

Infrastructure & Reliability Overview

Last Modified on 12/05/2023 4:25 pm EST

Overview

Resolver's technical infrastructure and product architecture are designed to meet customers' expectations for security, reliability, confidentiality, and integrity. Resolver is built on a technical architecture that incorporates redundancy and high availability. We regularly perform stress tests to discover and address points of failure and monitor every aspect to ensure peak performance using the below processes. For more information on Resolver's current status and history, see the [Status Dashboard](#) article.

Service Availability Targets

Resolver will use commercially reasonable efforts to make Resolver available 99.9% of the time during a 24/7 period, except during scheduled maintenance or any unavailability caused by any of the provisions set out in Section 17(i) of the [Terms of Service](#).

Scheduled Maintenance

Resolver conducts routine environmental maintenance monthly, which involves downtime for Resolver. We schedule this in advance to minimize disruption.

Resolver will communicate scheduled maintenance at least 72 hours in advance to authorized contacts via email notification, except for emergency maintenance. Emergency maintenance notifications will go to authorized contacts only when deemed necessary, with notification timing based on the severity of the maintenance.

Customers are responsible for providing us with current contact information for maintenance notifications. Changes or additions to contacts can be submitted to [Resolver Support](#).

Outage Reporting

Resolver will notify authorized contacts by email when an issue results in more than one hour of service unavailability.

Data Backup

All customer data in Resolver is backed up continuously and retained for 30 days. Requests to retrieve customer data to recover from customer-initiated actions may result in costs being chargeable to the customer.

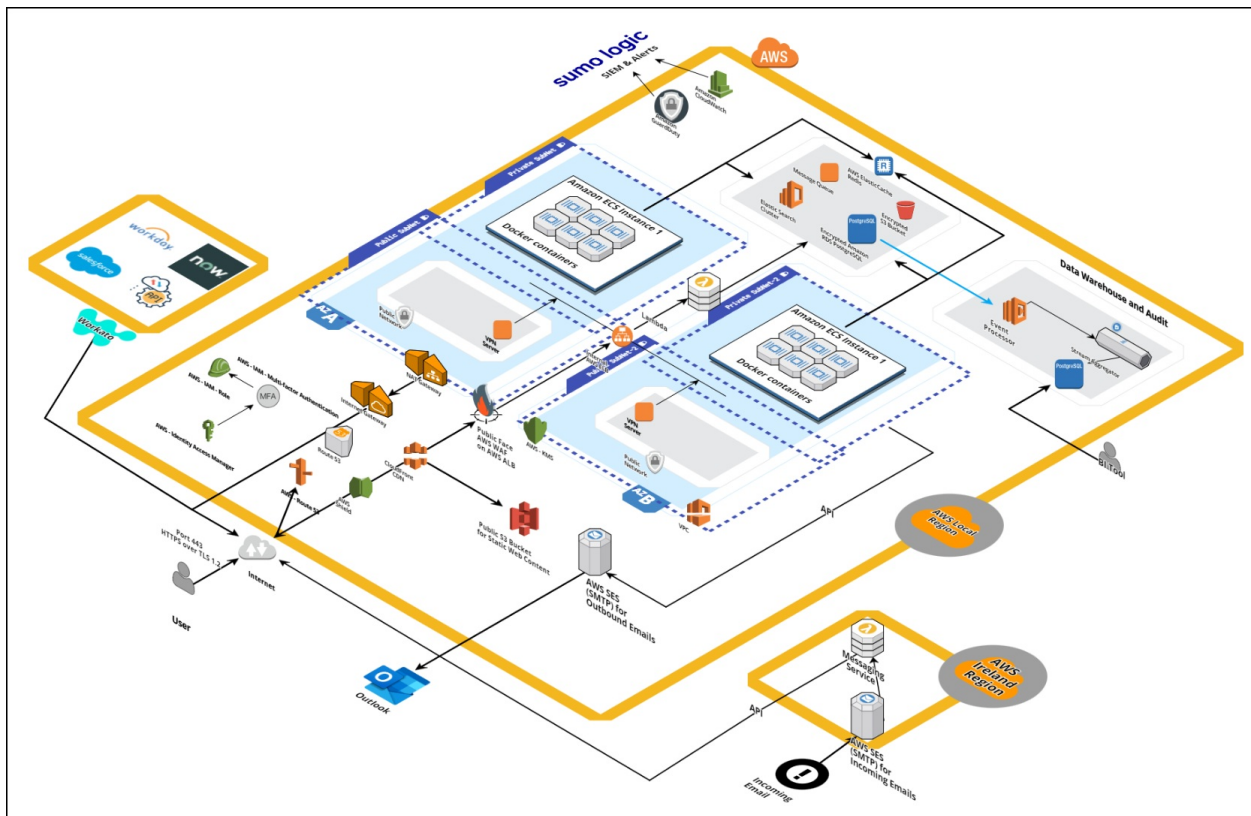
Disaster Recovery

In the event of a critical failure, the support team is committed to ensuring Resolver's complete recovery as quickly as possible. The below table sets out a list of the various types of failure Resolver users may experience and the recovery time objective (RTO) associated with that failure.

Failure Type	Recovery Target
Failure of a single compute or storage instance	Host capacity exists to create new instances 99.9% availability, 15-minute RTO.
Failure of a data center or multiple pieces of its infrastructure (E.g. multiple hosts, power and backup, network connectivity, etc.).	Redundancy between data centers within a region. 99.9% availability, 1 hour RTO.
Data deleted due to customer actions.	The RTO is a commercially reasonable effort during business hours. It could have associated costs.

AWS Architecture

Resolver is hosted on Amazon AWS. Resolver's deployment architecture is illustrated in the below diagram.



Resolver's Developer Architecture

