is the result in whole or in part from customer’s failure(s) to meet the foregoing requirements.

6 SUPPORT

6.1 Standard Support Program

The Standard Support Program is designed to help you maintain and support all of your Echo360-licensed solutions. All customers with an active license/contract or maintenance agreement will receive the following:

6.1.1 Designated Customer Contacts

With the Standard Support Program, you may designate a list of primary contacts to communicate with Echo360 Technical Support and may change the designated contacts for this service at any time.

6.1.2 Unlimited Incidents

Under the Standard Support Program, you'll have access to support for an unlimited number of incidents. Support incidents are managed with **Zendesk**, an industry leading incident management system.

6.1.3 Software Updates

The Standard Support Program offers free software updates for licensed Echo360 products. Customers with active licensing and support contracts with Echo360 are entitled to hotfixes, service packs, and major product releases.

6.1.4 Return Materials Authorization

Should you experience an issue with an Echo360 hardware device currently under warranty and we confirm that the unit is faulty; we will replace that unit at no cost to you.

6.1.5 Echo360 Service Status

Echo360 provides a tool enabling all customers to track active service degradations and disruptions. An Echo360 cloud-based active learning platform service degradation is defined as an impaired, but functional service. A service disruption is defined as an interruption in service functionality. Customers may go to status.echo360.com to see any active service disruptions or degradations as well as subscribe to receive notification of any future service disruptions.

6.1.6 Self-Service Resources

From getting started with your Echo360 deployment to learning the latest product features, we strongly believe in providing high quality support resources to enable your staff to be confident and successful users of our solution. By clicking Help from within the platform or accessing the Echo360 Resource Center at https://support.echo360.com, you have access to an extensive library of up-to-date product support content for administrators, instructors and students. From account and navigation information to onboarding resources, training videos and release notes, these resources offer detailed guidance on every aspect of our solution.

6.2 Target Support Response Times

All service requests logged with support are assigned a severity level from System Unavailable to Nominal Business Impact. The Customer determines the initial severity level when placing a request for assistance. Severity levels may be changed by an Echo360 Support Engineer after initial contact and assessment of the issue providing the Customer is in agreement with the change. The following table defines the severity levels and the targeted Initial Response Times. Target Resolution Times are based on commercially reasonable efforts and assumes Echo360 has the authority to resolve key issues. It is **critical** to clearly explain the business impact of your issue when contacting the Support Center.

Severity Level	**Severity 4 NOMINAL**	**Severity 3 MINIMAL**	**Severity 2 SIGNIFICANT**	**Severity 1 CRITICAL**	**Severity 1 UNAVAILABLE**
Description	No impact to product usage or customer's operations. Minor problem, question or enhancement request that does not affect the service function.	Non-critical functionality impacted resulting in a minor or intermittent loss of service operation.	Specific service feature is unavailable, but a workaround exists and the majority of the service is still useable.	Service is operating, but performance is severely degraded. Important service features are unavailable with no acceptable workaround.[1]	Service access not available within thirty (30) seconds due to hardware failure or sustained latency within the AWS facility.[1]
Target Response Time	2 Business Days	1 Business Day	4 Business Hours	2 hours (7/24/365)	2 hours (7/24/365)

[1] Level 1 System Unavailable and Critical service requests cannot be logged via telephone or email, they must be logged via our Resource Center at https://support.echo360.com

For more information on Echo360's support response times, please refer to Echo360's Support Case Severity Descriptions: https://help.echo360.com/hc/en-us/articles/360035038932

6.3 Issue Reporting

Your issues can be reported via our Resource Center, by email or telephone. For the fastest response time, issues should be reported via our Resource Center. You may also report support-related issues via email or telephone during normal business hours, excluding Echo360 company holidays. If you leave a voice mail, that message will be reviewed during normal business hours. Upon reporting your issue, you will receive an email reply that includes a case number for future reference.

6.3.1 Online

For the fastest response time, please use our Resource Center:
https://support.echo360.com

This method immediately injects your support request into Echo360's support queue for triage and assignment. You will receive an automated reply that includes a case number for future reference.

6.3.2 Email

You may also send any support related inquries via email to:

support@echo360.com

This will automatically generate a Severity 3 support case. Once you've reported your issue, you will receive an automated reply that includes a case number for future reference.

6.3.3 Telephone

Telephone support is available during business hours, excluding Echo360 company holidays. If you leave a voice mail, that message will be reviewed immediately during normal business hours.

Region	Hours of Operation	Phone Number(s)
North America **South America**	Monday – Friday 8AM – 8PM ET	+1 703 667 7500
Europe **Asia**	Monday – Friday 9AM – 5PM GMT	+020 3026 3473
Australia **New Zealand** **Asia**	Monday – Friday 8AM – 5PM AEST/AEDT	Australia: 1300 324 600 NZ: 0800 324 600 Outside Australia: +61 6180 2799

6.4 Support Escalations

All Echo360 cloud platform customers receive 24/7 support for Severity 1 issues.

To provide this service, Echo360 levearges both a follow-the sun support model and a paging system to alert our support staff during off hours. A special workflow has been established through the Resource Center to manage these types of issues. The Resource Center is functionally linked to our follow-the-sun process and automated paging system and as such, is the only trigger for off-hours assistance. If for any reason, an issue is not being resolved in a way that meets your business needs, an escalation path has been developed to ensure a timely and satisfactory resolution. This is initiated by contacting your Customer Success Manager who will ensure your issue is resolved accurately and in a timely manner.

6.5 Customer Communication

Echo360 will provide multiple channels of communication to the Customer to ensure proper information is available to best manage the Echo360 Software as a Service purchased.

6.5.1 Product Enhancements and Bug Fixes

Echo360 regularly updates the Service on a two week cycle. These updates include bug fixes, service optimizations and feature additions/enhancements. Echo360 has sole discretion to perform these updates when there is no required disruption of Service Availability. To keep customers informed of pending Product Enhancements, we invite all active customers to participate in our annual Active Learning Conferences held in the U.S., U.K. and Australia, as well as quarterly Product Roadmap Reviews. Led by our Product Management team, these sessions provide details on new features or capabilities that are upcoming. Echo360 also provides detailed release notes several days before the Service is updated via the Resource Center and to individuals who subscribe to these communications.

6.5.2 Scheduled System Downtime Notification

Echo360 updates the Service on a two week cycle – Scheduled System Maintenance. These updates usually entail no disruption of Service Availability or “zero system downtime”. If there is the need for Scheduled System Maintenance that requires a temporary disruption of the Service, Echo360 will provided Customers with five (5) business days advance notice. Echo360 will attempt to schedule these maintenance activities for evenings on the weekend not to exceed 3 hours per month. Echo360 in its sole discretion may disrupt the Service for Unscheduled System Maintenance if it is the only way to address a persistent degradation of performance, and in that event will attempt to notify Customer in advance, if possible. Such Unscheduled Maintenance will be counted against the uptime guarantee.

6.5.3 System Unavailable Reporting

If a service disruption impacting many customers in a region occurs, the Echo360 operations team will find out either through an automated alert or a customer report. The Operations team will immediately assess the issue to determine if it is system wide or region specific. Echo360's goal is to have a customer-facing status page update no later than two hours after Echo360 identifies or is notified of the incident. During a disruption, Echo360 will keep this status page updated at least once per hour until resolution. A final incident report with root cause analysis will be available within 5 days. Customers can subscribe to the Echo360 status page to receive alerts of new disruptions by email, SMS, webbook, or RSS feed.