
Reliability

Can Ex Libris guarantee a high degree of availability and response time?

Ex Libris' standard SLA commitment is to deliver service availability of at least 99.5%, measured over any calendar year. The cloud environment is monitored 24x7, with staff attending to all issues in real time. Based on our experience with thousands of institutions deployed in our cloud environment, a number of which are large consortia, actual system availability is significantly higher. In practice we have constantly exceeded the annual 99.5% uptime SLA since the first day of going live.

To deliver this performance, all aspects of the physical infrastructure are designed to have no single point of failure in any of the cloud layers: electricity, access to internal networking and switches, load balancers, firewall and storage, as well as edge routers and multiple Internet Service providers (ISPs) to the data center, ensuring high availability and accessibility to the Alma service in the cloud.

Ex Libris believes in full transparency and post all our Uptime Reports on the web – you can review quarterly reports from the past 3 years including Root-Cause-Analysis reports for downtime events at http://knowledge.exlibrisgroup.com/Alma/Product_Materials/Uptime_Reports

Is it possible to view uptime reports over a period of time (years)?

Uptime reports for our data centers from 2013 onwards can be found at http://knowledge.exlibrisgroup.com/Alma/Product_Materials/Uptime_Reports

Uptime has consistently exceeded our SLA commitment of 99.5% availability.

How is uptime calculated?

Uptime is based on the following calculation (as defined in the Alma Service Level Agreement):

"Uptime" means the total period in minutes during which the Service is available for access and use during this period.

"Uptime Percentage" means Uptime expressed as a percentage, calculated in accordance with the following formula:

$$\text{Uptime Percentage} = X / (Y - Z) \times 100$$

Where:

X = Uptime

Y = Last 12 months period

Z = The duration (in minutes) of any SLA Exclusions* during these 12 months

*SLA exclusions are defined in the contract SLA (e.g. Scheduled maintenance, etc.)

How is a high degree of availability and response time achieved?

All of Alma's components are redundant and provide high availability:

- All servers support active-active fault tolerance;
- Database components have automatic fail-over;
- Our data center is planned at any given time to provide at least over 15% of the required capacity. We maintain stand-by servers ready in case of multiple-server failure;
- The Ex Libris data center maintains a clear mitigation plan for any malfunction scenario (hot standby\cold standby, etc.); and
- The Ex Libris data center provides 24x7 support for all hardware components with our vendors, with an SLA for replacement hardware on site when needed.

Recognizing that even the most resilient system may occasionally experience downtime, Ex Libris operates a 24x7 hub (service center) that provides first-line support for any System Down event. The 24x7 hub team focuses on 24x7 monitoring, troubleshooting, correcting and/or escalating problems, assuring that issues are handled in accordance with our SLA and with the highest level of technical performance and quality. The team has developed rigorous procedures that are employed immediately in the event of a down-time alert.

The 24x7 hub team is also responsible for background maintenance of servers, operating systems, network devices such as switches and routers, security devices (e.g., firewalls, IPS/IDS), storage area networks, backup systems, disaster recovery environments, server side performance, end user browser performance, and all background batch jobs.

As part of our ongoing process of evaluating our high availability and improving our ability to handle potential issues, we continually monitor Alma's performance. For example, when we identify batch processes that are overloading the institutions' environment, we move them into a dedicated server area and by doing so, improve performance and reduce any risks presented by the additional loads.

As a cloud-based SaaS solution, Alma's actual user experience is based on several factors. Among these factors are the throughput of the Ex Libris Data Center, the server's performance, and the local institution network and bandwidth. The Ex Libris Cloud Data Center utilizes 1G bandwidth as its backbone, and works with multiple ISP vendors at every point in time. As with every SaaS vendor, we measure the performance of our servers; we do not, however, have control of the 'last mile' at the institution level. We can report that so far our customers report a high performance level for all operations that met web application expectations. We also monitor our application and server throughput 24/7 in order to make sure that transactions are handled and sent to the user browser as expected.

Ex Libris also has a web site which presents up-to-date system status for all Ex Libris multi-tenant environments. With the System Status site, Ex Libris customers can view the current status of their service, and they may sign up for email alerts when there are interruptions to the service. The site offers:

- Live and historical data on system status
- Scheduled Maintenance notifications
- An option to sign up for email alerts regarding interruptions to the service

The site, which is located at <http://status.exlibrisgroup.com/>, does not require login.

How does Alma monitor and report on system reliability and performance?

Deploying Alma as a SaaS offering, combined with the high uptime we commit to our customers, makes Alma monitoring a fundamental part of our operation. We separate monitoring into two main parts:

- 24x7 monitoring for real-time issues such as identification and resolution; and
- Proactive trend analysis and health check validation.

24x7 monitoring

The 24x7 hub Team provides 24x7 first-line support (in tandem with regular support personnel) for all cloud production environments, using advanced monitoring and dashboard capabilities.

Alma is designed with multiple layers of real-time monitoring of all facets of the solution. These layers include:

- Hardware (servers, storage, networking, FW, LB and all other devices);
- Operating system level monitoring;
- Database monitoring;
- Application level logs reflecting errors including database issues, software issues, and business level issues;
- Application performance as measured in the servers;
- Internet performance; and
- Security.

All components (physical and logical) are fault-tolerant with redundancy. All critical components work in active-active mode, such that loss of a component is invisible to users. In some cases, failover may result in very short outages (of a few minutes), and the return to full availability is completely automatic.

Monitoring teams are equipped with processes for provisioning decisions based on application performance monitoring and degradation of service scenarios.

Proactive trend analysis and health check validation

This type of monitoring is responsible for tracking system stability and forecasting provisioning needs to meet the system's uptime SLA:

- Run operational reports to track abnormal behavior and growth trends such as :
 - Total daily transactions
 - Per instance transactions/sec
 - Min hour transactions /sec
 - Max hour transactions /sec
 - Average transactions /sec
- Application health checks:
 - Validity of indexes, triggers, stored procedures
 - System job failures
 - External interfaces failures

- Running the operational BI
 - Provide uptime and performance reports for our customers
- Execute provisioning of new servers based on the analysis of system usage and capacity, taking into consideration lead time and required over-provisioning for meeting the uptime and performance SLA.

In addition, Ex Libris provides users with web access to view system status, including the ability to register for email alerts in case of performance issues or a system down event. The status page is available to all via the below link: <http://status.exlibrisgroup.com> and provides a live status on the system, with alerts in case of performance issues, planned maintenance or of a system down event.



How does Ex Libris record downtime and service breaches?

As a multi-tenant service, Ex Libris tracks downtime and SLA breaches for all the customers sharing the data centre. We provide a system status web page <http://status.exlibrisgroup.com/> where we provide alerts of downtimes and communicate any other service disruptions or scheduled maintenance events. In the event of downtime, Ex Libris updates the System Status Page on an hourly basis until resolution of such downtime event. In addition – we publish on a quarterly basis, an Uptime report that is available for all customers tracking the performance of our systems against the SLA. This Uptime report also share information about the maintenance period and any downtime events experienced during the quarter.

How are fixes and reported issues prioritized?

Ex Libris' response to cases is governed by the following SLA matrix:

Response Level?	Description	Initial Response
I	The Service is not available	1 hour
II	An inoperable production module	2 hours
III	Other production performance related issues, typically a module feature working incorrectly	1 business day

Response Level?	Description	Initial Response
IV	Non-performance related incidents, including: general questions, requests for information, documentation questions, enhancement requests	2 business days

How are releases, service management etc. managed in Alma?

According to our standard SLA, there is a standard maintenance window (Saturday 20:00 to Sunday 06:00 CET), though typically we use a much shorter window for maintenance and release upgrades. If we intend to use any part of this window, we notify our customers via email, and through a message that appears upon login to the application. In practice, this window rarely used for anything other than applying the monthly Alma upgrades. Ex Libris commits to, and has so far achieved 99.5% availability or higher, outside the above window.

Note also the system status page at <http://status.exlibrisgroup.com> Since Alma's launch in 2012, there have been monthly releases, with each new update released onto customers' sandbox environments two weeks before being released into the production environment.

Ex Libris follows the strictest possible change control in its cloud processes. Every request to change or do maintenance work on the cloud environment goes through approval processes and is audited and logged.

Moreover, all our software testers are ISTQB (International Software Testing Qualifications Board) certified. Ex Libris employs a Quality Assurance Director to lead quality assurance activities and ongoing improvement processes. Appropriate quality management is ensured by an application life cycle management methodology, which includes Quality Assurance during the development process and as part of our Quality Assurance methodology. During the Development process, unit testing is planned and conducted for every new development or bug fix. The QA team is involved in the planning and approval of each of the unit tests that have been planned. Before release for system tests, the Development team conducts a full cycle of sub-system tests, which includes testing every new development in an integrated environment and performing end-to-end system testing. The QA team is involved in the development process from the design stage, thus allowing for the preparation of detailed testing before the product is delivered to QA for system testing. QA team analysts are involved in the business design process and understand the business flow and product implementation at customer sites; this allows for a high guarantee of product quality. Collaborative testing with our customers is performed for major releases at the early stage of the system tests. Any major release is released first the beta customers. QA team records customers' testing experience and continually extends system tests scope.

Alma's monthly release occurs on the first Sunday of each month. Two weeks prior to the release, the release is available to customers on their Sandbox environment. This allows customers to try the new features prior the actual release (note, that - as Alma and Primo are SaaS solutions - there can be no user acceptance / non acceptance of any release.). It is important to mention that each release does not change any of the customers' data or configuration.

The two weeks before we also communicate to our customers through our mailing lists regarding the upcoming release. The communication includes:

1. Release highlights email listing the main new features of the release
2. Release notes – complete documentation of the entire release's new developments
3. List of resolved issues reported by customers

Our CRM uses automated clocks that alert both analysts and team leaders in the responsible support teams in the case that any target is at risk of not being met.

Ex Libris performs on-going analysis and reporting of whether we're meeting our SLA response targets. In case we fail to meet the target, the team leader will set corrective actions to ensure targets are met in future. Also there is a support escalation policy that customers can use to escalate specific calls.

Ex Libris performs regular support phone calls with customers. During these calls, they are asked for an "Overall satisfaction with support"-rating. In addition, we randomly call customers after closing support cases to perform a survey regarding the support experience the customer had.

Ex Libris' support team can be reached via our 24 x 7 Hub at any time, using various communication channels: CRM ticket, phone call, or emails. In cases of issues that requires the involvement of experts, we use our on-call staff; they are always available in order to start working on such issues within few minutes


What redundancy and load balancing features are in place?

Alma has been built with high-availability capabilities to ensure that the system can handle any type of disruption. This approach applies to all components of the system, and complies with the fundamental guidelines of business continuity:

- Full redundancy
- Load balancing
- Failover

To deliver this, the Alma environment has been designed so that all aspects of the physical infrastructure have no single point of failure. In all the cloud infrastructure layers, a fault-tolerant capability with failover mechanism ensuring true high availability has been built. This includes:

- Redundant power system including generators
- Redundant cooling/humidity control system
- Network bridges resiliency
- Multiple Load balancers
- Firewall redundancy
- Storage redundancy
- Edge routers redundancy
- Multiple Internet Service providers (ISPs) to the data center, ensuring high availability and accessibility to the Alma service in the cloud.



In addition, all of Alma's components are redundant and provide high availability:

- All servers support active-active fault tolerance
- NetApp clustered storage with snapshot backups
- Database components have automatic fail-over utilizing cluster solution
- Firewalls operate in an active-active mode ensuring availability with security
- The Ex Libris Data Center is planned at any given time to provide over 15% of the required capacity. Stand-by servers are maintained and ready in case of multiple-server failure
- The Ex Libris Data Center maintains a clear mitigation plan for any malfunction scenario (hot standby/cold standby, etc.)
- The Ex Libris Data Center provides 24x7 support for all hardware components with our vendors, with an SLA for replacement hardware on site when needed.

Total views:

3798