

Alteryx Analytics Gallery

An Alteryx Server Use Case

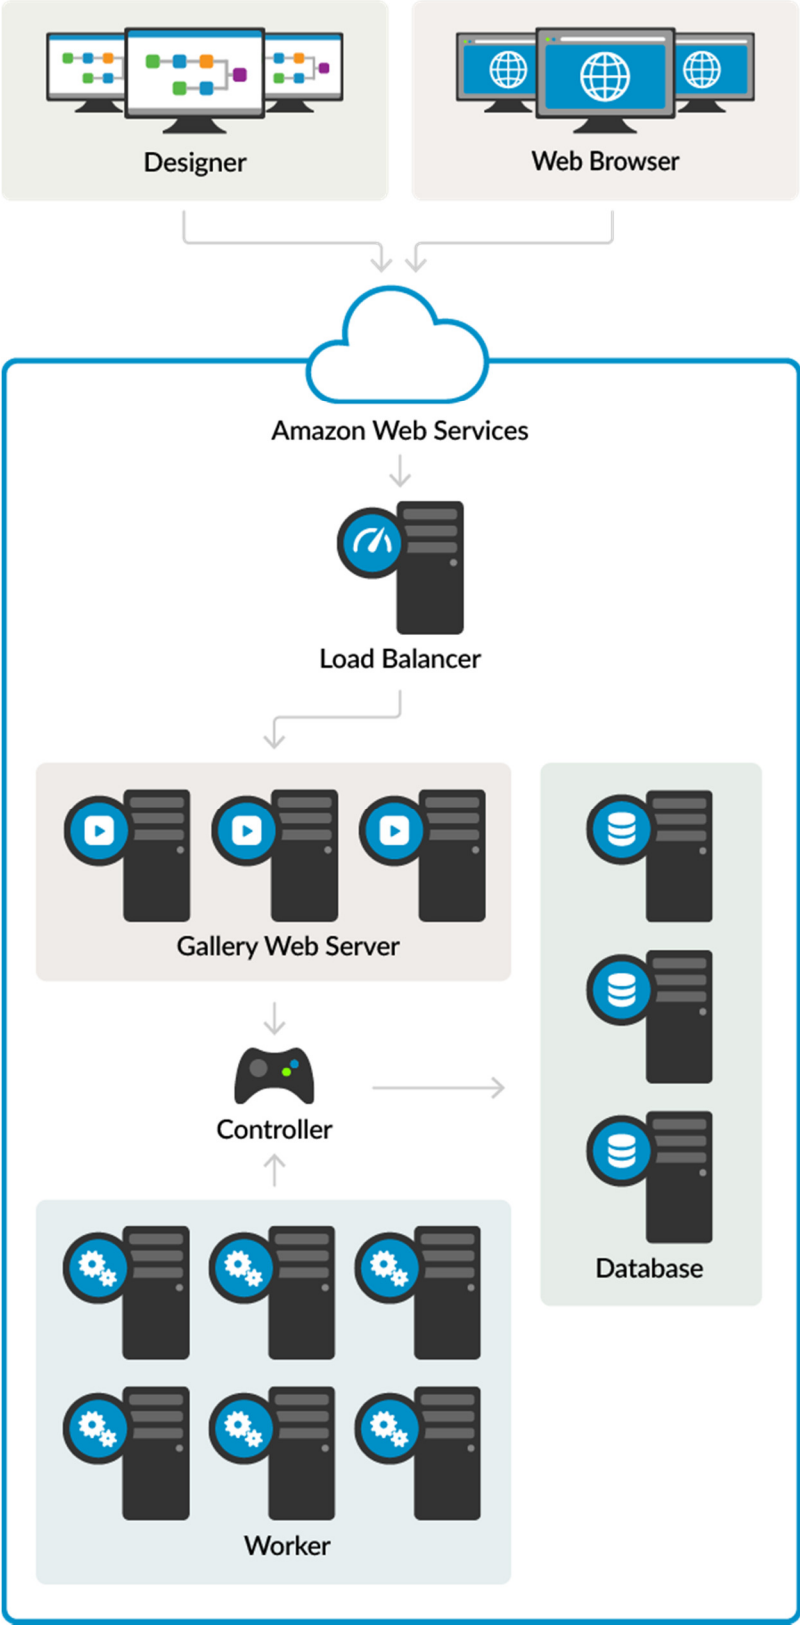
Overview

Alteryx Server provides a comprehensive and scalable server-based analytics solution that allows users to schedule and run workflows built via the Designer as well as to publish, share, and run analytic applications in a web-based environment via a Gallery. With Alteryx Server, companies can offer a private Gallery to their internal users hosted on their own server infrastructure. Alteryx uses Alteryx Server to offer a public Gallery (the Alteryx Analytics Gallery) so that any user can sign up and have access to sharing workflows, apps, and macros.

Architecture

The Alteryx Server can be configured with all components running on the same machine, or it can be configured as a multi-node instance by enabling certain components on multiple machines. The Alteryx Analytics Gallery (<https://gallery.alteryx.com/>) architecture includes the following:

- **Host** – Amazon EC2 (Amazon Elastic Compute Cloud) provides scalable computing capacity in the Amazon Web Services (AWS) cloud. With Amazon EC2, Alteryx can launch as many virtual servers (known as instances) as needed, configure security and networking, and manage storage.
- **Load Balancer** – AWS Elastic Load Balancing automatically distributes incoming application traffic across multiple Amazon EC2 instances. Alteryx uses a load balancer in front of the Gallery web server nodes to serve up the web pages and API functionality and to easily manage the varying levels of web requests.
- **Gallery Web Server Nodes** – The Alteryx Analytics Gallery is configured to work across three web server instances. The web server nodes take the incoming web requests from the load balancer and pass them on to the Service Controller when there is work to be processed. They are also responsible for the management of user accounts and the sharing of workflows and apps.
- **Service Controller** – The Controller manages the delegation of work to the Service Workers. Although Alteryx Server is deployed across multiple servers, only one of the servers acts as the Controller. When someone wants to run a workflow or app published to the Gallery, the web server communicates to the Controller to schedule the job to be run immediately.
- **Service Workers** – The Worker receives work from the Controller and is responsible for executing workflows and apps and retrieving the output. Alteryx has four Worker servers and each is configured with two processing threads (one thread per physical core for a total of 8 threads) in order to run more concurrent jobs. Alteryx also has two workers dedicated to processing map tile requests.
- **Database** – Alteryx Server includes a persistence layer to store information such as Alteryx application files, the job queue, and result data. Alteryx uses a standalone MongoDB Enterprise database for the Alteryx Analytics Gallery that is configured in a three node replica set. Data is written to the primary node and read from either the primary or one of two secondary nodes.



Technical Specifications

	Web, Controller, and Worker	MongoDB
vCPU*	4 x 2.6 GHz Intel Xeon E5-2670 V2	4 x 2.6 GHz Intel Xeon E5-2670 V2
Memory	15 GB	30 GB
Attached Storage Disk Size	600 GB (Controller and Worker only)	2 TB
Operating System	Windows Server 2008 R2 or later	Amazon Linux
Amazon EC2 Instance Type	m3.xlarge	r3.xlarge
SSD Ephemeral Storage	2 x 40 GB	2 x 80 GB

*A virtual CPU (vCPU), also known as a virtual processor, is a physical central processing unit (CPU) that is assigned to a virtual machine (VM). By default, virtual machines are allocated one vCPU each.

A vCPU in an AWS environment represents half a physical core.

vCPUs	Physical Cores	Worker Threads
1	1/2	1
2	1	2
4	2	4
8	4	8

Usage Data

Totals (for past 3 years since Gallery inception)

# of users who signed up	59,286
# of users who ran apps	3,739
# of apps that were published	3,949
# of apps that were run	147,612
# of times apps were run	245,711

Monthly Totals (for September 2015)

	Registered users	Alteryx employees	Totals
# of sessions	8,312	8,291	16,603
# of unique users who logged in	1,247	207	1,454
# of unique users who ran apps	451	89	540
# of new users*	3,704	44	3,748

*New users include anyone added to the Gallery (i.e. every license created by Fulfillment or via Trial and every user who signs up).

	Users who sign in	Anonymous users	Totals
# of times apps were run	10,628	1,845	12,473
# of unique apps run	353	113	466

Daily Totals (for September 2015)

Average # of users who logged in (sessions)	48
Average app run duration (in seconds)	24

Performance and Monitoring

Alteryx uses several tools for monitoring the Alteryx Analytics Gallery performance.

- **New Relic** – Provides visibility into critical system metrics, including CPU usage, physical memory, network activity, processes, and disk I/O utilization and capacity.

Alteryx uses New Relic to gain insights into the performance of the web server nodes. Dashboards in New Relic display real-time server processing metrics and transaction response times to assist with understanding usage patterns. Alteryx is alerted if response times are high or errors occur, such as with user login failures. New Relic’s availability monitoring helps Alteryx ensure the Alteryx Analytics Gallery site is always up and running and MongoDB query performance is maintained.

- **Amazon CloudWatch** – Collects and tracks metrics, collects and monitors log files, and sets alarms in order to provide system-wide visibility into resource utilization, application performance, and operational health.

Alteryx creates custom scripts to publish metrics to CloudWatch and then creates alarms within CloudWatch to send notifications for sustained state changes. For example, CloudWatch monitors the number of jobs in queue and invokes an alarm if the number of jobs exceeds the defined threshold. Additionally, CPU utilization is monitored – if it exceeds a specific threshold for a given period of time, an alarm notification is sent. Once notified, Alteryx can take action as appropriate.

Redundancy and Recovery Strategy

Amazon EC2 is hosted in multiple regions in different geographic areas and each region has multiple, distinct locations known as Availability Zones. Instances and data can be hosted in separate Availability Zones to protect against the failure of a single location. Alteryx has distributed server instances across multiple Amazon EC2 Availability Zones in the US East region. This ensures that if one instance fails, an instance in another Availability Zone can handle the requests.

With Amazon EC2, Alteryx can quickly increase or decrease server capacity within minutes as needed. If a Worker instance fails, the Controller can simply delegate jobs to another Worker or, if processing needs increase, Alteryx can easily spin up a new instance. If a Controller ever fails, a Worker can easily be configured to become a new Controller.

Alteryx follows best practices for deploying a MongoDB replica set to maintain data and provide high availability and redundancy. Alteryx uses Amazon EBS (Elastic Block Store) to make a backup of the database each night and keeps four rolling days of data for rollback.

Troubleshooting Resources

Alteryx Server contains functionality to produce logs of events such as services being started, shut downs, execution requests, etc. which can be helpful for troubleshooting issues. Alteryx systematically syncs the log messages from all of the servers every five minutes and saves them on Amazon S3 for easy retrieval if necessary.

Configuration Settings

General

Number of Servers:	3 web nodes, 1 controller, 4 workers, 2 map workers, 1 primary DB, 2 secondary DBs
Failover Strategy:	If a server fails, a new one is created with a backup clone.
Scaling Strategy:	As usage increases, a new server is spun up.
Data and Plugins:	Alteryx Internal Data Sets & R Packages (Installed on Service Layer Machines)

Controller

Storage:	SSD Ephemeral Storage is used for log and temporary files.
Persistence:	A standalone MongoDB database configured in a three node replica set is used.

Worker

Storage:	SSD Ephemeral Storage is used for log and temporary files.
Run As User:	All workflows and apps are run as a standard level user account.
Mapping Options:	Two servers are configured to act as a Map Worker to render map tiles.

Gallery

Storage:	SSD Ephemeral Storage is used for log and temporary files.
SSL:	An SSL certificate is installed on the Load Balancer – SSL is not enabled on each web server.
Authentication Type:	Built-in Authentication is used and users can log into the Gallery with an email address and password.
SMTP Settings:	Amazon SMTP Service is used for email-based features and notifications in the Gallery.

Engine

Processing Threads:	Each Worker server is configured with two processing threads (one thread per physical core for a total of 8 threads).
---------------------	---