



Ever wondered what's under the hood of Advanced Analytics?

Looking for more information before making the leap? Learn more about the system architecture of Advanced Analytics and the benefits it serves as an integrated solution.

A laboratory's day-to-day operations are recorded through STARLIMS, accumulating data over time. This data includes sample lifecycle information and metadata.

A range of insights can be measured from this data. Here are just a few examples:

- Calculate the average time it takes to perform a task.
- Analyze the breakdown of patient condition by age group.
- Do data comparison, trending, and forecasting.

The best way to get insights out of this data is by rendering it visually, intuitively, and interactively. This is exactly what Advanced Analytics does. It adds the right toolkit to the STARLIMS v11 solution, allowing you to intuitively visualize and interact with your data to maximize the total value of your data.

INTEGRATED SOLUTION

Architecture

There are four components to the STARLIMS Advanced Analytics solution.

The first two components are already a part of the base STARLIMS ecosystem. As Advanced Analytics is designed to complement your STARLIMS solution, the first component must be STARLIMS itself. As Advanced Analytics is data-driven, the second component is the Database server.

The core component, the Advanced Analytics server, converts raw data into visual representations. Using decoupled processes like rendering, user access, backgrounders, and other supportive processes it does all the heavy lifting required in the process, including data aggregation, statistical calculations, and other advanced calculations. Because these processes are decoupled, the Advanced Analytics server can be configured in a high availability and/or scalable environment.

The fourth and last component is the Advanced Analytics builder, a desktop application used to design the dashboards that will be used by the end-user. This tool creates charts, calculates statistics, and builds interactive dashboards, all through an intuitive interface, so the user can concentrate on the task at hand – analyzing the data.

Viewing Content as an End User

The first workflow is the act of viewing and interacting with a workbook. This is the end-user experience of Advanced Analytics.

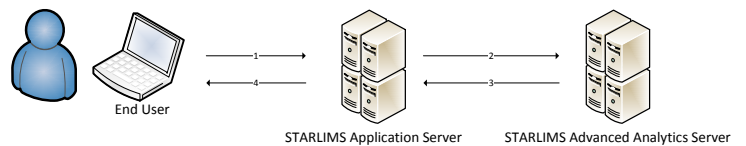


Figure 1 – View Workbook Overview

1. The end-user requests access to a workbook.
2. STARLIMS checks with the Advanced Analytics server, and determines whether the end-user has access.
3. Advanced Analytics sends the workbook URL to STARLIMS.
4. STARLIMS sends the Advanced Analytics server Workbook URL to the end-user, which is rendered in an IFRAME

This process is thoroughly detailed in the security section under trusted authentication.

Publishing New Content

The second workflow is the act of publishing workbooks for others to view. This is when the builder is used to create new content.

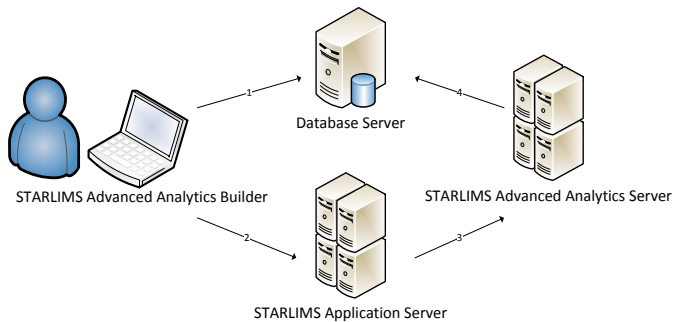


Figure 2 – Publishing Infrastructure Overview

1. The user connects to the STARLIMS database, through the builder, to create a workbook.
2. The user adds the workbook to STARLIMS and configures access per role.
3. STARLIMS synchronizes the workbook with the Advanced Analytics server.
4. The Advanced Analytics server queries the database to refresh the data.

SECURITY

Single Sign-On

Viewing content and publishing content is always done through STARLIMS, using your STARLIMS credentials. All users are managed through STARLIMS. Even in the case of LDAP authentication, each user will have a corresponding user record in STARLIMS. This is done through a new feature association module which allows an administrator to license users by specific features. Advanced Analytics end-user licenses are assigned through this module.

Once a user is licensed, a user is created automatically on the Advanced Analytics server with a randomly generated password (36 characters long), making it very difficult to hack into that user account. Users do not need to connect to the Advanced Analytics server since STARLIMS bridges the gap through trusted authentication.

As such, a licensed Advanced Analytics user has direct access to Advanced Analytics by logging into STARLIMS.

Trusted Authentication

The Advanced Analytics server is configured to trust specific servers, generally limited to the STARLIMS Application Servers. When STARLIMS sends a request on behalf of a licensed user, the Advanced Analytics server trusts the sender, knows the user, and generates a one-time usage ticket used in the URL to render the visualization. The request/authentication process goes through 6 steps:

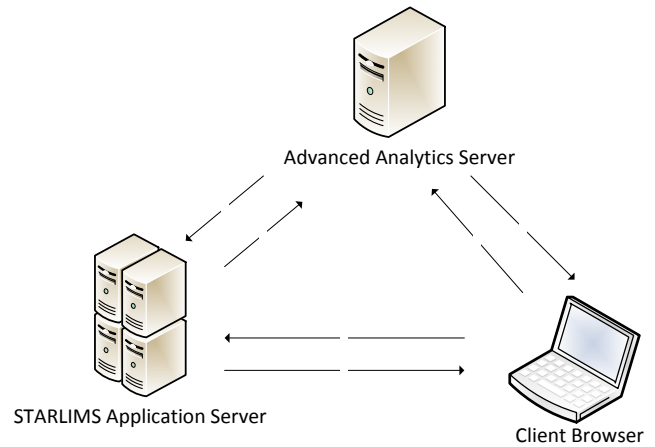


Figure 3 – Trusted Authentication

1. The user attempts to view a dashboard. Only a user authenticated in STARLIMS with rights to access the dashboard will be able to proceed beyond this step.
2. STARLIMS sends a request to Advanced Analytics Server. STARLIMS knows whether the user is licensed, authenticated, and if access to that dashboard should be granted. If the user is licensed, authenticated, and has permission to access that dashboard, STARLIMS will send a request to the Advanced Analytics server.
3. Advanced Analytics Server creates a ticket. The AA server checks the IP address or host name of the STARLIMS application server. If it is set up as a trusted host, then a ticket is created in the form of a unique 24-character (URL-safe, Base64-encoded) string. It responds to the request with that ticket. If there is an error and the ticket cannot be created the Advanced Analytics Server responds with a value of -1, often indicating an untrusted host or unlicensed user.
4. STARLIMS passes the URL to the client. STARLIMS constructs the URL for the view and sends it back to the client for rendering. The ticket is included (e.g., `http://aaserver/trusted/<ticket>/views/requestedviewname`).
5. The browser requests the view from the Advanced Analytics server. The client web browser sends a request to Advanced Analytics server using the URL with the ticket.
6. The Advanced Analytics server redeems the ticket. The server sees that the client requested a URL with a ticket and redeems the ticket. To prevent session hijacking, tickets must be redeemed within three minutes after they are issued. Once the ticket is redeemed, the Advanced Analytics server authenticates the user, removes the ticket from the URL, and sends back the final URL for the embedded view. The user's browser will then retrieve and display an interactive dashboard.

Managed within STARLIMS

Advanced Analytics enables customers to leverage their STARLIMS investment. As such, data visualizations added to STARLIMS (referred to as workbooks) are considered STARLIMS components and benefit from every standard feature the system provides to UI components, ranging from source control to transportability.

Source & Version Control

As with every other STARLIMS component, the Designer allows you to:

- Check-in and check-out dashboards
- Add comments to a dashboard upon check-in
- Revert to previous versions
- View check-in history

Multi-Language

Support of multi-language within information systems is always a challenge, especially with user interfaces. For example, you may have to accommodate for languages which do not read left-to-right, top-to-bottom. Advanced Analytics addresses multi-language support of the user interface (the labels, not the data) through the following methods:

- When adding a dashboard to STARLIMS, it is added to the system's default language and becomes the default version which will be viewed by everyone.
- Designers have the ability to create a copy of the dashboard, change the layout, relabel everything, and add it to a different language within STARLIMS.
- End users logged in to the system will see their workbooks based on their configured language.
- Each language has its own source-control, allowing designers to update the workbooks for specific languages without affecting others.

Transportable (Export/Import)

There are typically three separate environments to support development, staging, and production. It is crucial to have seamless transition of newly developed workbooks from development to staging to production.

Developers who use the STARLIMS Designer are already able to export and import STARLIMS components as STARLIMS Deployment Packages (.SDP). Advanced

Analytics Workbooks can also be exported along with other components using STARLIMS deployment packages using the same method.

Role-Based Access

Using the full-featured role-access system included in the STARLIMS's core, administrators can:

- Grant and deny access to workbooks per role.
- Define default workbooks per role and device type (desktop, tablet, phone).

ANALYTICS AS A WEB CONTROL

Application-Driven Analytics

STARLIMS applications can be complemented with visual analytics related to the current data, allowing the user to make better decisions based on data visualizations. When seeing transaction oriented data, you can see analytics on the big-picture data alongside the granular transactions details.

For example:

- A run with a lot of samples and results could have these result plotted on a chart allowing the user to spot any outliers as early as possible.
- An instrument or method could display historic data, so the end-user can spot trends and patterns.
- When assigning samples, a lab manager could see the historic and expected test and sample distribution in the lab, and make his assignments based on availability

Analytics-Driven Applications

In addition to embedding data visualizations in STARLIMS applications, Advanced Analytics also provides the ability to launch a STARLIMS application from a dashboard view. Using analytics as an entry point into an application can be powerful. For example, QA can look at the last few month's out of specs results, and, upon seeing an upward trend, select a specific peak area, opening the STARLIMS application with the selected point's data. (Figure 5)

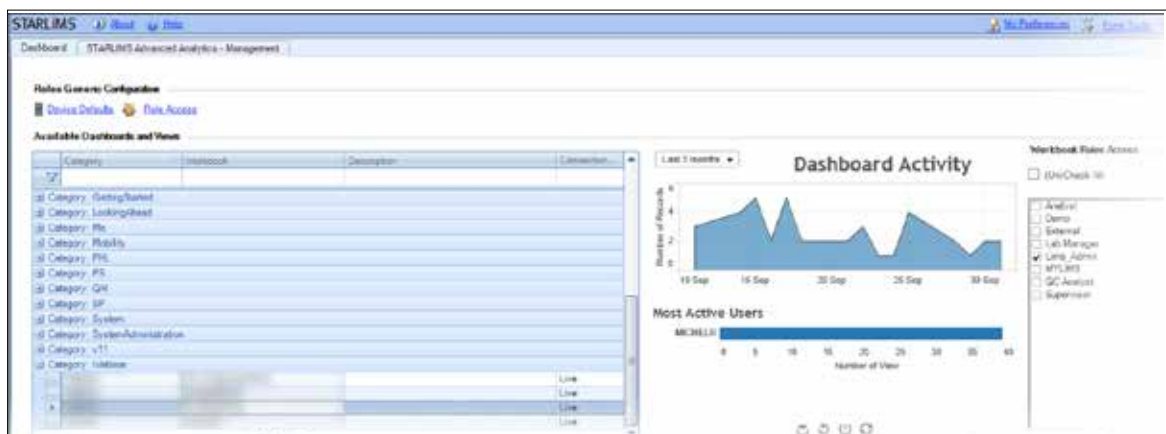


Figure 4 – Application-Driven Analytics – Dashboard Activity Example

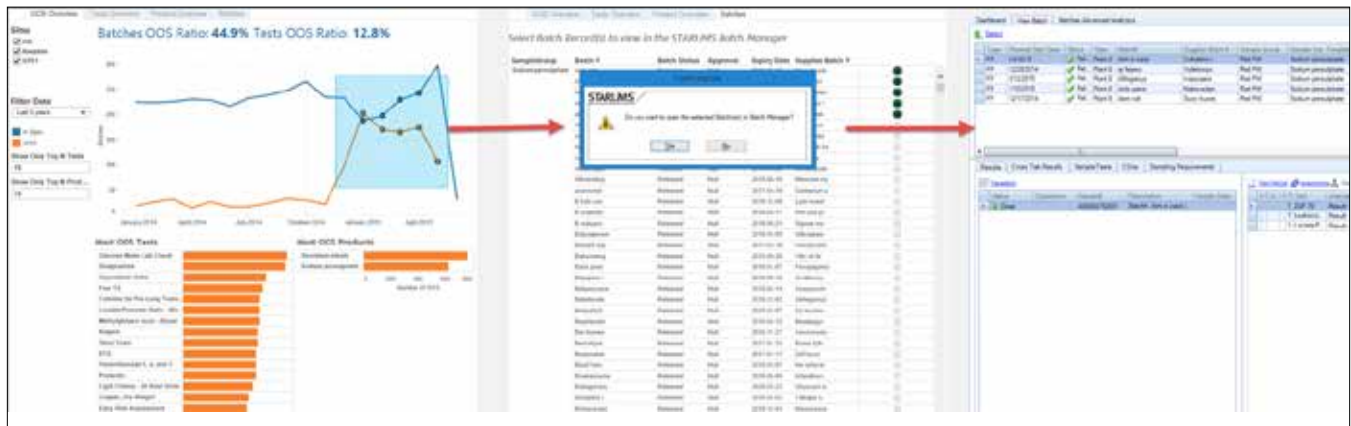


Figure 5 – Analytics-driven Applications

Mobile Analytics

Sometimes insights from analysis need to be shared and communicated. Whether this needs to be done in a conference room or while traveling, users can access the dashboards from any device as long as it is connected to the same network as STARLIMS and the Advanced Analytics server. By appending /AA to the STARLIMS application URL, users can connect with their STARLIMS credentials and dive straight into their dashboards. (Figure 6)



Figure 6 – Advanced Analytics on a Tablet

LEVERAGE STANDARD STARLIMS FEATURES

These features allows us to enhance the benefit of Advanced Analytics with minimal effort by integrating existing STARLIMS functionality into the system.

Scheduler

- You can extract data from STARLIMS and store it within Advanced Analytics' proprietary data engine. This speeds up the data access, but makes it persistent. This data can be refreshed periodically through the standard STARLIMS scheduler, configurable per workbook.

- Users can receive notification messages with PDF exports of the workbooks at a scheduled time. For example, directors may want to receive a monthly overview of the lab in their mailbox.

Email Distribution

Administrators and configurators can create email distribution lists from the Advanced Analytics Management application. These list can be re-used and modified as necessary, making it more convenient to configure multiple dashboard distributions. Recipients get a PDF export of the dashboard, allowing them to be aware of the current state of the lab without having to connect to the system.

Trigger User-Defined Actions

Sometimes scheduling of a refresh might not be enough. It might be necessary to perform specific actions periodically before refreshing extracts or emailing exports. More advanced users are able to write custom scripts, using the STARLIMS scripting language. These scripts have priority and will run before refresh and email distribution. A good use case for this would be when an extract requires STARLIMS to run a script to import data from an external system.

ADVANCED ANALYTICS SERVER

Hardware Requirements

STARLIMS Advanced Analytics hardware requirements vary based on the number of concurrent users. To better understand the hardware requirements, some definitions are required.

Licensed User - A user who has been licensed for Advanced Analytics. That user exists on the AA Server.

Concurrent Users - A licensed user who is currently interacting with a visualization. This means the user is filtering, hovering or selecting data points, loading a dashboard, or any other activities beyond simply viewing a dashboard without moving the mouse.

In our recommendations, we assume that roughly 10% of the licensed users are concurrent. Each organization may have a higher or lower ratio of concurrent users, so use this as a guideline only.

While both 32-bit and 64-bit versions of the AA server are available, STARLIMS recommends using the 64-bit version on a 64-bit OS.

Environment	Hardware Configuration	Concurrent	Licensed Users
Proof of concept, development or evaluation	4GB RAM Dual Core 2.0GHz 2GB free disk space	1-2	5 – 10
Server – Small	8GB RAM Quad Core CPU 2.0GHz 5GB free disk space	Less than 25	10 – 200
Server – Medium	32GB RAM 2xQuad Core CPU 2.0GHz 50GB free disk space	Up to 100	200 to 1000
Server – Large *	64GB RAM 8 Cores 2.0GHz 100GB free disk space	Over 100	Over 1000

* For enterprise deployments, contact STARLIMS for a more precise estimate.

Trusting STARLIMS Servers

For Advanced Analytics to run, the network on which it runs must allow the communication between the client, the STARLIMS server, and the Advanced Analytics server. By default, the Advanced Analytics Server will run on port 80, or, if it is not available, it will use port 8000.

When trusted authentication is configured, the Advanced Analytics server will perform lookups on the sender host (or IP), and will grant tickets only when such requests come from a trusted host. This mechanism is what enables single-sign-on between STARLIMS and Advanced Analytics and is a mandatory configuration in the installation process. All steps are described in the installation manual.

It is possible to enhance the security of the Advanced Analytics server by enabling SSL. A valid SSL certificate from a recognized authority is recommended. Self-signed certificates can also work but are not recommended.

Scalability & High Availability

The Advanced Analytics server consists of multiple processes dedicated to specific tasks, such as the graphic renderer to convert data to charts, the web server to expose the features over HTTP, the backgrounder for all scheduled tasks and jobs, the cache server for optimal performances, the data server for talking to the databases, the data engine for manipulating the data in a uniform way, the repository to manage the workbooks, and the gateway to manage communication. Each serve a different purpose and will be under different stress. Depending on how Advanced Analytics is used in an organization, it might be necessary to scale up all or some of these processes. It is then possible to add new servers to the cluster and configure them to run only specific processes.

For high availability, the entry-level is to install a three-node system where there is a gateway server, a primary server, and a backup server, which will run if the primary goes down. The gateway's only purpose is to route the users to the right place and do the easy tasks like handle licensing, host the search engine, and coordinate the synchronization of the backup server.

The next level in high availability is to add a load balancer, where all the three nodes have their own gateways.

Lastly, a backup primary adds a safeguard to your system. It will be on standby, and, if the primary system goes down, the backup can be brought online quickly.

PERFORMANCE

In-Memory Data

The technology behind Advanced Analytics leverages the hardware to its maximum by detecting the hardware it runs on, allowing utilization of the complete memory hierarchy from disk to L1 cache.

Columns vs. Rows

The column-based representation of the data uses compression that supports execution of queries without decompression. Every data connection, either live to SQL server or Oracle database or data extracts, uses the same columnar structure of the Advanced Analytics data engine, and thus benefits from the performance boost the engine provides.

Live Connections vs. Extracts

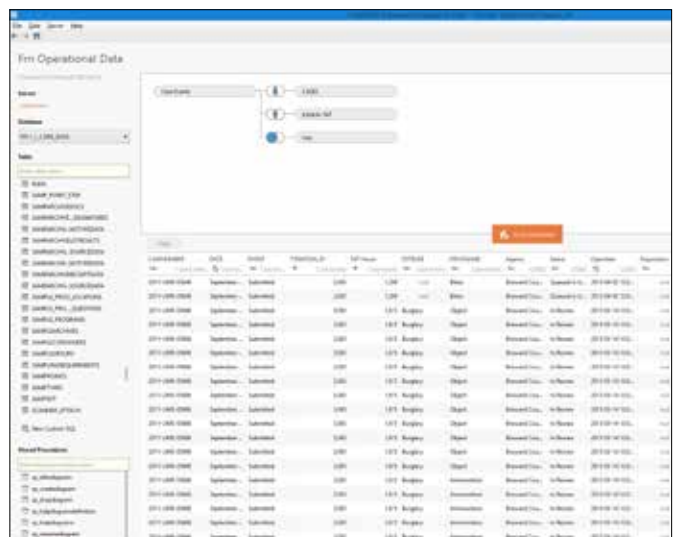
As to not affect day-to-day operations, it is possible to flag a data connection to use an extract. An extract is a copy of the data residing directly in the data engine, providing faster response time as no communication to the database happens until the extract is refreshed. This facilitates usage of huge amounts of data with instant response time. This data can be refreshed periodically.

Sometimes real-time data is necessary, requiring a live connection to the database. Even then, the data will be transformed and cached in such a way that the visualizations will be responsive and near instantaneous.

ADVANCED ANALYTICS BUILDER

Connecting to STARLIMS Data

The Builder provides a built-in query builder, which will allow you to connect directly to your STARLIMS database



The query editor allows you to build the data sources intuitively. The data connection builder allows the usage of:

- Tables
- Views
- Stored Procedures
- Custom SQL Queries

Each of these objects can be dragged on the editor and joined one with another, expanding the flexibility of a typical query builder.

The query builder will not let users connect to non-STARLIMS databases as STARLIMS is tailoring the Advanced Analytics solution to the STARLIMS software and is targeting laboratory processes, metrics, and analysis. However, STARLIMS is a very flexible platform that can easily import data from any external system, as such, making the Advanced Analytics Server able to use data from virtually any other system.

Connecting to External Data

Advanced Analytics connects only to STARLIMS databases. If an organization needs to connect to external data sources, there are a few options.

If the external data resides on the same kind of server (SQL server or Oracle), then a “bridged view” can be created in the STARLIMS schema. That view will be available while building a query as any other table would.

Otherwise, it is possible to create new sets of tables using the STARLIMS designer, and write a migration script to transfer the external data into the newly created STARLIMS tables. These tables will be available in the Advanced Analytics builder at design time. These scripts can also be scheduled within the STARLIMS Advanced Analytics management application to run before refreshing the workbooks, the right timing for the refresh to happen.

Summary

In this whitepaper, we explained the security concepts used with Advanced Analytics, the advantages of our integrated solution, the requirements behind the server (and the methodology to get to these numbers), and the data connection possibilities.

starlims.com

STARLIMS

Tel: +1 954 964 8663

4000 Hollywood Blvd, Suite 333 South, Hollywood, FL 33021-6755 USA

UNITED KINGDOM

Tel: +44 161 711 0340

GERMANY

Tel: +49 2302 915 245

FRANCE

Tel: +33 1 61 37 02 00

SPAIN

Tel: +34 91 902 50 69

NETHERLANDS

Tel: +31 72 511 8100

ASIA PACIFIC

Tel: +852 2793 0699

LATIN AMERICA

Tel: +1 954 964 8663

CANADA

Tel: +1 888 455 5467

AUSTRALIA

Tel: +61 3 9670 0678