Brightly

Assetic Cloud Platform

FAQs

February 2024

Contents

Assetic Overview **3**
Customer Environment **3**
Software and Network Requirements 3
Data Load and Implementation 5
Production and Sandbox Environments 7
Customer Support 8
Access and Administration 10
System Integration 17
Data Ownership/Sovereignty/Privacy 22
Assetic Cloud Platform **26**
Infrastructure Overview 26
Application Architecture 28
Software Uptime 30
Security Management 31
Load Management & Performance 38
Database Management 39
Release Management 40
Backups and Disaster Recovery 41

Assetic Overview

1. What is Assetic?

Brightly Assetic™ is our world-class, cloud-based asset management system that protects your community's essential infrastructure and is designed to increase efficiency, satisfy compliance, and optimise spending.
Below is an overview of available modules and functionality within the Assetic SaaS platform:

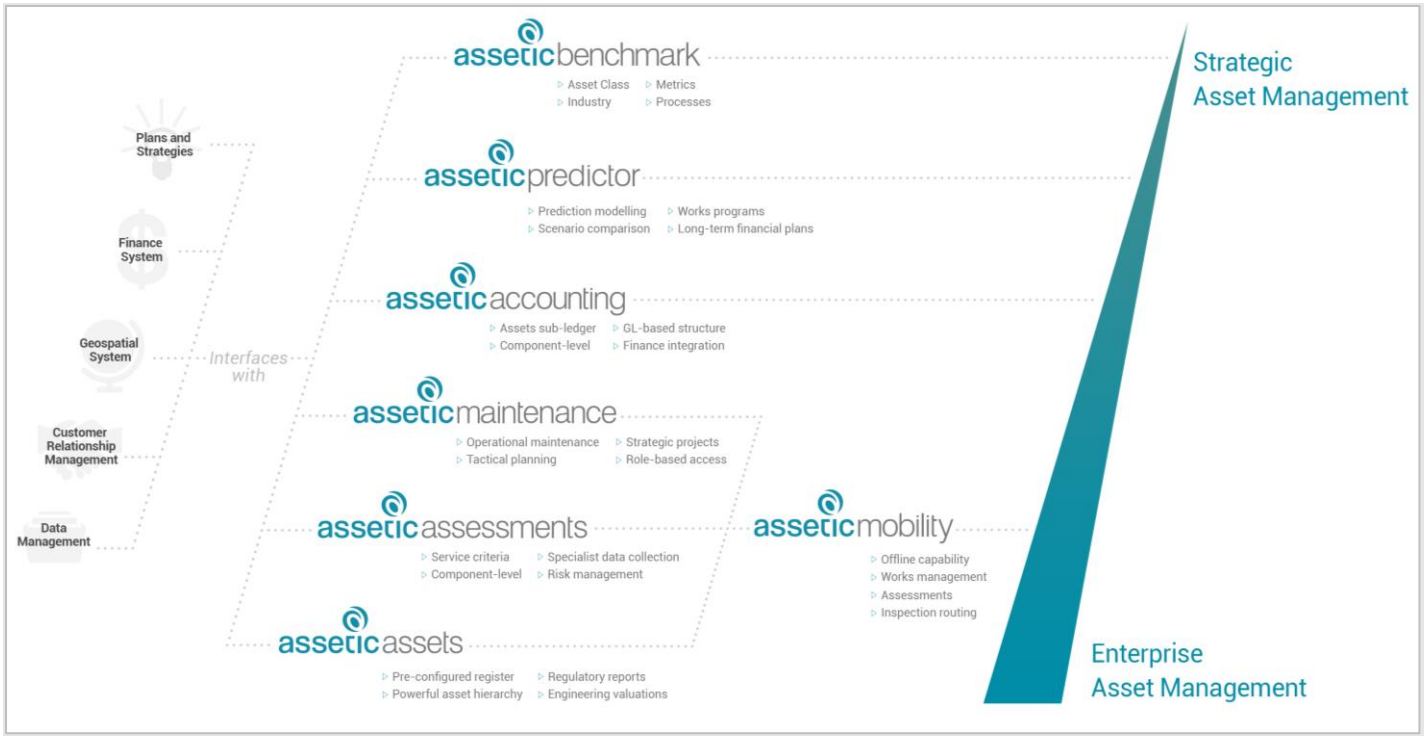

More detailed information on product and module functionality is available at:
https://www.brightlysoftware.com/en-au/products/assetic

Customer Environment

Software and Network Requirements

2. Does Assetic's service or web application provide browser-based access?

Assetic is fully browser based for all modules, except Assetic Mobility, which is an offline native app for Android/iOS.

3. What do I need to install on my local computer to use Assetic's cloud platform?

Assetic is a web-based application with no additional plugin requirements. It requires a modern and up-to-date browser that is HTML5 compatible, such as Chrome, Microsoft Edge, Firefox, or Safari. We also support IE11 running on Windows versions that are supported by Microsoft.

4. Is Assetic a responsive web application?

Yes, the Assetic application is built using responsive web design (RWD). We achieve RWD using the Bootstrap and KendoUI frameworks, in conjunction with our own custom code. As such, the web application supports tablet touch screens and finger clicks (in addition to the Assetic Mobility software, which is a native mobile application).

5. Can Assetic run on Windows 10 mobile/iOS/Android?

Assetic Mobility is a native mobile application that is available for Android (via the Google Play Store) and for iOS (via the Apple App Store), is designed to support offline work and has a simplified user interface for field use. The app is designed for the mainstream versions of Android and iOS. For Windows 10 mobile (or other mobile operating systems), the full Assetic application is accessible via a web browser.

6. Do Brightly's mobile applications work offline and support the ability to save as you go so that data entered is not lost when a device error or shutdown occurs?

Assetic Mobility has been designed to work fully offline and it caches photos, documents, and data that it requires locally. It syncs this information with the Assetic Cloud using your phone's mobile data (or can optionally be restricted to only sync when connected to a Wi-Fi network). Photos can also be compressed automatically to limit data storage and usage.

Utilising this localised storage approach and being a native mobile application, all data that is collected and processed via the Assetic mobile application is not lost in the event of a power outage or failure on the mobile device.

7. Can the entire Google map for the district be cached on the mobile app to allow for offline map browsing?

Yes, anywhere you browse on the map whilst online gets cached for offline access when out of range.

8. Does Assetic make considerations for colour blindness, vision impaired, and those that have difficulty using touch-screen technology?

Yes, Assetic follows web content accessibility guidelines (WCAG) when designing and implementing user interfaces.

9. How does Assetic support accessibility facilities, such as vision impairment and literacy levels (e.g., simple icons, language, text-to-speech, etc.) for field workers in the application? Does the application support speech recognition specifically for mobility?

The Assetic Mobility software has been designed utilising native, standard UI components on Android and iOS devices and follows the iOS and Android human interface development guidelines.

Assetic Mobility is very simple and intuitive to use and employs simple language and icons, as well as field-specific capability to ensure easy and rapid adoption by field users.

Text-to-speech is supported via the native voice dictation software available on both Android and iOS. No third-party software is required.

10. What are the bandwidth requirements for running Assetic and how many users can access the software?

We would recommend a minimum of a symmetric fibre connection of 10Mbps per 100-150 users. However, even a basic cable/DSL2 type connection of 10Mmbps/1MBps may be sufficient for up to 30 users (assuming a spread of different user roles — that is, not all data collector roles with heavy photo and document upload requirements).

The Internet connection speed required for accessing Assetic will naturally be affected by the bandwidth of your existing internet connection, usage, and number of network users. If your existing internet connection is sufficient for serving up access to other online sites and software, then it will likely already be sufficient for accessing Assetic.

11. Are there any data storage/transfer limits?

Depending on the modules purchased, you will be allocated a 250GB or 750GB base data storage allowance for each environment in use. Additional 500GB data blocks can then be purchased on an annual basis if required. Data transfer is not charged for.

Data Load and Implementation

12. Are there bulk import tools available for data, documents/attachments, etc.?

Yes, there is a built-in bulk import tool called "*Data Exchange,*" which allows for bulk upload of data files into Assetic. Data Exchange works as follows:

- **System fields (e.g., the attributes for a building's asset category) are mapped to an external data file (e.g., a CSV file)**
- **The mappings can optionally be saved as a Profile for re-use in the future**
- **The Data Exchange import is sent to Assetic's service bus to run as a "*Background Worker*" job (that is, a background process)**

It is also possible to use Assetic's REST APIs to perform bulk data imports/actions (e.g., for automated/scheduled data synchronisation with other systems).

For bulk document uploads, there is a Knowledge Base (KB) article with working sample on how to do this. The process involves creating an Excel file with all document paths that you wish to upload and running a software script to upload all documents in the Excel file for you.

13. How do I manage environments during a normal implementation? Can I get additional environments that I can use for testing or training?

Two environments are typically provided during the implementation phase (one for Testing/Training and one for Production).

A single Production environment is provided once delivery is completed, unless an ongoing subscription is required for other environments. Additional information regarding environments is provided in the Production and Sandbox Environments section below.

Production and Sandbox Environments

14. Can Brightly make different Assetic environments available for production, testing and training, and other purposes?

Yes, as part of the implementation of the Assetic cloud platform, customers have the option of accessing a number of different environments. The following information outlines the nature of these environments and rules governing their use.

Production
The Production environment is where day-to-day operations occur. As part of any data migration and system configuration by Assetic, this is the environment that is initially populated. It will never be overwritten by another environment. Availability, back-up, and disaster recovery for this environment are explained in other sections of this FAQ document.

Sandbox
Sandboxes are multipurpose environments. They can be used for testing of new processes and data (including integrations) and the delivery of training. As part of a standard deployment, one Sandbox environment is made available during the implementation phase and up to three months after it goes live.

The utilisation of the Sandbox environment can be extended at the client's request for a nominal annual fee.

This environment can be overwritten (usually from the Production environment) with permission from the client. Brightly will aim to complete this process within 24 hours. Clients can request this environment be refreshed up to six times per annum.

Note:

- *It should be assumed that to revert to a clean state/recover a Sandbox environment that this will be done via a copy of the Production environment*
- *When Sandbox environments are refreshed, the user documents and SAML configurations are not copied*

Preview
Brightly can make a pre-production environment available for specific users to test and review software before deployment to Production. The computing resources and availability of this environment are significantly less than for **Production** or **Sandbox** environments. Additionally, this environment will have regular outages to accommodate work-in-progress features, as well features upcoming in the next release.

This environment is initially configured with a copy of the data in **Production**; however, it will only be updated from **Production** once every six months at the most. It is a requirement of clients who wish to access a Preview environment that they supply Brightly with details of their testing methodology and resources prior to being given access.

As Brightly operates on a release schedule of once every two weeks, this environment is only available in the week before release, after the issuing of the Draft Release Notes. All issues logged from a Preview environment must only be related to items detailed within the Draft Release Notes supplied. Any other features/issues should be tested and lodged from Production or Sandbox environments.

This environment is not part of a standard deployment and has to be requested specifically and will be made available for an additional annual fee.

> *Note: Brightly are currently developing a separate Preview Access Agreement to provide more detail on both the client's and Assetic's obligations around environments of this nature.*

Utilisation
Brightly periodically reviews the utilisation of all environments. Should this reveal low or no usage of non-Production environments, Brightly may contact the customer to determine if these additional environments are still required.

Customer Support

15. Once the application is live, how do we receive support from Brightly?

The simplest way to receive support is within the Assetic cloud platform itself. Select the help menu at the bottom left of your screen after logging in:

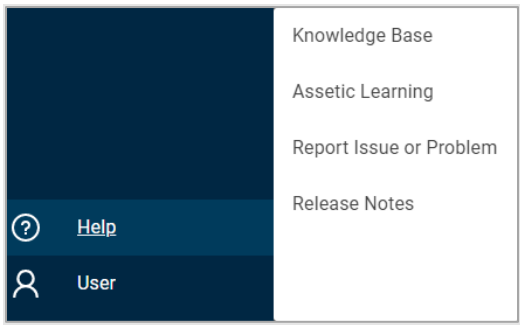

The help menu provides access to several key support and training resources:

Knowledge Base
The Knowledge Base contains over 700 articles, including:

- **Detailed explanations on how to use features in Assetic**
- **Best practice asset management and configuration principles**
- **Technical and integration resources and samples**
- **Product Updates and Assetic Release Notes (updated fortnightly)**

Many of the Knowledge Base articles include videos and business process maps to quickly show you how to navigate through key processes.

Assetic Learning

Brightly Assetic Learning hosts a catalogue of on-demand, online training courses. Study at your own pace, set your own schedule, and undertake training wherever and whenever you choose. The online training is provided in addition to face-to-face training, helping to accelerate on-boarding of new users and provides product refreshers as required.

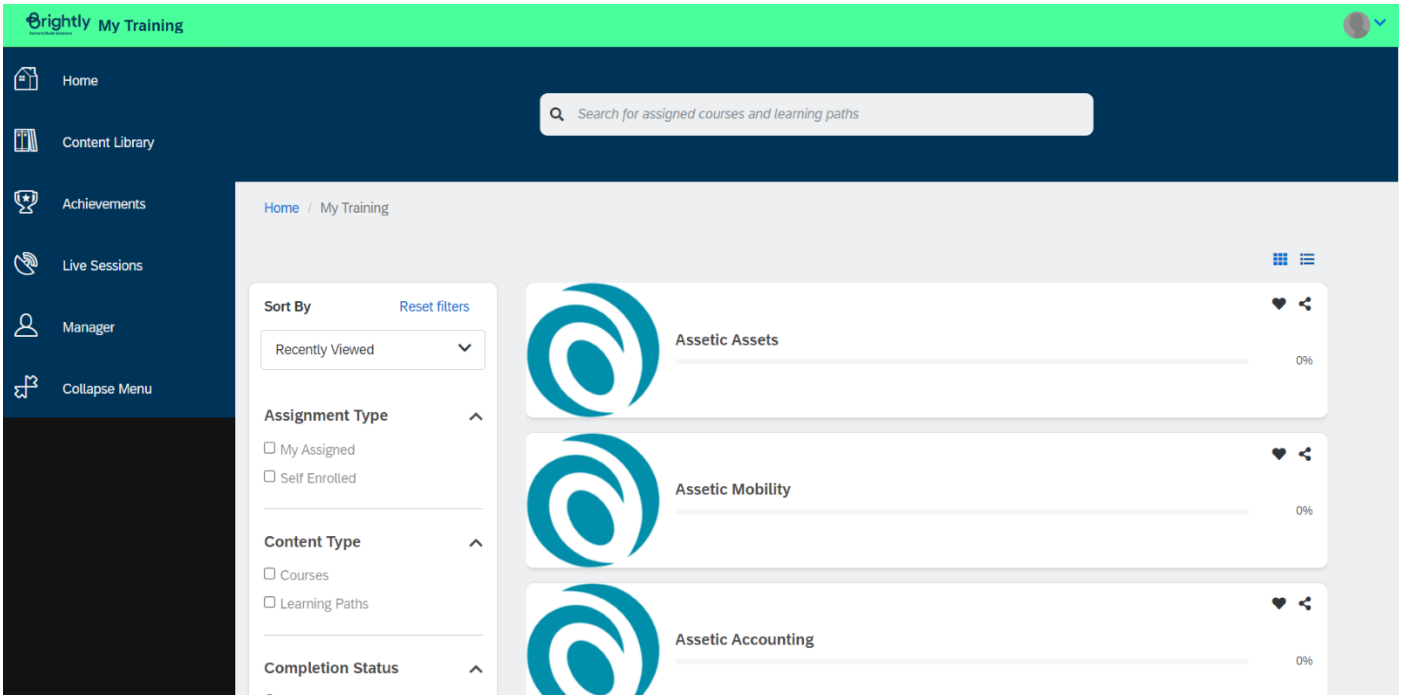

Support Tickets

If you can't find an answer in our Knowledge Base or Assetic Learning, then our support team is there to assist with your queries and suggestions for your production environment.

You can submit support requests/feedback via the **Report Issue or Problem** link. Any previous support requests you have submitted will appear in the **My Support Tickets** screen.

Brightly receives detailed diagnostic information when support requests are submitted through the application itself and this is the most efficient way to lodge support requests.

Further information regarding support response and resolution times can be requested by contacting Support, Consulting or Account Management.

16. What is the process for customer notification in the event of an application outage?

We provide system information and notifications via the Assetic Status Page: http://status.assetic.net

The status page provides an operational summary, notices of upcoming releases, and any degraded performance/outage issues being experienced. You can subscribe to the status page to receive these notices automatically via email.

In terms of unscheduled outages, the platform has been designed with a high degree of fault tolerance and it is actively monitored from a network, performance, and security point of view with real time alerts in place.

Assetic undertakes software releases every fortnight. For these scheduled outages, Brightly communicates details of major upcoming releases one week in advance of deployment via both the Status Page and Knowledge Base. The release notes provide a summary of the release highlights, as well as a detailed breakdown of other smaller features/fixes included in the release.

Access and Administration

17. Does Assetic face the public Internet? Do you protect access through VPN?

Yes, the application faces the public internet. Connections to Assetic are only allowed through HTTPS (TLS 1.2) on port 443. Port 80 is redirected to port 443. In terms of backend administration of the Assetic cloud platform, this is restricted via fixed IPs and is done over VPN.

18. Does Assetic support Microsoft Integrated Active Directory authentication? If not, do you support another LDAP authentication or local authentication only?

Yes, we support directory-based authentication with a range of identity providers that support SAML 2.0, including Windows Active Directory (ADFS), Microsoft Entra ID, and Okta.

19. Does Assetic ensure all sensitive data is encrypted both in transit and at rest?

Yes, the data is encrypted both in transit (TLS 1.2), and whilst at rest (AES).

20. Does Assetic support SSL? If so, will it support Customer SSL Certificates?

Yes, SSL access is a requirement. For a standard deployment, the end user will access Assetic by accessing the following link where <**customer-environment-name>** represents the customer's name e.g. https://brightlysoftware.assetic.net:
https://<customer-environment-name>.assetic.net

Custom domain hosting (e.g., https://assetic.custom-domain.com) is available for an additional annual fee. The customer will need to supply SSL certificate for such option.

21. Does Assetic provide tools for managing user accounts, security settings on data and/or applications?

Assetic provides a comprehensive administrative module, including managing account and security settings, in which users are assigned Assetic-specific roles and work groups.

Users, passwords, and authentication can be configured to work via a corporate Identity Provider, or alternatively Assetic provides a stand-alone user management setup (whereby users and passwords are fully managed within the Assetic admin module). Environments can also be restricted in the admin module to only allow one of these authentication modes.

22. Does Assetic provide a full password security process based on roles and groups?

Yes, system authorisation is managed via a combination of Licensing, Roles, and Work Groups to limit a user access to various parts of the platform.

Roles

Roles apply on a platform- and module-level to limit a user's ability to view, edit, and delete in specific parts of the platform.

For example, an accountant cannot assign work requests, or a works supervisor cannot create assets without the appropriate user role.

However, as a user's role within their organisation can be multi-tasked, the roles can be stacked up to provide greater access to the Assetic platform.

Example: Wendy has the Data Collector, Assessments Admin, and Assets Manager roles.

The following diagram gives her editing rights in both the Assets and Assessments modules, as well as read access to all the other modules. It is worth noting that her data collector role is not needed as she already has higher permissions via the Assessments Admin Role.

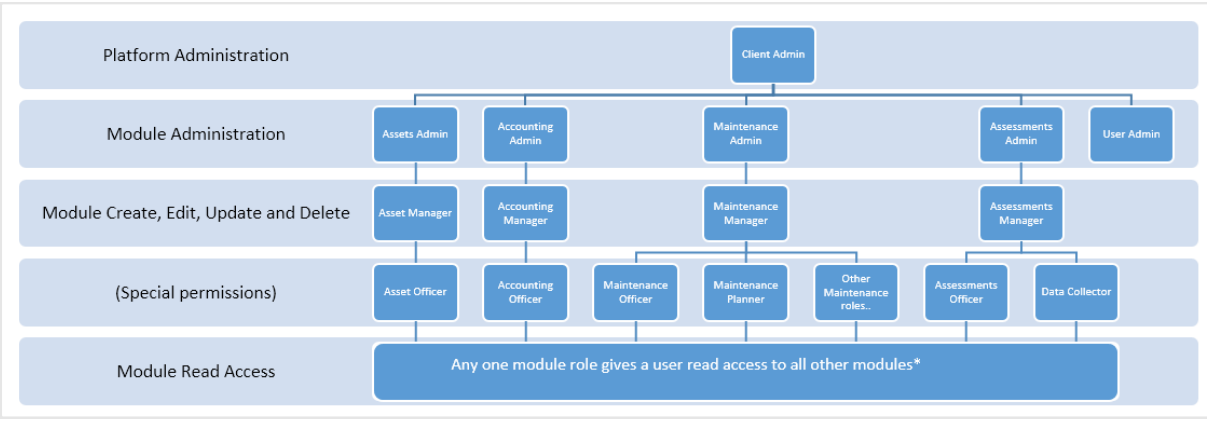

** Module read access is for embedded module content like work order history in the assets area. It does not give a user access to the module itself.*

Claims

Each role has a series of permission claims. These are also cascading in that if a role has the delete claim, it also has the update, create, and read claims.

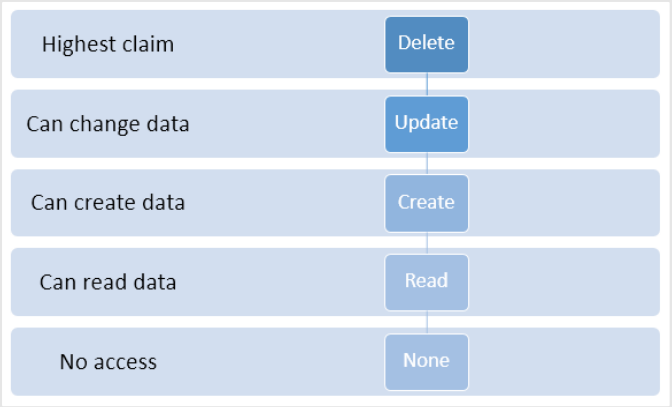

Work Groups

Work groups are assigned on a per-asset or work order basis and provide a way to filter data by regions or work zones.

Example: Joan has been assigned the Red work group and Arthur has been assigned the Yellow work group.

Based on the following image, this means that Joan can only see assets that have a work group of Kingsweston and Southmead, whereas Arthur can only see assets that have a work group of Henleaze and Lockleaze. Additionally, they can only view work orders and other system data that is related to these assigned work groups.

Note: Both of them can see complex assets that do not have any work group assigned.

A user can be assigned to multiple workgroups. For instance, Joan and Arthur have the same manager, Betty, who is assigned both yellow and red workgroups so that she can see both regions.

When the manager assigns technicians to tasks on work orders, she can only assign resources belonging to a matching work group.

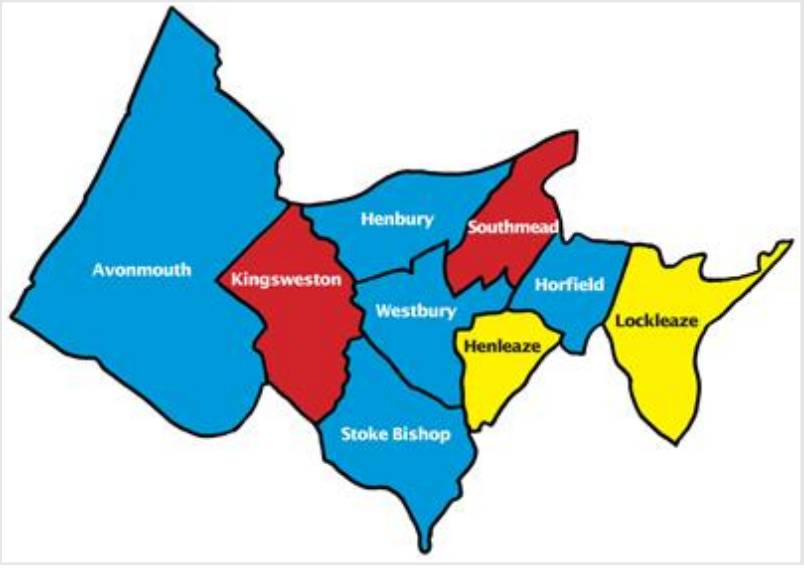

23. Does Assetic provide a flexible and secure security management process for assigning privileges and rights?

We provide comprehensive pre-packaged roles that are built up from a logical set of system rights and privileges (which Assetic pre-configures). We also utilise work groups as another layer of access control on top of this (record-level control).

24. Do you provide an audit trail of all system activity, including by user, date, and time?

Yes, an Audit Trail subsystem is tracking all system activities and the logs can be viewed at a record/sub-module level (i.e., contextual filtered views). The admin module also has a centralised view of all audit log data.

25. What auditing is available on configuration and data read/write/delete?

All system events are tracked via an isolated Audit Trail subsystem. The logs are made available within the Assetic platform itself (in the admin area and also filtered at a record level). It provides a log of configuration changes and record read/write/delete (i.e., who made the change, when the change was made, what module and what record was changed, old and new value, and further system information outputs where applicable).

26. Does Assetic support record locking to prevent simultaneous updates of the same information?

The web application uses optimistic locking to detect data that may be accidentally overwritten by another user. A warning message is displayed if multiple users are editing a major entity record concurrently (e.g., asset, work order, work request, etc.).

Additionally, the system audit trail provides a log of system changes, including the date-timestamp, the user, the entity, and the old and new values. This data is also attached at a record level against Assets and Maintenance items. The audit trail feature means one can review changes that has been made to a record.

27. Do you support session persistence and session portability across different devices or logon sessions?

No, we do not support session portability/persistence, nor do we limit the user to a single session.

28. Does Assetic provide an automatic log-off feature after a specified period of inactivity?

Yes, after 60 minutes of inactivity, users will be logged out of their web browser session by default. The session timeout value can be changed by the admin.

29. What technology components are used to enable Assetic's operational reporting architecture and are they third party reporting solutions?

Assetic includes comprehensive search, reporting, and dashboard capability that allows tabular,

chart, KPI, and pivot reporting. The inbuilt reporting functionality has been built on top of Elastic search.

Assetic also provides a read-only OData endpoint for reporting, which adheres to the version 4.0 specification. This allows third-party tools such as Microsoft Excel and Microsoft Power BI to readily connect to the Assetic platform and create custom reports. There are also third-party ODBC drivers for OData available in the marketplace that can be utilised if required.

Additional extracts and reporting can also be achieved via utilisation of Assetic's APIs. We provide an extensive set of Knowledge Base articles and SDKs to support this.

30. Can Assetic support configurable, dynamic criteria like "Last X weeks", "Next X days," and support wild card searching?

Yes, via both the in-built reporting module and APIs.

31. How can clients customise layouts and add new user-definable fields on data entry screens?

Assetic is pre-configured with a complete suite of asset classes that have been utilised across many different industries. Assetic is pre-configured with 22 infrastructure classes across 173 specific asset categories, including roads, bridges, drainage, facilities, open space, buildings, water/sewerage, and more. Each asset class has fields, algorithms and service criteria that cater for common industry standards and strategic asset management out of the box. The various attribute group boxes can be hidden and removed on a per-user, or organisation-wide, basis.

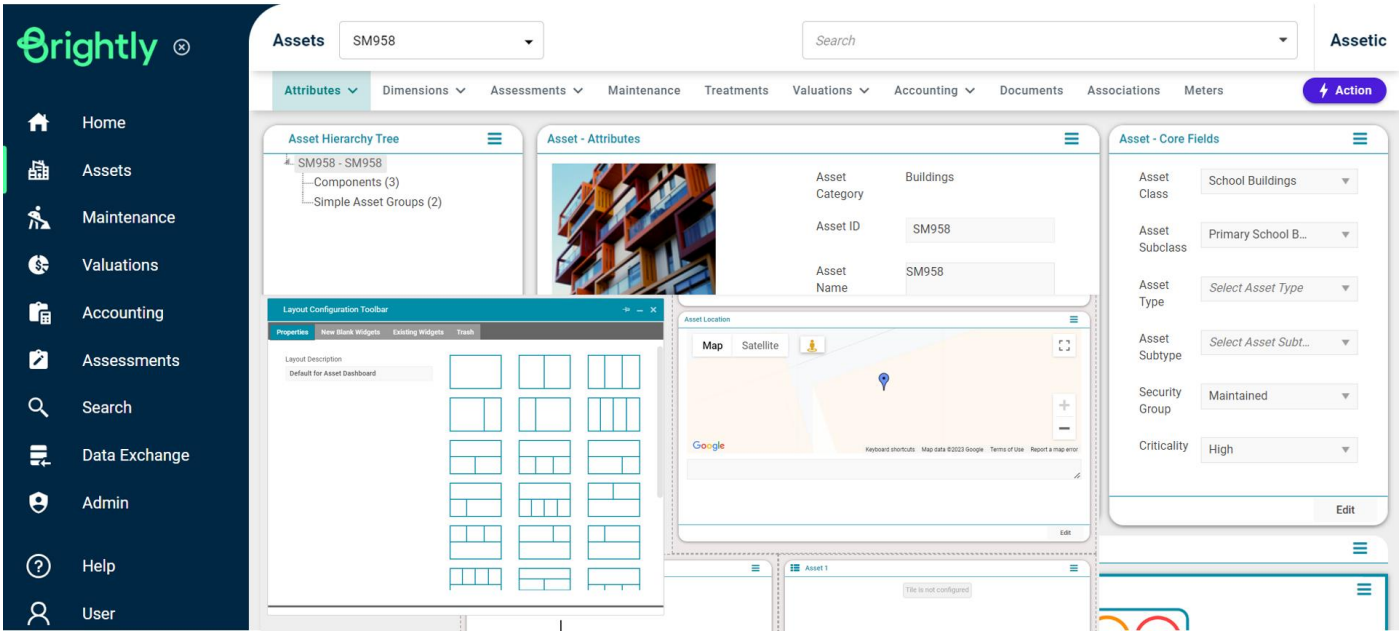

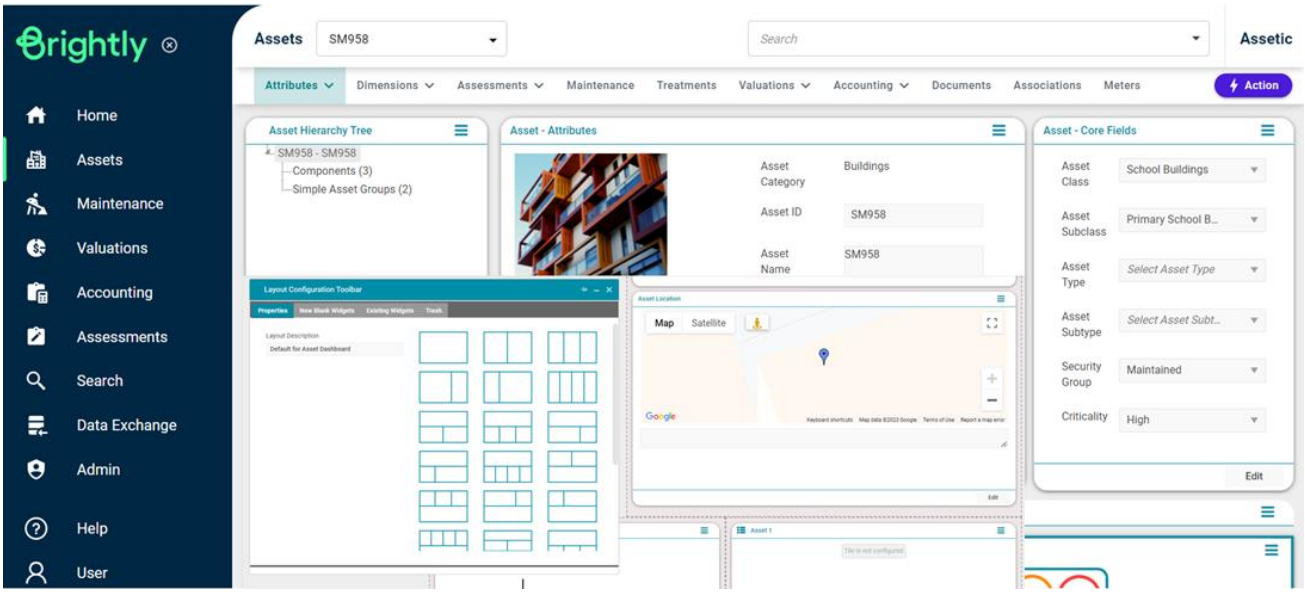

You can create various dashboards via the Reporting module, including tabular data, graphs, KPIs, and more. These dashboards and layouts are customisable on a per-user or organisation-wide basis. Additionally, where a page is contextual (e.g., a Buildings Asset Category), the dashboards can have a contextual filter applied to limit records presented based on that building's asset category context.

The Assetic Assessments module includes a form builder that allows the UI, tabs, and form elements to be fully customised and user defined. Assessments can then be attached to Assets, Work Orders, Work Requests, or other necessary areas.

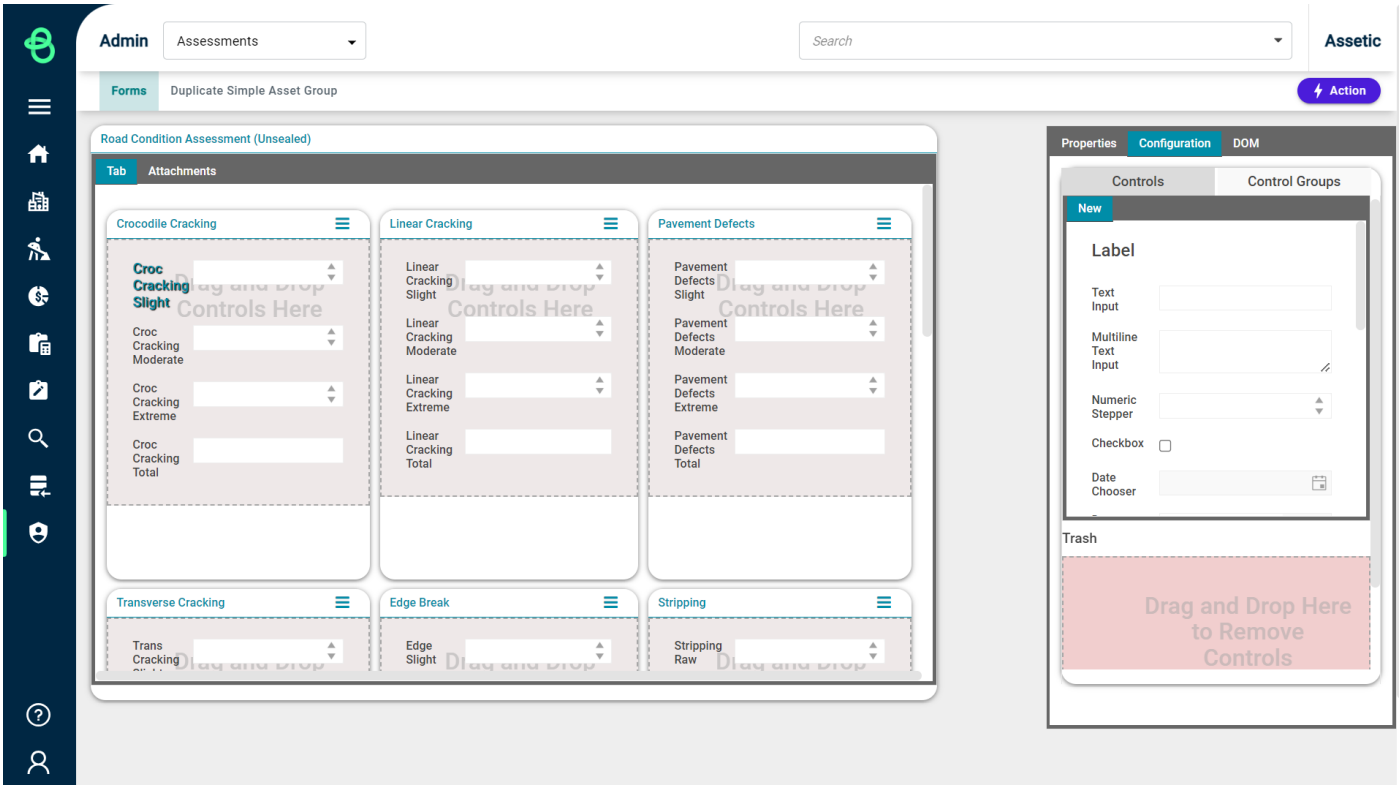

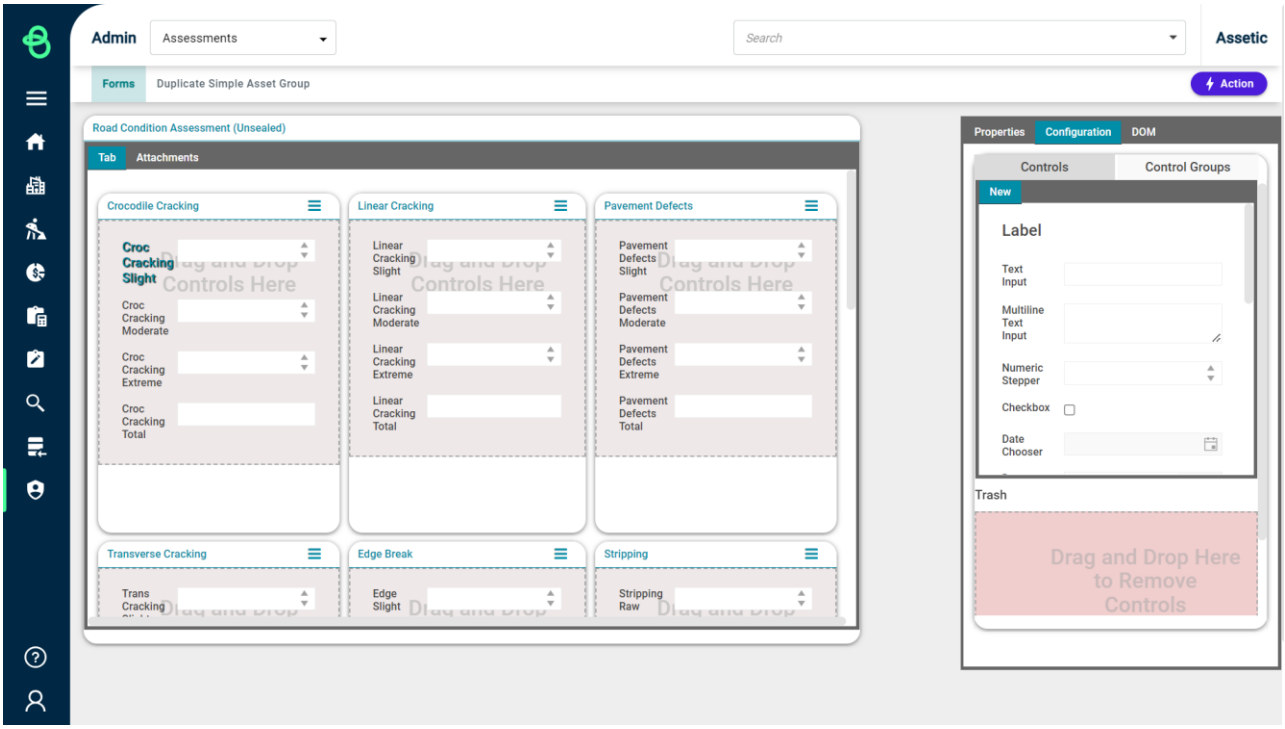

Many system entities are based on data structures that allow customisation and addition of any number of records, including for asset components, service criteria, measurements, valuation patterns, treatments, meter readings, and simple assets.

32. Does Assetic provide the ability to purge a range of data based on a specified retention schedule?

The underlying data and files are generally not permanently deleted (to prevent accidental data loss/inability to recover).

There is a facility to permanently purge data in bulk (Bulk Erase), which can be enabled for 30 days during the initial implementation phase.

Database log file trimming is done daily to reduce storage and offsite backups are also purged periodically.

33. Has Assetic been designed to ensure the requests and responses between client and back end are optimum in terms of managing bandwidth requirements? What approach was taken?

Yes, Assetic's SaaS platform has been designed with optimal performance in mind.

There are a number of technical decisions that we have made to ensure optimal network performance:

- **Reducing message passing overhead by sending business object references (rather than having to pass large datasets)**
- **Ensuring that data views are limited to a sub-set of records (e.g., 10,20,50,100 per page)**
- **Utilising indexing technology to improve the performance search and reporting**

- **Utilising a cache layer for frequently accessed data**

34. By using AWS hosting, do end customers receive access to tools like QuickSight, Amazon API, Mobile hub and analytics, and others to query our data?

Not directly, as it presents a security risk.

Assetic provides built-in reporting features that allow dashboard to be created by end users within the platform, including search queries, data aggregation, cascading grouping, charts, KPIs, pivot tables, and more.

We also publish APIs, OData (v4.0 compliant), and SDKs that enable data extraction and manipulation (e.g., for third-party system extracts/integrations). For custom reporting, OData makes it easy for end-users to consume Assetic data in tools like Power BI, Microsoft Excel, and others.

System Integration

35. Authentication Options: Does Assetic support single sign-on of users with our corporate Active Directory Federation Services (ADFS) through SAML integration? Does Assetic support integration with Customer Identity Management Systems such as Microsoft B2C, Salesforce Identity, or other third-party solutions?

Assetic supports either of the following authentication methods:

- **SAML for Single Sign On (SSO) authentication against an external identity provider, for example,**
  - **Windows Active Directory Federation Services on an organisation's network**
  - **Cloud identity providers like Microsoft Entra ID/Okta (all of which can integrate with an organisation's AD)**
- **User accounts managed within the Assetic platform**

The external identity provider can be configured with additional security layers (e.g., multi-factor authentication).

Assetic's KB provides detailed documentation on how to configure a local Windows Active Directory Federation Services (ADFS) server and integrate this with Assetic's cloud platform for central user account management. In this setup, the customer admin can also enforce SSO authentication for the end users.

36. What technical options are available for integration with other customer systems (e.g., finance systems, GIS, CRM, etc.)? How will the data be exchanged and what methods are proposed to protect this data in flight?

All data in transit must go over HTTPS (TLS 1.2).

The Assetic platform provides an extensive set of REST APIs, and we publish SDK that can be used for system integrations.

Assetic also provides a System Connector in the software with an intuitive UI to launch third-party systems contextually from within Assetic. We also have a number of powerful in-built Google Maps functions available.

Brightly provides the following technical resources:

- **Integration & Administration Knowledge Base (KB)**
- **REST API Documentation**
- **Source code for the SDKs in Python**
- **Published SDK for Python (available as a PyPI package)**

There is a dedicated Integration section in our KB that details how to get started with development and integration of the Assetic cloud platform.

The KB covers topics like how to access the APIs and SDKs, how to authenticate via API keys, how to handle common integration use cases like automated data feeds, GIS, CRM, finance and document management system integrations.

The REST API documentation is available from your cloud environment after you have authenticated. It is available at https:// <customer-environment-name>.assetic.net/apidocs. The API docs include functionality for inline testing of endpoints.

To speed up the development, the Swagger definition of the API can be used to generate sample client for major programming languages. The sample implementation for Python is available. We also publish and maintain an SDK for Python in PyPI package that is ready to use.

Assetic also provide an OData version 4.0 read-only endpoint for accessing your data via tools such as ODBC for SQL, Microsoft Excel, Microsoft Power BI, and others.

Brightly

Assetic > Integration

Search

Integration

Articles describing the use of the Assetic Cloud Platform REST API's including sample Json payloads

ASSETIC REST API INTRODUCTION

- Using Postman to test Assetic API's
- REST API Introduction
- Search Filters and Pagination
- Using Python Requests Library to access Assetic REST API
- Using .Net to access Assetic REST API

BULK PROCESSES

- IRIS Export (Main Roads WA)
- Bulk Export via Advanced Search Profile
- Data Exchange Integration

SEARCH

- Search Introduction
- Using Advanced Search Profile
- OData endpoint

ASSET INTEGRATION

- Asset Integration Overview
- Asset Configuration
- Attribute Fields - Asset and Functional Location
- Asset association with Functional Location
- Asset To Asset Association

GIS INTEGRATION

- Launch Assetic from GIS
- Launch GIS from Assetic
- Create or Update Asset Spatial Data via GIS
- Process: Create New Assets from GIS Data
- Repairing Spatial Data
- Dispose of Assets

See all 7 articles

MAINTENANCE INTEGRATION

- CRM Integration Overview
- Create Work Request
- Get Request Details
- Update Work Request
- Create and Retrieve Work Request Comments
- Get Work Order

See all 8 articles

RESOURCES

- Get Resources
- Create Resource
- Update Resource

DOCUMENTS

- Document Integration Overview
- Document Metadata Search
- Retrieving Files
- Uploading Files
- Update Document
- Delete Document

ASSESSMENTS

- Create Assessment Project
- Get Assessment Tasks
- Create Assessment Task
- Update Assessment Task
- Link Assessment Form Result to Assessment Task
- Get Assessment Form Result

See all 12 articles

FINANCE/ACCOUNTING INTEGRATION

- Import Chart of Accounts
- Export Journal Entries to General Ledger

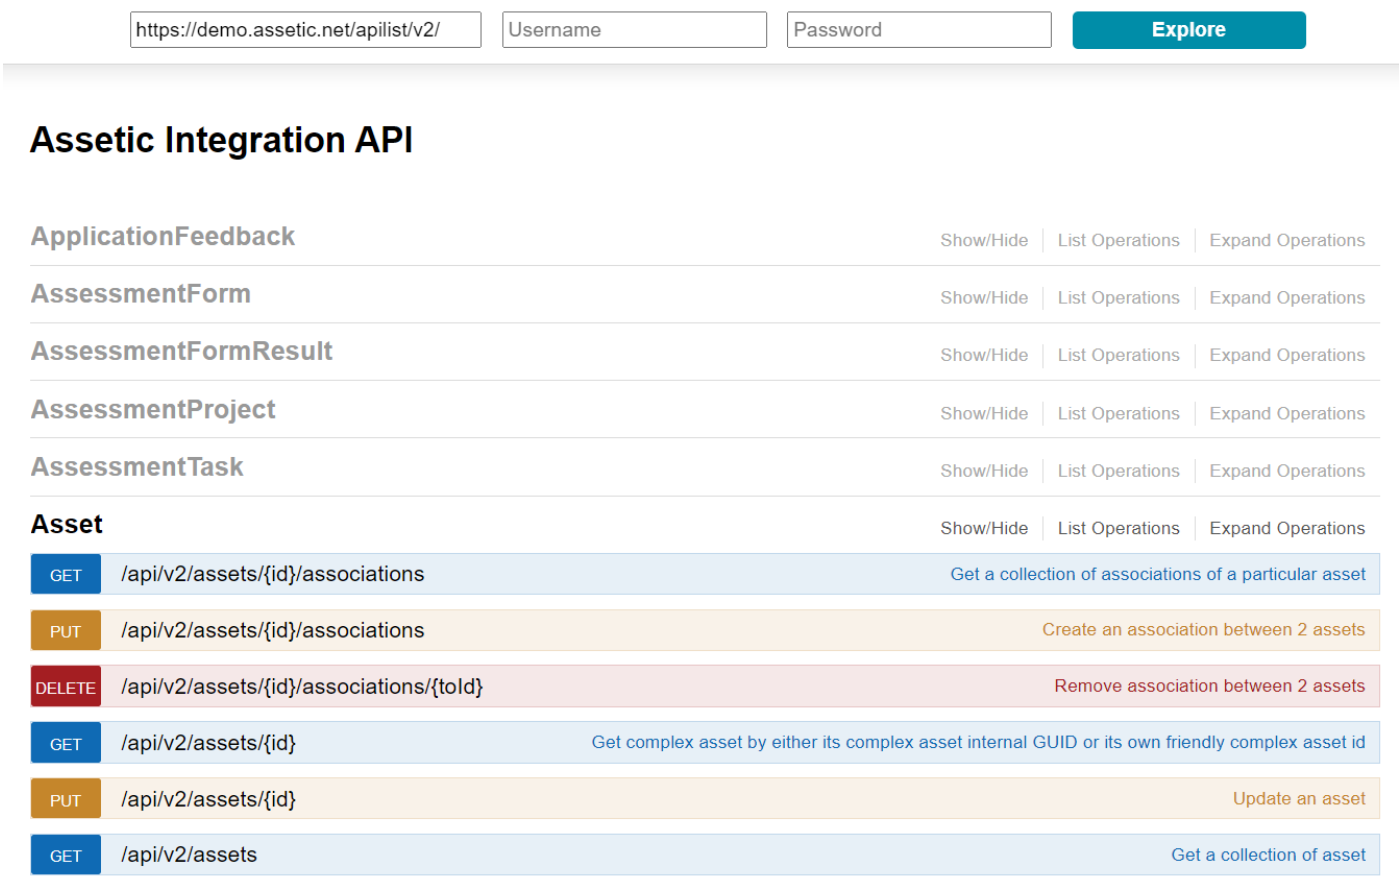

37. Are system integrations a server-side or client-side process? For example, when connecting to Assetic, does the loading of data for an integrated on-premise system (like GIS using GIS layers) load locally, or does it load to Assetic in the cloud, then load back to the client?

There are two typical Integration methods employed and they're both often employed simultaneously for any given system integration:

1. **Data-Level integration**
2. **Application-Level Integration**

Both methods of Integration are setup server side (i.e., within Assetic and the third-party system being integrated).

The Assetic application is a server-based application and can be readily integrated with other server-based applications via the **data-Level** integration (i.e., data flowing between the third-party system and Assetic, with data physically stored in both systems). One common integration use case is the transfer of core asset fields (including Asset ID and Asset Name) and polygon data between the Corporate GIS and Assetic.

For other integration use cases, it makes sense to adopt an **application-Level** integration, which is easily configured at a server-level within Assetic's admin area and gets executed at web browser-level. Typical use cases include things like:

- GIS launch button in Assetic to run a network-trace on the current asset (a similar URL link is inserted into the GIS to launch Assetic for viewing detailed strategic asset data)
- Document management system search based on current Asset ID/Asset Category/Insurance #, etc.

38. Will a copy of the database be made available for integration purposes until integration points are developed? Will it be a full copy of the database, subset of normalised data, or some other method?

Brightly supplies data exchange templates (including a data dictionary) so that data upload mappings can be developed and saved for future reuse. Similarly, data can be extracted via reports and data exports within the front end of the software and saved as a profile for ongoing reuse.

Data extracts/updates must go through Assetic's web application or through Assetic's REST APIs. This is a requirement so that security, business logic, and workflow processes are adhered to.

The Assetic Knowledge Base provides a detailed article on how to use the Assetic Python SDK utility to maintain a synced copy of the Assetic data in a local SQL server database. It automates the bulk export capability of advanced search profiles via the REST APIs. The export file is downloaded and synced with the existing version of the data in the site's local SQL Server database. This export and download process is run as scheduled tasks on a site's local server.

Assetic also provide an OData version 4.0 read-only endpoint for accessing the Assetic data via tools such as Microsoft Excel and Microsoft Power BI.

39. Is Assetic Open GIS Consortium (OGC) Standards Compliant?

Yes, Assetic uses Well Known Text (WKT) format for uploading spatial data and for export. Assetic also offers a GeoJson search option since this is another common standard.

Data Ownership/Sovereignty/Privacy

40. What is Brightly's approach to data ownership?

Brightly takes very seriously your rights to all data supplied. Brightly 's Terms and Conditions of Business and Privacy Policy formally set out Brightly's approach to data ownership, sovereignty, and privacy. The customer retains legal ownership of all of their data stored within the platform.

41. Does the customer retain legal ownership of the information, or does it belong to the Cloud Service Provider (CSP)? How will the customer be consulted if it is to be shared with third parties beyond the provider? Will it be considered an asset for sale by liquidators if the Cloud Service Provider (CSP) declares bankruptcy?

The customer retains legal ownership of the information. Our Terms and Conditions of Business only contemplate using the information for providing the Assetic cloud service. If there was a requirement to share information with third parties (i.e., a legal requirement), then we would inform the customer in writing and limit information provided wherever possible. No, customer data would not be considered an asset for sale by liquidators in the event of CSP bankruptcy.

42. Where is the provider's registered head office? Which countries does the provider deliver services from? Does the provider depend on any third parties? If yes, where is their registered head office and their services delivered from? How will the Customer be consulted if these third parties change?

Brightly's registered head office is in Cary, North Carolina, United States. We have offices and staff in the following countries:

- **Australia**
- **Canada**
- **India**
- **U.S.**
- **U.K.**

For international work we do in Asia, UAE, and South Africa, this is provided by Australian staff.

Assetic is hosted on Amazon Web Services (AWS) Cloud infrastructure, which is made up of regions and availability zones around the world. Assetic software is hosted in availability zones enabling the platform to operate applications and databases with fewer faults and higher availability than from a single data centre.

Assetic currently operates from either Australia, Canada, U.K. or U.S. AWS regions and your information is only hosted in the AWS region closest to you. Offsite backups are encrypted and stored in Google Cloud in their equivalent local data centres.

In the event that Brightly were to change cloud providers, this would be communicated via our usual communication channels – Status Page, Knowledge Base, and Technical FAQs documentation. Additionally given this is substantive change, it would require extensive customer

consultation and formal written advice to you. This type of change would only be considered in the context of moving to a similarly reputable and large-scale cloud provider.

43. Information Requests: Under what circumstances will customer information be shared with external entities (e.g., governments, law enforcement and regulatory agencies, etc.)? How does the provider handle requests for customer information?

As per our Terms and Conditions of Business, we will not share customer information with third parties unless required to do so by law or any regulatory authority. We would do so in consultation with the customer and ensure any such disclosure is made on a confidential basis to the maximum extent possible.

44. If the customer wanted to migrate to a different provider or insource the service, what are the procedure, policies, costs and/or penalties that apply to the customer in the event of such a request?

Assetic provides several data extraction and migration options, including:

- **OData - for ODBC/Microsoft Product adapters**
- **REST APIs**
- **SDKs**

Our Knowledge Base provides extensive documentation and tools to setup and automate various types of data extractions. Data extractions are common for a number of use cases (e.g., backups, BI reporting, system integrations, data migrations).

The costs to terminate the service would be covered under the project contract. Typically, this would be based on a payout of the subscription amount and term agreed in the contract.

45. Can data be retrieved from redundant locations and be destroyed in compliance with privacy laws at termination (e.g., Australia Privacy Act or Canada Privacy Act)?

Privacy laws including Australia/New Zealand/Canada Privacy Acts are concerned with personal information, not asset data. However, on termination, Brightly will destroy all data in your environments on a specified date at which the data should be destroyed.

Our approach to secure deletion and disposal of systems and information assets, so that information stored cannot be reconstructed when no longer required for business purposes, is summarised below and covers both digital and physical media:

- **Secure deletion of all data must be completed (using The United States Department of Defense 5220.22-M method must be used to overwrite the hard drive three times);**
- **Company-specific files and configurations (such as routing tables, firewall rules, network and VLAN information, port settings) must be removed; and**
- **Physical destruction by use of a purpose-built device, drilling through, or otherwise physically destroying the media; or**

- **Third parties who have a secure disposal service, including certification of destruction of the media**

46. Privacy and confidentiality — personnel vetting. Consider vendor, procurement, hosting and sub-contractors' personnel in relation to data security, privacy and confidentiality. What is the personnel vetting and employment process? What measures are taken to restrict data and systems access? Are any third-party certifications available?

Reference and personnel checking is carried out for all Brightly employees. All employees receive security training as part of our security management policies and procedures. Employment contracts and security policies and procedures ensure privacy, confidentiality, security, and acceptable ICT use issues are all documented and agreed to.

All Brightly system access requires MFA and is restricted to departmental groups. Users, groups, and MFA is all centrally managed via an Identity Management System (Okta) to ensure that access granted/revoked propagates through all systems from a central source. Access is further restricted by system-level roles within the various internal systems.

Third-party training and certifications are in place from a dev/admin point of view (e.g., Microsoft, AWS, etc.). Personnel regularly attend a number of dev meetups, security, and AWS events. Production stack access is only available to a small group of engineers that have had police checks and extensive AWS training/certification.

47. How will customer data be separated from general data if/when benchmarking?

Customer data is logically separated from other organisations and data at rest is encrypted.

For benchmarking purposes, organisational data will be non-identifiable before analysis and reporting (i.e., analysis and reporting are done based on cohorts only).

48. What is Brightly's approach to international privacy laws?

Assetic is hosted on Amazon Web Services (AWS) cloud infrastructure which is made up of regions and availability zones around the world. A region is a physical location where the availability zones exist as distinct data centres. Brightly software is hosted in availability zones enabling the platform to operate applications and databases with fewer faults and higher availability than from a single data centre. Brightly currently operates from either Australia, Canada, or U.S. regions. AWS regions and your information is hosted in the AWS region closest to you. Hosting can be provided through other AWS availability zones globally, if required: https://aws.amazon.com/about-aws/global-infrastructure/

Data protection and privacy are key to Brightly and we ensure any personal information is protected in accordance with applicable data protection laws and Brightly 's Privacy Policy. We

understand that through use of the AWS Cloud, customers may be concerned about applicable privacy laws where an AWS data centre does not exist in a particular country.

Most jurisdictions where we do not have data centres have laws concerned with the protection of information about an identifiable individual as outlined below. Brightly collects information about assets, not information from which an individual customer may be identified. Accordingly, the requirements of privacy legislation relating to an identifying individual are not applicable to the service we offer.

Australia
Protection of personal information is governed by the Privacy Act and the Privacy Principles. It defines personal information as any information or an opinion about an identified individual, or an individual who is reasonably identifiable:

- **whether the information or opinion is true or not; and**
- **whether the information or opinion is recorded in a material form or not**

The Australian Privacy Principles specifically provide under Principle 8 that before an entity may disclose any personal information to an overseas recipient, reasonable steps must be taken to ensure the overseas recipient does not breach the Australian Privacy Principles. This adds another layer of protection to an individual's personal information.

New Zealand
Protection of personal information is governed by The Privacy Act 1993. It defines personal information as information about an identifiable individual. It needs to identify that person or be capable of identifying that person.

Canada
Protection of personal information is governed by The Privacy Act 1983 and the Personal Information Protection and Electronic Documents Act 2000. It defines personal information as any information about an identifiable individual.

U.K.
Protection of personal information is governed by The Data Protection Act 1998. It defines personal information as personal data; this is data which relates to a living individual who can be identified:

- **from that data: or**
- **from that data and other information, which is in the possession of, or is likely to come into the possession of, the data controller**

Assetic Cloud Platform

Infrastructure Overview

49. What is the overall structure of the Assetic cloud as implemented in AWS?

The following diagram represents the high-level infrastructure architecture for Assetic's cloud platform:

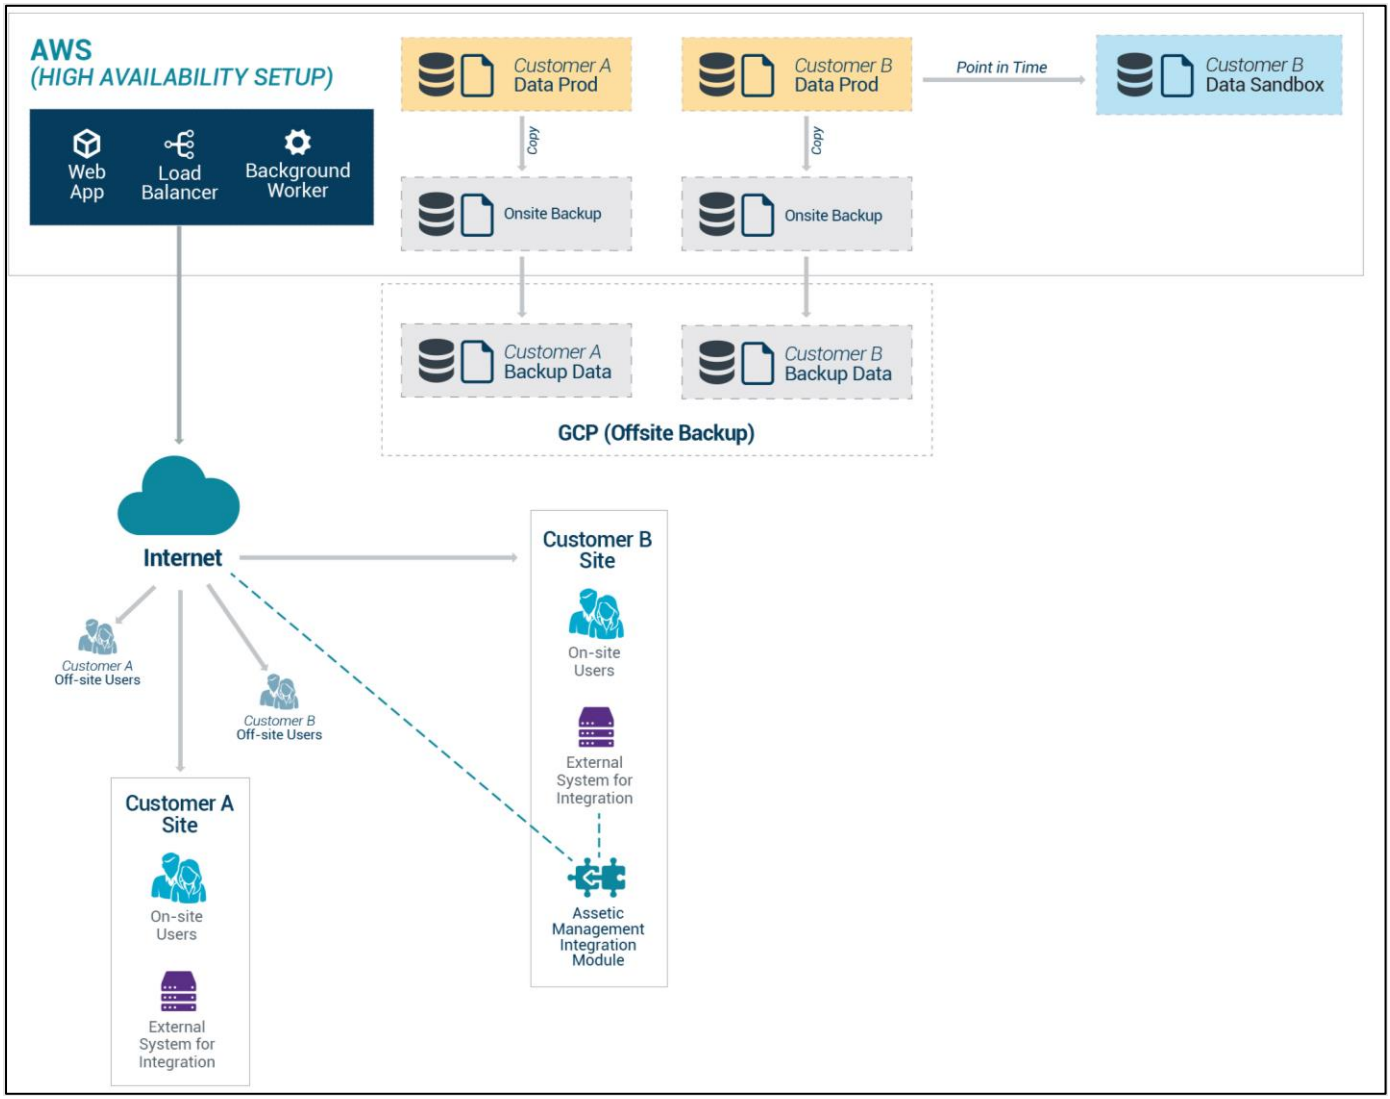

- Customers are responsible for internet connectivity and any External System Integrations
- Brightly is responsible for the back-end infrastructure and backups running across multiple data centres (High Availability Setup)
- Access to application data is only accessible over HTTPS (TLS1.2) via a browser interface, REST APIs, or OData. Brightly provides the following resources:
  - **REST APIs (Swagger documentation via the API docs page)**
  - **Base SDKs for C# and Python (other languages can be provided on request)**
  - **OData version 4.0 endpoint**
- Customers will have one production environment and multiple sandbox environments can be created as a copy of production on request (e.g., training):
  - **All production databases and file objects are backed up**

- **The DB and file objects must be backed up at a consistent point in time**
- **The initial environment originates as and remains the production environment**
- **Additional sandbox environments may be initiated via a point-in-time copy of the production environment, or can be an empty environment**

Application Architecture

50. What is the general overview of the application architecture, in terms of layers?

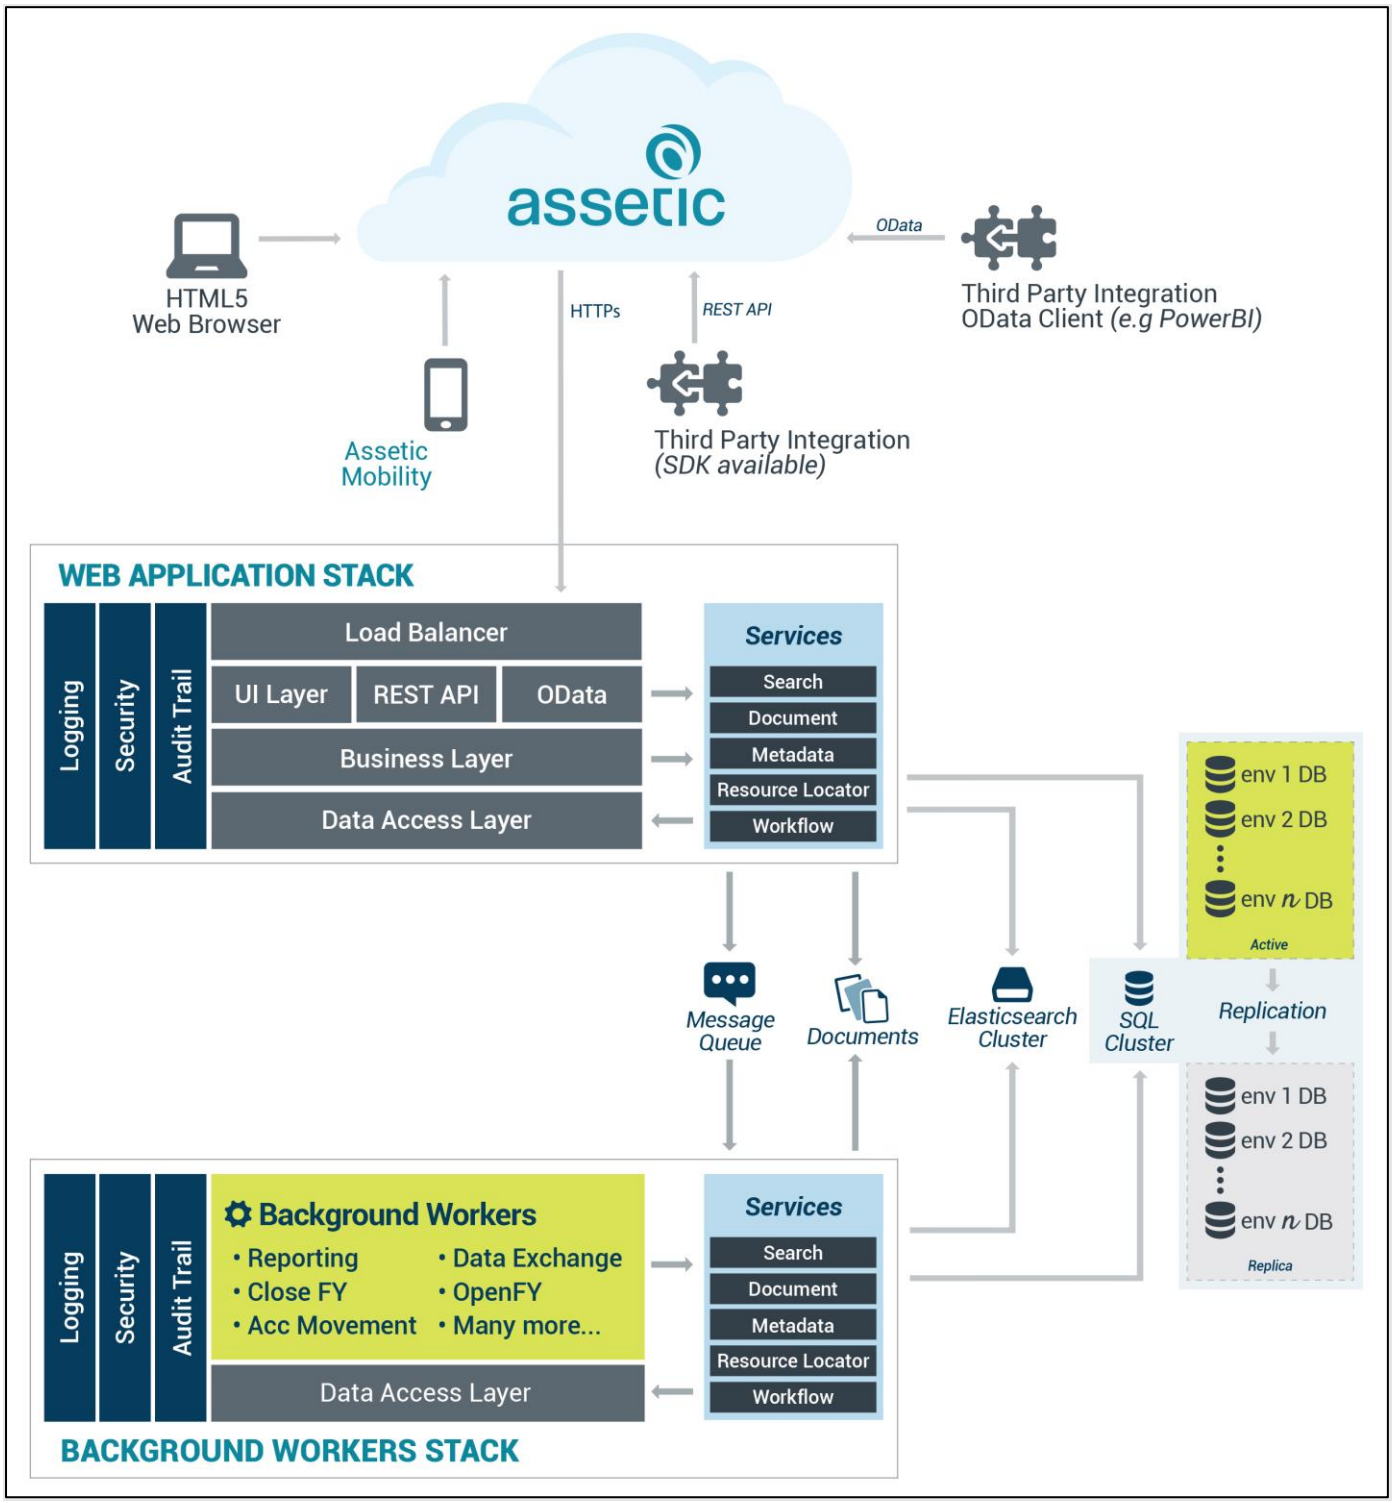

The above diagram shows how Brightly has applied a Service Oriented Architecture (SOA) to its SaaS platform. The five horizontal layers to an SOA are:

- **Consumer Interface Layer (GUI)**
- **Business Process Layer (Business use cases)**
- **Services (Used to serve up several server components)**
- **Service Components (Technical building blocks to make a service – e.g., libraries/interfaces)**

- **Operational Systems (Data & Technological Platform)**

Cutting across this are the four vertical layers (Quality of Service, Integration, Business Information, and Governance), which are applied to each of the horizontal layers. These are addressed as follows:

- **Quality of Service**
  - **Security**: The system has been built with secure development processes in mind and annual penetration audits are conducted
  - **Availability**: A scale-out architecture and a leading global cloud provider (AWS) are utilised in order to ensure the SaaS platform is highly available
  - **Performance**: The SaaS platform has been designed with performance at its core
- **Integration** - The platform has been fully built on REST APIs in order facilitate integration across any part of the system
- **Business Information** – This is at the core of how and why we build software. Brightly, and our Assetic product owner group in particular, have deep technical and asset management expertise
- **Governance** – At its core, our IT Strategy (particularly as it relates to Brightly's SaaS platform) is about achieving:
  - **Industry Excellence**: Asset Management best practice needs to be built into our SaaS platform and disseminated through our Learning Management System
  - **Quality**: Performance and scalability needs to be at the core of our SaaS platform
  - **Security**: Ensuring we deliver industry best practice security within both the platform and within Brightly

51. How are the distinct client databases segregated?

Each client environment is in a separate database within a SQL Instance. There are also associated and fully separated Elasticsearch Indexes for each environment within an Elasticsearch Cluster.

It is also worth noting that Clients in different global regions will also be in separate production stacks and thus separate SQL instances and Elasticsearch Clusters respectively.

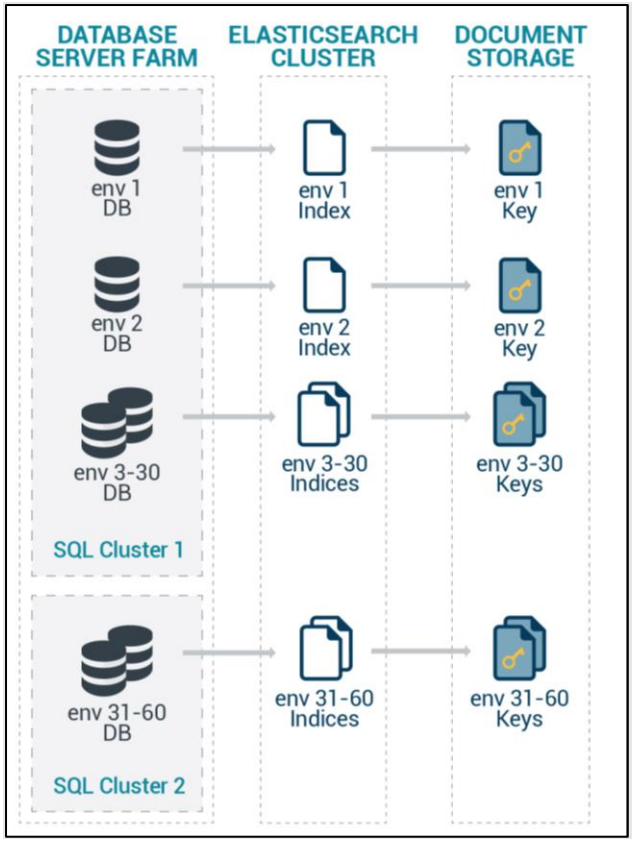

52. Does Assetic make use of CDNs?

No, Assetic does not currently use CDNs. We are evaluating the possibility of moving static content into CDN network for faster page load and reducing load on our infrastructure in a future release.

Software Uptime

53. What is Assetic's targeted service availability and how is this achieved through availability methods and geo-diversity? Are there any relevant business continuity methods that are used to maintain services?

Assetic SaaS platform has been developed utilising a scale-out architecture, which is a common approach for modern cloud platforms. A scale-out architecture enables Assetic's application, services and database layers to be highly scalable, due to the fact that additional computing resources can be continually added to expand the available processing and memory power of the application.

For application and services, stacks we utilise separate AWS Auto Scaling groups to maintain application and service availability. AWS Auto Scaling allows Assetic to scale EC2 server capacity up or down automatically. AWS Auto Scaling is configured to automatically increase EC2 instances based on demand spikes, and additionally a larger number of fixed EC2 instances are by default available during typical usage hours (6:30 a.m. – 8 p.m.). Reporting is configured via auto-alerts within AWS Auto Scaling, and we also utilise SumoLogic for performance, security, and error logging and analysis.

For the database layer, we utilise a SQL database cluster and a normalised database schema.

Brightly utilises at least two AWS availability zones to ensure the Assetic application is highly available. These zones operate as follows:

- **Each availability zone runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable**
- **Common points of failures like generators and cooling equipment are not shared across availability zones**
- **Additionally, they are physically separate, such that even extremely uncommon disasters such as fires, tornados, or flooding would only affect a single availability zone When one zone fails, the application is failed over to the other one automatically**
- **Full, daily data backups are encrypted with AES and stored remotely (along with the entire SaaS platform) in the event both Amazon availability zones are down**

Brightly provide system information and notifications via the Assetic status page: https://assetic.statuspage.io/

The status page provides an operational summary (including uptime) and provides notices of upcoming releases and any degraded performance/outage issues being experienced. In terms of unscheduled outages, the platform has been designed with a high degree of fault tolerance and it's actively monitored from a network, performance, and security point of view with real time alerts in place.

Brightly performs scheduled maintenance every fortnight and we provide detailed release notes one week in advance of deployment within the Knowledge Base and status page. The release notes provide a summary of the release highlights, as well as a detailed breakdown of other smaller features/fixes included in the release.

Brightly targets 99.9% uptime 24x5 and perform scheduled maintenance activities in out of hours periods over weekends.

Security Management

54. What is security like on Assetic?

Brightly takes security seriously and we engage a leading security firm for an annual security audit and penetration testing in order to ensure that Assetic operates in accordance with industry best practice guidelines, including ISO27001, ISO27017, NIST 800-144, CSA 3.0 & ASD.

Assetic software is deployed on Amazon Web Services (AWS) high availability zones to ensure industry-leading security and reliability standards are met (including ISO27001 certification).

Brightly has strong security policies and procedures in place, and we run security training throughout the organisation. From a development standpoint we perform secure code reviews and attend various security focused seminars/training events and also coordinate development tech talks internally.

We have setup up an internal RSS feed with security content from various news sources (e.g., Newsnow, http://www.theregister.co.uk/security, http://slashdot.org), which we disseminate to our Development and Dev Ops teams. The development team also references the Open Web Application Security Project (OWASP), CERT, Build In Security and SANS/CWE to support their knowledge and awareness of secure coding practices.

55. Will a 3rd party security organisation be performing penetration tests?

Yes, an annual security audit and penetration testing engagement is in place. There are weekly automated security scans on the web application. We perform monthly internal audit to check for any known security vulnerabilities on the tools and libraries used by the platform.

56. How is access to Assetic secured? Is there domain authentication with integration into an organisation's domain required through a firewall, is it limited by outgoing IP of our network etc.? How does this work for mobile clients?

All data transmitted between browser/mobile client applications and Assetic's cloud platform is encrypted via HTTPS (TLS 1.2) and all data at rest is encrypted (AES).

Technologies/deployment utilised for securing the Assetic platform include a narrow network surface area, firewalls, load balancing and auto-scaling, network monitoring tools, anti-virus, content filters, and separation of services.

Secure development process which adopted by following OWASP best practice, monitoring security bulletin from CERT/SANS/CWE, internal security code reviews, and annual penetration testing via leading, external security firms are all key aspects of how we ensure the Assetic application itself is secure.

Assetic supports either of the following authentication methods:

- **SAML for Single Sign On (SSO) authentication against an external identity provider, for example**:
  - **Windows Active Directory Federation Services on an organisation's network**
  - **Cloud identity providers like Windows Active Directory Federation Services (ADFS)/Microsoft Entra ID/Okta (all of which can integrate with an organisation's AD)**
- **User accounts managed within the Assetic platform**

The external identity provider can be configured with additional security layers (e.g., multi-factor authentication).

Assuming an on-premise deployment is chosen, an organisation can limit access to Windows Active Directory Federation Services (ADFS) as per integration with other Microsoft services: https://support.microsoft.com/en-us/kb/2510193

Microsoft's best practice deployment approach is that an ADFS proxy is configured on the DMZ, which then talks to the ADFS server. The ADFS proxy server can also be locked down for user

devices on the firewall (e.g., Allow TCP 443 for allowed user devices based on IP address).

Assetic's Knowledge Base provides documentation on how to configure a local Windows Active Directory Federation Services (ADFS) server and integrate this with Assetic's cloud platform for central user account management.

57. Does Assetic have appropriate build and hardening standards that meet appropriate security requirements?

All servers must be configured to a standard operating environment (SOE) build in conformance with industry accepted hardening standards, including at minimum the following requirements:

Config

- **Servers must have only one primary function (e.g., hosting a service);**
- **Servers must be deployed with a firewall separating them from the internet;**
- **Antivirus software must be installed and enabled;**
- **Only ports and services with a defined business requirement are to be enabled**

Management

- **All enabled functions are to be documented and securely configured;**
- **All non-standard functions for which there is no business requirement are to be uninstalled or disabled, including drivers, subsystems, file systems, scripts, and other features**

58. Are third party resources referenced and versioned?

Only stable versions of third-party libraries are used. These are developed and tested (in regression testing) within dev and preview environment before ever being released into production.

59. How is privileged access to your servers managed?

Connections to the web application are only allowed through HTTPS (TLS 1.2) on port 443. Port 80 connections are re-routed to port 443.

Production stacks are completely isolated from development stacks (refer Release Management documentation) and only a small number of engineers have production stack access. Production administration is restricted via fixed IPs and is done over VPN.

60. What parties are involved in providing service (including suppliers and subcontractors), their roles in the delivery and maintenance of Assetic, and their operating locations (with special attention to any overseas supplier, technician, or support function with any access to any infrastructure, software, or services containing customer information)?

Brightly is the only provider involved in the development, delivery, and maintenance of the Assetic cloud platform. All Brightly personnel, including software managers, software engineers, QA/testers, DBAs, Integrators, and Dev Ops engineers are based in our offices in Melbourne,

Australia and Noida, India. Brightly supports a work from anywhere culture, meaning employees may also be based from their home office in other locations within Australia.

Production stack access where customer information is held is only available to a small group of engineers in the Dev Ops team that have had police checks and AWS training/certification.

We utilise AWS for the cloud infrastructure and Google Cloud for offsite backups (only local AWS and Google data centres are used if you are based in Australia/U.S./Canada/U.K.).

We engage a leading and independent third party for regular security assessments and penetration testing. They do not get access to production systems.

61. Do you have intrusion detection/prevention mechanisms in place alerting you to possible attacks against your web service or application?

We utilise a number of monitoring and security vulnerability scans technologies to monitor the cloud platform including SumoLogic and AWS CloudWatch.

We use SumoLogic (log aggregation tools) extensively for multiple purposes, from performance management to intrusion and anomaly detection. Logs such as traffic logs, system event logs, and application logs from different servers are all collected and sent to SumoLogic. These logs are then aggregated for dashboarding and also used to compare against with server health and utilisation information for anomaly detection. In addition, there are regular searches set up for sending notification for top security events and potential malicious requests.

Here is an example of our operation dashboard that is IDP related:

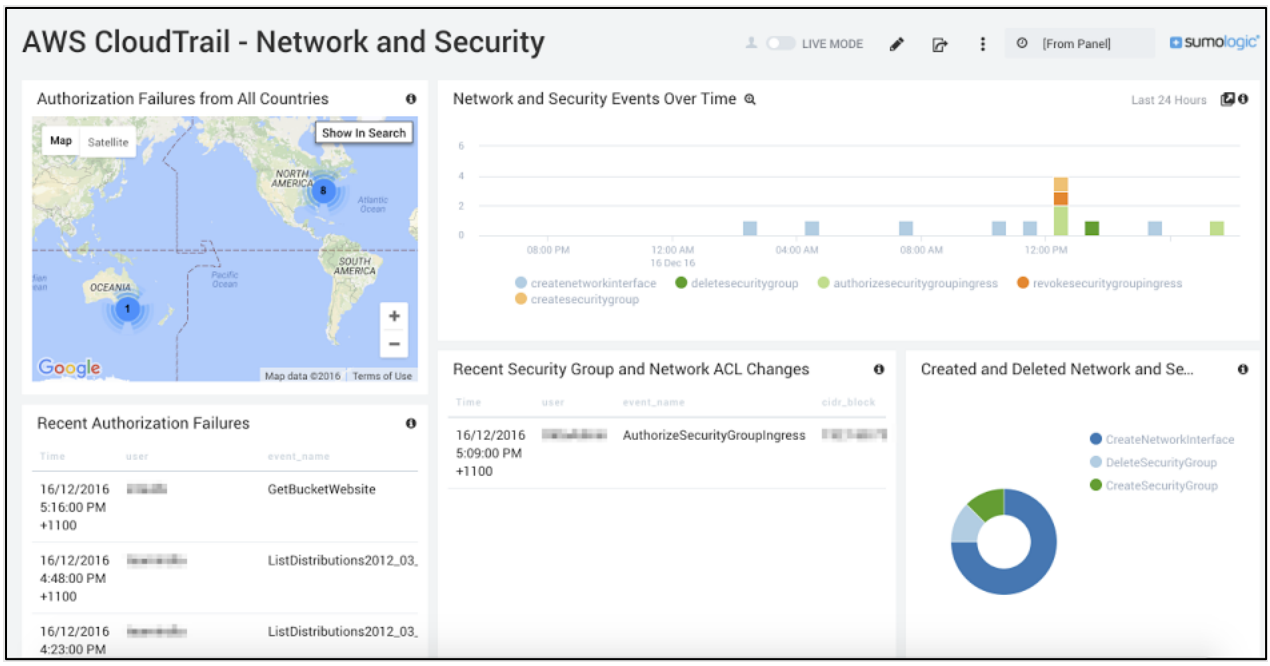

62. Is remote monitoring, management, or administration of the data centre performed by the vendor, and if so, is it from foreign countries? Does the same apply to any hosted software or services?

AWS does network monitoring and management (e.g., against DoS attacks), but the Assetic system is also designed to scale out in these types of circumstances, and we do our own monitoring as well. AWS monitoring, management, and administration is done on a regional basis (https://aws.amazon.com/about-aws/global-infrastructure).

Brightly utilises SumoLogic and AWS CloudWatch for remote monitoring of the Assetic platform. We utilise AWS CloudFormation and related tools for full automation and deployment of the application.

63. Do you review security related audit logs? If yes, how frequently? Is this automated by a Security Incident Event Management (SIEM) program?

Daily monitoring of audit and security logs is carried out. Real-time alerts are also configured and automated for any major spikes in abnormal activity.

64. How is logging and auditing applied to the actions of any staff (whether directly employed or contracted through a supplier) with any access to systems, solutions, or infrastructure containing customer information?

All AWS activities are logged using AWS CloudTrail to a completely isolated AWS Audit account. The Assetic platform itself (in the admin area and also filtered at a record level) has audit log features, which record configuration changes, and record read/write/delete changes (i.e., who made the change, when the change was made, what module and what record changed, old and new value, and further system information outputs where applicable).

65. How are encryption keys managed?

We use AWS Key Management Service (KMS) to assist encryption key management on our cloud infrastructure. AWS KMS centrally manage cryptographic keys and control their use across a wide range of AWS services such as RDS, EBS, and S3, used by the Assetic cloud platform. AWS KMS is a secure and resilient service that uses hardware security modules that have been validated under FIPS 140-2, or are in the process of being validated, to protect the keys. AWS KMS is integrated with AWS CloudTrail, which provides logs of all key usage for compliance purposes.

66. Is Assetic hosted in a Tier 3 or Tier 4 data centre? Is an independent SSAE16 or CSAE3416 audit performed against the controls? Is Assetic's gateway environment certified against Australian Government standards and regulations? [e.g., several major Cloud Service Providers (CSPs) in Australia use gateways certified by the Australian Signals Directorate (ASD)]

Hosting is provided through AWS: https://aws.amazon.com/compliance/

Each AWS datacentre is Tier 4, and every availability zone (AZ) is made up of a cluster of connected datacentres. Audits are conducted against AICPA: AT 801 (formerly SSAE 16).

AWS is one of the cloud providers in the ASD Certified Cloud Services List: https://www.asd.gov.au/infosec/irap/certified_clouds.htm

67. What routine maintenance actions or operational procedures are required to assure system reliability? Are there updates/patches/hotfixes for the application, interfaces, and operating systems and how are they handled?

All Infrastructure deployment and maintenance is fully automated via AWS CloudFormation. All EC2 server instances are regularly refreshed (outside of business hours) and are deployed using AWS base images (which include the latest OS patches).

For the Assetic application, we will release hotfixes/minor improvements on a fortnightly cycle and major new features typically every quarter. Please refer to the Release Management section for additional information.

68. What is your strategy in the event of a security breach?

Brightly is committed to protecting the security of its customers' information and we take all reasonable precautions to protect it from unauthorised access, modification, or disclosure.

In the event that a physical, network, or application-level security breach occurs, Brightly will stop the breach as quickly as possible and at the first reasonable opportunity, advise customers' whose information is lost, stolen, accessed, used, disclosed, copied, modified, or disposed of by any unauthorised persons or in any unauthorised manner. Brightly has identified the relevant law

enforcement and regulatory authorities whom Brightly may need to contact in the event of a security incident.

Subsequent to a security breach, Brightly will undertake a review applying best practice forensics in investigating the circumstances and causes of the breach and make long-term infrastructure changes to correct the root cause of the breach to ensure that it does not recur.

Load Management & Performance

69. How is load management dealt with and reported?

Brightly's SaaS platform has been developed utilising a scale-out architecture, which is a common approach for modern cloud platforms. A scale-out architecture enables Assetic's application, services, and database layers to be highly scalable, due to the fact that additional computing resources can be continually added to expand the available processing and memory power of the application.

For application and services stacks, we utilise separate AWS Auto Scaling groups to maintain application and service availability. AWS Auto Scaling allows Assetic to scale EC2 server capacity up or down automatically. AWS Auto Scaling is configured to automatically increase EC2 instances based on demand spikes, and additionally a larger number of fixed EC2 instances are by default available during typical usage hours (6:30 a.m. – 8 p.m.). Reporting is configured via auto-alerts within AWS Auto Scaling, and we also utilise SumoLogic for performance, security, and error logging and analysis.

70. How does Assetic provide event and error logs for troubleshooting and root cause analysis?

The log collected by SumoLogic is correlated with other system metrics to detect system and performance issues. The information can also be used to perform root cause analysis.

71. What metrics are in place to determine the source of performance issues?

Some of the key metrics that we monitor across all servers are CPU usage, memory usage, and number of requests per second.

We also actively monitor for any exceptions being generated in order to proactively address any system issues being detected.

We utilise several layers of cloud monitoring tools in order to carry out monitoring. These include a combination of AWS CloudWatch and SumoLogic.

Some of this information is provided in our system status page: (http://status.assetic.net)

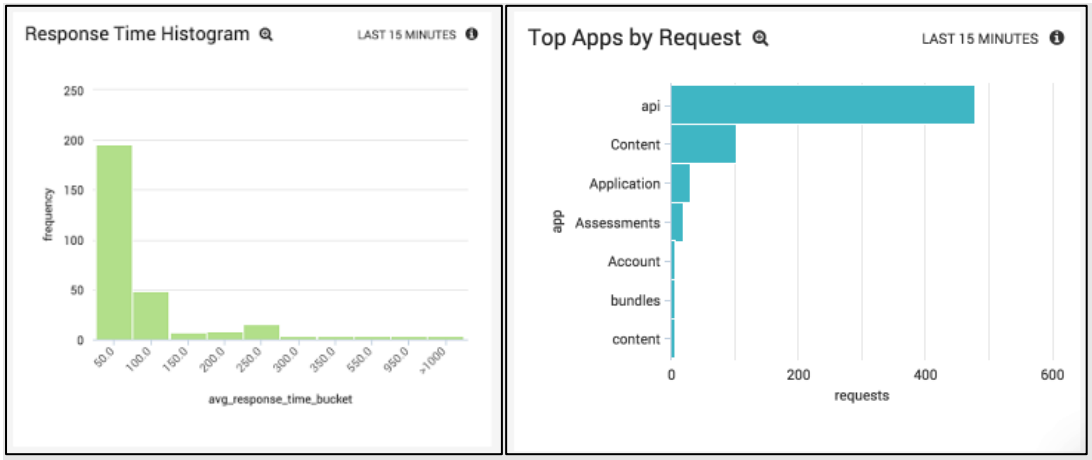

Overall response times and requests across all services in the last 15 minutes

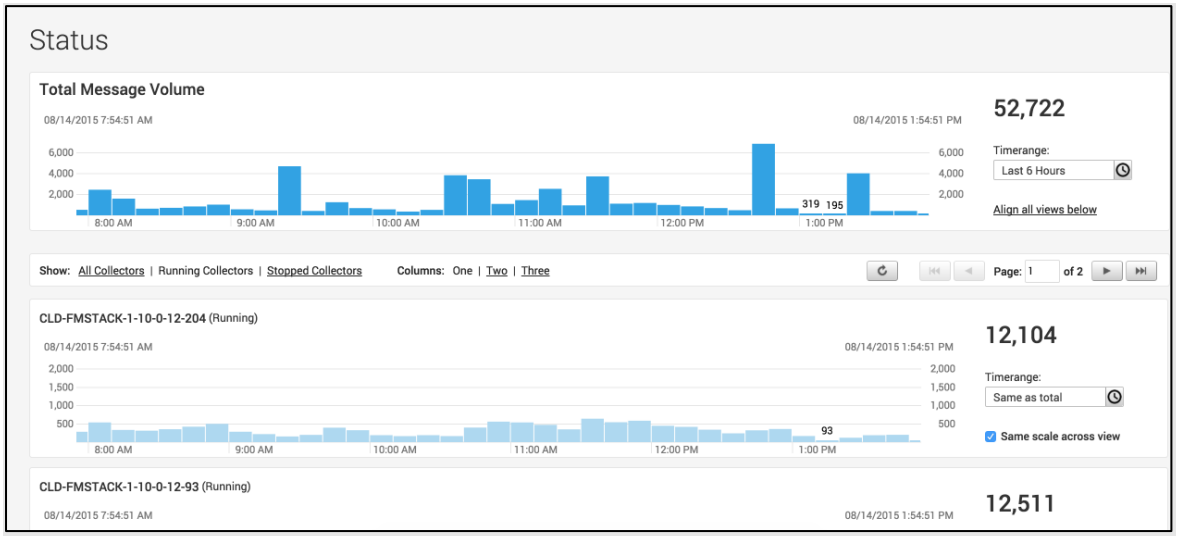

Summary of log messages across all services in the last six hours

Database Management

72. What is Assetic's schema approach and platform?

Assetic's cloud platform utilises an SQL database cluster and a normalised database schema. The database cluster runs across multiple availability zones to ensure real-time backup of data and automated database failover.

We also maintain multiple searchable index (Elasticsearch) that is used by leading cloud providers like Twitter and Wikipedia. This enables Assetic to provide extremely fast performance when carrying out system searches and aggregation.

73. How is database maintenance managed (Consistency checking, index defragging, statistics updates etc.)?

Database upgrades are managed through DACPACs and seeding data. These form a core part of our release process (refer Release Management Process). DACPACs are used for controlling schema upgrades. A data-tier application (DAC) is a self-contained unit for developing, deploying, and managing data-tier objects. DACs make it easy to monitor database changes that get made and also make it easy to package database changes alongside new features developed.

Seeding data is a series of SQL scripts that Brightly maintains in order to keep core system data up to date.

Industry-specific seeding SQL scripts are also utilised by Assetic so that when a new environment is commissioned, we can provide a fully packaged initial industry configuration that customers can then maintain to suit their individual requirements.

Consistency checking is applied to the system seeding data that Brightly has the sole right to fully maintain. The SQL MERGE function is employed to perform an update, insert, and delete of data based on the primary keys of both the source and the destination table to maintain data consistency.

All other consistency checks are based on the primary and foreign key constraints in the user tables.

Index defragging and statistics updates are done on a nightly basis via a SQL server agent job. The threshold to trigger these processes is when the index defragmentation or statistics are outdated by more than 10%.

The database log file is regularly trimmed on a nightly basis to keep its size to about quarter of the data file size.

74. Are the SQL Server instances in EC2 as opposed to RDS?

Database instances are set up in RDS with mirroring (active-passive) across multiple availability zones.

Release Management

75. How are upgrades managed? Will release notes be provided before changes are applied? What is the expected release schedule to the system (frequent small changes, quarterly revisions, etc.)? Will changes be applied to a test environment first?

Updates to Assetic are made fortnightly (i.e., frequent small changes) as this helps to minimise software errors and any security vulnerabilities over the long term.

Note that major new product features are managed through alpha/beta feature flags in the software. Brightly's asset management staff and key pilot customers are consulted throughout development and testing of these new major new features before they ever make it through to production release.

Upgrades are managed through a four-stage deployment process:

1) Build and deploy on a **nightly** environment during software development
2) Deploy to a **preview** environment for Assetic product owner/customer beta testing (weekly basis)
3) Deploy to a **staging** environment to run end-to-end automated unit and integration tests
4) Deploy to **production** once all automated and manual Assetic QA, testing, and sign-off occurs

When the SaaS platform is deployed within the nightly environment, we run a continuous integration process that automatically runs all unit tests across the entire system (object-level testing) and also integration tests across the entire system under various roles (UI-level testing). This enables us to rapidly release new features, as well as ensure that bugs haven't been introduced

that affect existing software features.

We also run the integration test suites whenever a new release is deployed within the preview, staging and production environments. This is a fully automated process in order to ensure that deployments have been applied correctly. Manual QA and testing by Brightly also occurs throughout all deployment phases as a safeguard.

We provide notice of upcoming releases one week in advance via the status page (http://status.assetic.net). Release notes are also provided for all releases of the software within the Assetic Knowledge Base one week in advance of the release.

The diagram below outlines how the release process works:

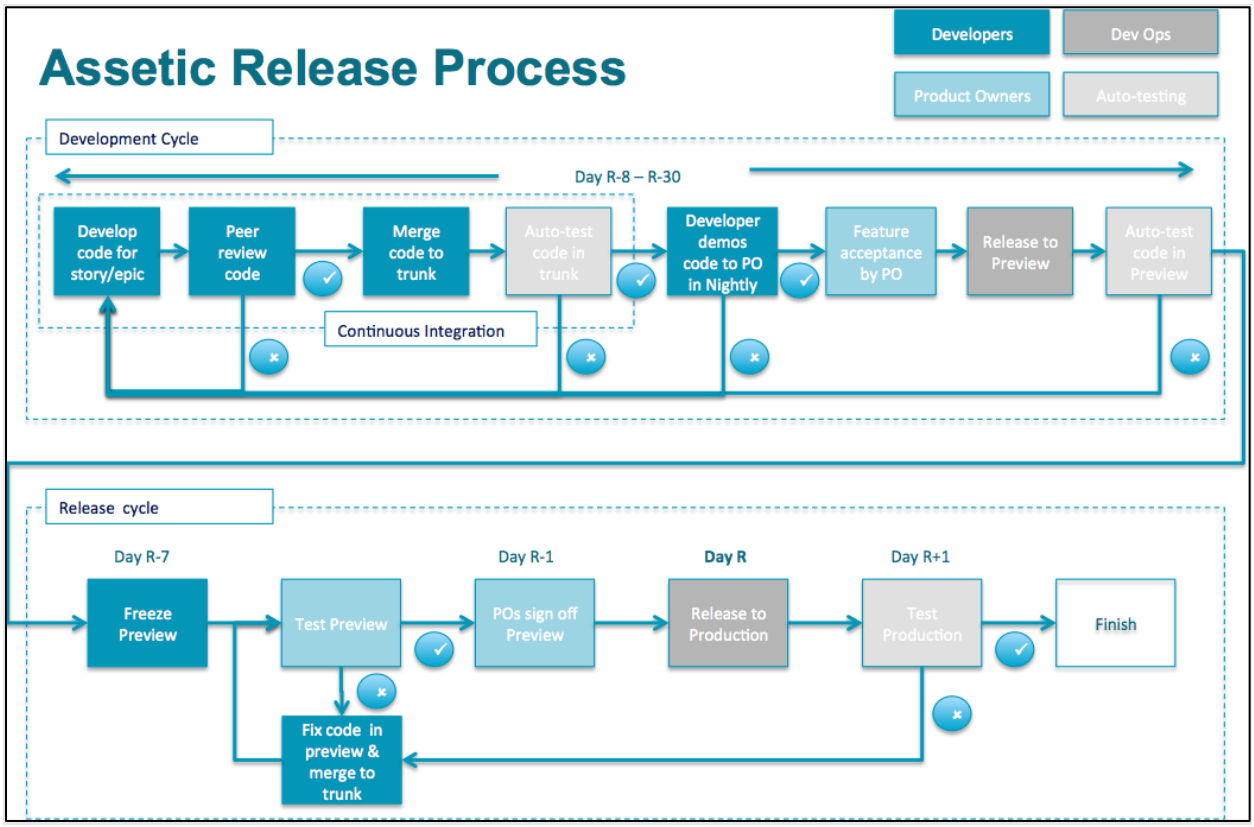

Backups and Disaster Recovery

76. What backup policies and procedures are in place (backup frequency, retention periods)?

Brightly maintains backups as follows for all production environments:

- **Databases are deployed in a mirrored setup (active-passive) across multiple availability zones. In the event of a node/data centre failure, the cluster will automatically failover to the backup server node**
- **Nightly backups of the entire database server are taken**
- **Individual database backups are created at a file-level every 24 hours and stored on AWS S3**
- **All document objects are stored on S3, which has high availability to cover node/data centre failures**

- **Database and document backups are copied to Google cloud platform daily as an offsite backup**

Full daily database and file backups are persisted with AES encryption. The daily backups are kept both at AWS and also stored remotely on the Google cloud platform; they are transmitted via HTTPS (TLS 1.2) offsite and stored with AES encryption.

The full daily backups are retained for at least the past 14 calendar days.

Brightly also provide mechanisms for agencies to download and backup information locally and we make it clear that agencies own their data.

77. What is the process for recovering deleted data (database and documents)?

- **An end-user administrator can recover deleted records via the front end**
- **Documents are not permanently deleted (only a database reference gets deleted)**
- **There are also a number of layers of data protection that are put in place:**
    - **The system has been designed at a role and schema level to minimise the chance of data corruption/deletion (i.e., many functions are transactional, such as condition assessments, asset accounting transactions, etc. and access control mechanisms minimise likelihood of data corruption/malicious damage)**
    - **Assetic provides an audit log of all changes made within the system. The log includes system processes/functions executed and details of records added/edited/deleted. It logs who made the change, when it happened, old and new values, the module that called it, and whether it's a user/system/integration related change**
        - **The system audit log can be searched at a whole-of-system or record level (e.g., when viewing an asset)**
        - **Records are disabled (not deleted), so can be recovered by client administrators**
- **Brightly can upon special request recover from daily database/document backups**

78. What is the process for rolling back a change if there is an error encountered, for example, with a bulk data upload?

At an individual transaction level where serious data integrity issues may occur (e.g., accounting transactions where several DR and CR transactions must succeed for the entire database transaction to be valid), we utilise RDMS to ensure transaction rollbacks occur in the event of any one DB add/update failing.

For bulk accounting transaction there is a pre-close process that allows all transactions to be checked before any journals are posted. Both the accounting pre-close and close processes can be resumed in the event of the bulk upload transaction failing midway (e.g., service outage).

For bulk transactions handled through the data exchange module (e.g., for updating/populating an entire asset register), the software caters for errors as follows:

- **Allows you to insert values into dropdown lists (or you can trigger errors in the event that data doesn't match pre-configured dropdown values)**
- **Allows you to save the import as a template so that it can be reused**
- **Runs as a background process and processes data in batches (showing data upload progress)**
- **If all data uploads successfully, a confirmation message is displayed**
- **If there are errors with particular data rows in the upload file, then:**
    - **A CSV error file is created as per the original upload file format (i.e., same column structure)**
    - **The error file is a subset of the original file that only shows rows that contained errors (also includes a column at the end describing the errors encountered)**
    - **The CSV error file can be downloaded, the errors corrected, and then the import can be run again**

You can also process the entire bulk upload again if there were no data validation errors encountered but you wanted to correct wholesale errors that existed in the original data file for some reason.

There are many examples where you want data to reside in the system, even if that data has subsequently been corrected, which is where a combination of software functionality, the system Audit Trail and appropriate database relationships are utilised, for example:

- **In the accounting module, there is a reversal transaction if you want to correct data for an asset that has had its financial data closed and reported on in a previous period(s)**

Service criteria information (e.g., condition) is structured tabular data that shows changes in service criteria data over time, but the most recently approved data is used for reporting and valuations purposes.

79. Can Assetic be failed over to servers/databases already running at a remote site, thus eliminating the need for any sort of restore procedure?

Yes, there are multiple high availability zones running in parallel for fail over (i.e., no restore procedure is required). These are across physically separated data centres.

Failover to another AWS region is possible via AWS CloudFormation and we have other regional stacks running that are available for DR.

80. What DR strategy do you have in place for Assetic and how is it actioned?

Assetic has been designed as a high availability cloud application to meet business continuity best practice. All application layers have been designed to scale out across multiple AWS high availability zones (i.e., physically separated data centres), which provides protection against data centre failures/outages and also server/service failures.

Assetic utilises at least two AWS availability zones to ensure the Assetic application is highly available. These zones operate as follows:

- **Each availability zone runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable**
- **Common points of failures like generators and cooling equipment are not shared across availability zones**
- **Additionally, they are physically separate, such that even extremely uncommon disasters such as fires, tornados, or flooding would only affect a single availability zone. When one zone fails, the application is failed over to the other one automatically**
- **Full daily data backups are encrypted with AES and stored remotely (along with the entire SaaS platform) in the event both Amazon availability zones are down**

In the event of a catastrophic failure across both AWS availability zones, the following process applies:

- **Utilise the Assetic SaaS application stack at another global AWS high availability zone (there are other 'warm' stacks available and stack deployments can be fully automated via Cloud Formation if required)**
- **Recover from the current full daily database and document backups stored remote from AWS**

In the event of a catastrophic failure of AWS, the following process applies:

- **Spin up the Assetic SaaS application within Google cloud platform data centres.**
    - **The platform has previously been tested to run on other cloud providers (e.g., Microsoft Entra ID and Google Cloud Platform)**
    - **Some services may need patching to restore 100% of Assetic's functionality (i.e., any that are AWS-specific)**
- **Recover from the current daily full database and document backup stored remote from AWS**
- **Import a backup of AWS Route 53 DNS data within Google Cloud DNS**