

BigML Private Deployments Overview

The BigML Team

Version 3.2



MACHINE LEARNING MADE BEAUTIFULLY SIMPLE

Copyright© 2019, BigML, Inc., All rights reserved.

info@bigml.com

BigML and the BigML logo are trademarks or registered trademarks of BigML, Inc. in the United States of America, the European Union, and other countries.

No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without the express written permission of BigML, Inc.

Last updated September 17, 2019

About this Document

This document provides an overview of the Private Deployment options offered by BigML.

This document is part of a series of documents that explain all aspects of BigML Private Deployments:

- **Overview:** A high-level overview of BigML Private Deployments and the deployment options.
- **Pre-Installation:**
 - **BigML-Managed:** A pre-installation checklist for BigML-Managed Private Deployments.
 - **Self-Managed:** A pre-installation checklist for Self-Managed Private Deployments.
- **Self-Managed Installation:**
 - **Single-Instance:** Installation steps for a Self-Managed Single-Instance Private Deployment.
 - **Multi-Instance:** Installation steps for a Self-Managed Multi-Instance Private Deployment.
- **Self-Managed Maintenance:**
 - **Single-Instance:** Maintenance guide for a Self-Managed Single-Instance Private Deployment.
 - **Multi-Instance:** Maintenance guide for a Self-Managed Multi-Instance Private Deployment.

For further details or questions please contact us at info@bigml.com

Contents

- 1 What is BigML? 1**
- 2 BigML Private Deployments 2**
 - 2.1 Choosing BigML Lite vs. BigML Enterprise 2
 - 2.2 Choosing a Private Deployment Type 3
 - 2.2.1 Managed VPCs 4
 - 2.2.2 Self-Managed VPCs 6
 - 2.2.3 On-Premises 7
 - 2.3 Choosing a Cloud Provider 8
- 3 BigML Architecture 9**
 - 3.1 Web-Based Front End 9
 - 3.2 Middle End (REST API) 10
 - 3.3 Distributed Machine Learning Back End 10
 - 3.4 Storage Types 11
 - 3.5 Authentication 12
- 4 Deployment Planning 13**
 - 4.1 Deployment Sizing 13
 - 4.2 Deployment Options 14
 - 4.3 Pre-Deployment Steps 14
 - 4.3.1 Managed VPCs 14
 - 4.3.2 Self-Managed VPCs 15
 - 4.3.3 On-Premises Deployment 15
- 5 Deployment Steps 16**
 - 5.1 Managed VPCs 16
 - 5.2 Self-Managed VPCs 16
 - 5.3 On-Premises Deployment 16
- 6 Post Deployment Steps 18**
 - 6.1 Application Testing 18
 - 6.2 Documentation 18
 - 6.3 Maintenance Training 19
 - 6.4 Remote Access 19
- 7 Support & Training 20**
 - 7.1 Support 20
 - 7.2 Training 20
 - 7.3 Critical Fixes 20
 - 7.4 Updates and Upgrades 21
- 8 License 22**
 - 8.1 Definitions 22

8.2	Subscription License Grant	22
8.3	License Term	22
8.4	License Fee	23
8.5	Ownership	23
8.6	Restrictions	23
8.7	General Conditions	24
8.8	Confidentiality	24
8.9	Warranty Disclaimers	25
9	Pricing	26
9.1	BigML Lite	26
9.2	BigML Enterprise	26
9.2.1	Single Instance - Max 1 Server / 8 Cores	27
9.2.2	Multiple Instance - Max 12 Servers / 96 Cores	27
9.2.3	Multiple Instance - Max 36 Servers / 288 Cores	27
9.2.4	Unlimited Server License	28
9.2.5	Unlimited Deployment License	28
9.2.6	Custom Deployment	29
9.2.7	Estimated AWS Cost	29
	List of Figures	32
	List of Tables	33

What is BigML?

BigML is consumable, programmable, and scalable Machine Learning platform that makes it easy to solve Classification, Regression, Cluster Analysis, Anomaly Detection, Association Discovery and Topic Modeling problems using several patented and patent-pending technologies.

BigML helps you address these problems “end-to-end”. That is, you can seamlessly transform data into actionable models that can be used as remote services or locally embedded into applications to make predictions.

BigML can be used either:

- Interactively through an easy-to-use web interface with an intuitive workflow and several patent-pending visualizations.
- Programmatically using a REST API with open-source bindings for a multitude of programming languages such as Python, Java, Node.js, and even an open-source command line tool called BigMLer.

BigML provides its software as a turnkey, multi-tenant Machine Learning service that is fully managed, auto-scaled, maintained, updated, upgraded, backed-up, and kept running smoothly and securely.

BigML combines the power of comprehensive Machine Learning with the benefits of a cloud-based infrastructure, such as Amazon Web Services (“AWS”), enabling businesses to easily build a multitude of predictive applications cost-effectively with high scalability, flexibility, and reliability.

In BigML’s Software-as-a-Service offering at <https://bigml.com>, everything listed above, including future improvements, is bundled into either:

- A simple monthly subscription, with discounts if the service is contracted using quarterly or annual plans.
- Pay-as-you-go pricing, with no minimum fees and no term-based contracts required.

For companies with stringent data security or privacy requirements, BigML offers Private Deployments that can run in a Virtual Private Cloud (VPC) or on-premises with dedicated servers to meet enterprise-grade requirements.

BigML Private Deployments

As a complement to BigML's multi-tenant Software-as-a-Service (SaaS) delivery model, BigML also offers private, single-tenant deployments. BigML Private Deployments include all the functionality of the BigML multi-tenant version and can be implemented in three ways: Managed VPC, Self-Managed VPC, and On-Premises.

2.1 Choosing BigML Lite vs. BigML Enterprise

BigML provides two options for Private Deployments: BigML Lite and BigML Enterprise. BigML Lite is a fast track to value option for companies ready to implement their first production use cases in a robust and efficient way. BigML Enterprise offers full-scale access for companies ready to adopt Machine Learning across departments with unlimited users and organizations. Before explaining the main differences between these two options, it is important to note that all Private Deployments of BigML (Lite or Enterprise) are available in any public cloud, private cloud, or on-premises, with fully managed or self-managed options. These deployment types are explained further in [Section 2.2](#). Regardless of the type chosen, all BigML Private Deployments include the following:

- Unlimited tasks.
- Regular updates and upgrades of new features and algorithms.
- Priority access to customized assistance.
- Easy upgrades to bigger deployments.

The following [Table 2.1](#) summarizes the differences between BigML Lite and BigML Enterprise.

Feature	BigML Lite	BigML Enterprise
Number of users	5 users	Unlimited users
Organizations access	1 organization	Unlimited organizations
Customization	BigML Standard theme	Personalized theme and logo Prioritized feature request Auto-scaling
Support	Standard 8x5 via email & chat (48-hour max. response time)	Customized direct email & chat (24-hour max. response time)
Configuration	1 server (8 cores)	1 server (8 cores) - unlimited serves/cores
Price	\$10,000/yr or \$1,000/mo	Starting from \$45,000/year + \$10,000 setup fee

Table 2.1: Private Deployment Types Feature Comparison

2.2 Choosing a Private Deployment Type

Managed VPC

For companies that want a secure private deployment but are otherwise comfortable with the SaaS delivery model which eliminates infrastructure and maintenance tasks, and can instead focus on using BigML's Machine Learning service.

Self-Managed VPC

For companies that already operate their own cloud-based infrastructure and feel comfortable managing a new cloud-based service inside their own private cloud.

On-Premises Deployment

For companies that wish to manage and maintain a BigML installation behind their corporate firewall. BigML will be installed on your servers and your IT Staff will have full system control.

A subscription to a BigML Private Deployment includes all the regular improvements and bug fixes released by BigML for the backend, API, front end, and command line layers. Improvements and fixes are regularly made available, except for critical fixes which are provided immediately. All the new features, including new modeling algorithms provided by BigML on its multi-tenant version are implemented immediately for managed VPCs and made available quarterly for self-managed VPCs and on-premises deployments.

A BigML Private Deployment includes the full API making it easy to develop custom applications or integrate BigML directly with any in-house or third-party systems. HTTP requests to BigML's API are SSL-encrypted, so your data is always secure.

Any of the three Private Deployment options can be complemented with technical support plans that offer different levels of assistance and training. Please see [Chapter 7](#) on page 20 for more details.

When choosing one of the three deployment options, businesses should consider not only the total cost of the solution, but also the customization needs, scalability, redundancy, and integration requirements. The following [Table 2.2](#) summarizes the differences between the BigML Private Deployment options.

Feature	Managed VPC	Self-Managed VPC	On-Premises Deployment
Pricing	Yearly Subscription	Yearly Subscription	Yearly Subscription
Hosting	Securely hosted by BigML	On your cloud	On your servers
BigML Updates	Automatically included	Every month	Every 3 months
BigML Upgrades	Automatically included	Every month	Every 3 months
BigML Critical Fixes	Immediately applied	Immediately provided	Immediately provided
Requirements	No customer specific technology requirements beyond a modern browser and high bandwidth Internet	An account with a BigML supported cloud provider	Hardware and software that is physically installed on your own company's computer systems
Access Restriction	IP List	IP List or VPN	IP List
Customization and configuration	1-5 business days	5-10 business days	7-15 business days
Maintenance Fee	In accordance with the number of instances	In accordance with the number of instances	In accordance with the number of instances
Support	From 24x7 to next business day in accordance with your specific support agreement	From 24x7 to next business day in accordance with your specific support agreement	From 24x7 to next business day in accordance with your specific support agreement
Cloud Operation Fee	Monthly billed	N/A	N/A

Table 2.2: Private Deployment Feature Comparison

2.2.1 Managed VPCs




Figure 2.1: Managed VPC

BigML's Managed VPCs provide transparent access to resizable computing capacity in the cloud and virtually unlimited cloud-based data storage without needing to install or configure any hardware or software.

With Managed VPCs, BigML will use the best methods to safeguard your data at the application, access, and data center layers. We not only rely on AWS' extraordinary security, but we take extra steps to ensure that your data is kept private and safe including:

- BigML ensures that access to VPCs is restricted and constantly monitors them to lock them down in case of threat or attack.
- Only a small group of authorized personnel in BigML have access to VPCs and they will not look at your data under any circumstances.
- BigML implements authentication, access control processes, and audit controls that minimize the risk of external compromise.
- BigML uses exclusive S3 buckets, IAM users, and SSH keys for each customer.
- **Note:** In case of disaster, it will be possible to rebuild the VPC including the data and application, in another location within a reasonable amount of time.

BigML will directly pay for your AWS costs and bill them monthly to your business.

PROS

- No initial investment required in infrastructure.
- For multi-instance deployments, catastrophic events will have almost no effect on the service operation. This is because your BigML VPC will be hosted using redundancy and other special safeguards within multiple AWS availability zones.
- Cluster topology is elastic and dynamically adapted, making this solution, in many cases, the most cost-effective.

CONS

- Integration with other on-premises applications can be more difficult.

2.2.2 Self-Managed VPCs

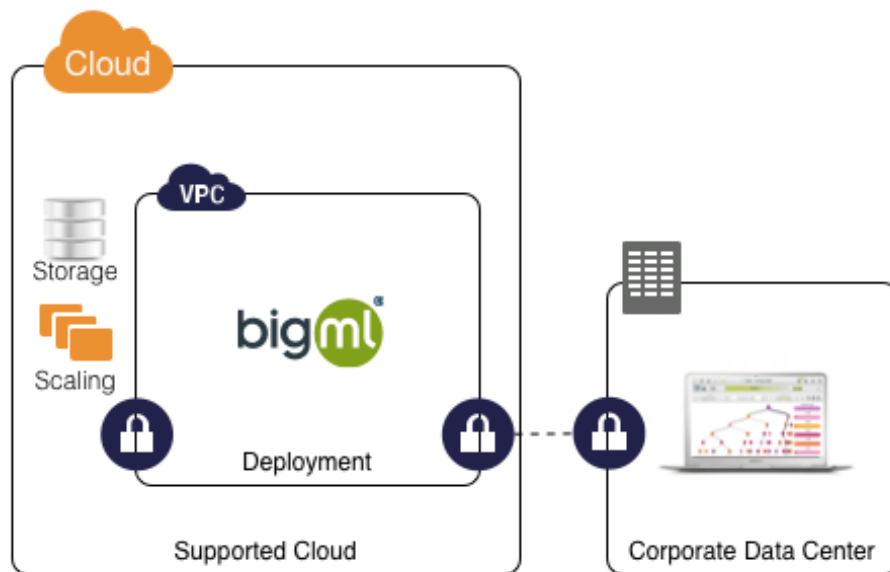


Figure 2.2: Self-Managed VPC

Self-Managed VPCs are a way to leverage your pre-existing cloud infrastructure, policies, and investment. In this case, your company pays for the usage of BigML within your own cloud, managed according to your own requirements. BigML will assist you during the VPC configuration and deployment steps.

PROS

- Everything resides under control of your own IT Staff.
- No separate monthly billing needed for cloud services.

CONS

- Extra services are required to monitor, operate, and maintain.
- Extra steps for initial installation and configuration are required.
- Extra time for regular updates and upgrades needs to be allocated.

2.2.3 On-Premises

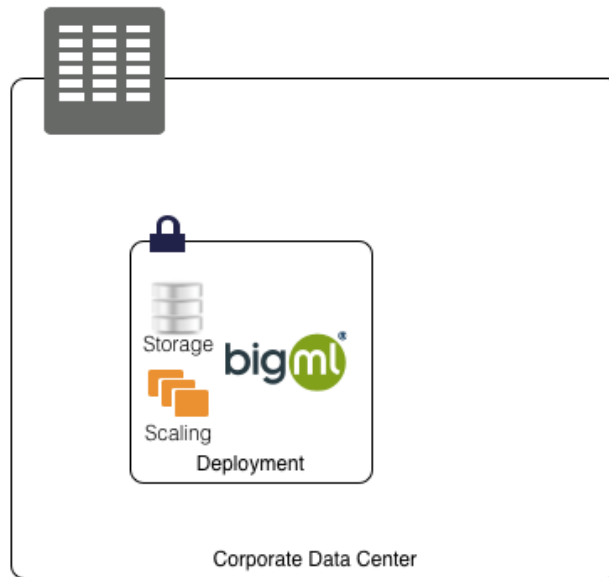


Figure 2.3: On-Premises

BigML offers an On-Premises deployment for companies wishing to configure, manage and maintain their BigML installation behind their corporate firewall.

BigML's On-Premises deployment includes all the functionality of BigML VPC deployments, but this approach also gives your IT Staff control and validation over third-party products that many regulated businesses require. Consider BigML's On-Premises deployment if:

- Your users will have no or limited Internet access.
- Your business requires that all company data resides behind your own corporate firewall.
- On-Premises implementation is a requirement to meet compliance guidelines governing company data management.
- You want the flexibility to customize and control the system or integrate it with other on-premises enterprise applications.
- The economy of scale of acquiring and maintaining the infrastructure remains a more attractive proposition.

PROS

- Better control of systems and data
- Easier integration with existing software
- Lower ongoing monthly costs

CONS

- Higher up-front capital costs for hardware
- Extra steps for initial installation and configuration
- Higher long-term maintenance costs since system must be maintained either by your staff or an outsourced vendor
- Harder to manage peak demand without over-provisioning
- Extra backup systems, redundancy, and failover solutions might add to the initial cost and maintenance

2.3 Choosing a Cloud Provider

BigML is an Amazon Technology Partner, and our default cloud environment which we use for Managed VPCs is AWS.

Self-Managed VPCs can be implemented on any of the cloud providers currently supported by BigML which are: **Amazon AWS**, **Google Cloud**, and **Microsoft Azure**.

While we are working on expanding the list of supported cloud providers, it should be noted that BigML can currently be deployed on *any cloud provider* by treating it as an On-Premises deployment. Of course this will remove some of the benefit of running in the cloud, but can be a good option when those benefits are not important, for example when running BigML in a single instance, or when otherwise constrained to a cloud provider that BigML does not currently support.

There are three main advantages of choosing a BigML supported cloud provider for Self-Managed VPCs:

Storage

In a supported cloud provider, BigML stores source, datasets, model, and other potentially large objects in the default cloud storage option. This means that problems regarding fault tolerance, scalable storage space and performance are shifted to the cloud provider. In an unsupported cloud provider or On-Premises deployment storage must be managed, ensuring that adequate fault-tolerant disk is available, etc.

Flexible Topology

In a deployment with multiple instance, BigML manages the allocation of machines to ensure that there is always adequate resources to process the compute load. In a supported cloud provider, BigML has a lot more flexibility in the types of topologies that can be run, whereas in an unsupported cloud or On-Premises deployment the number of machines will be fixed and BigML can only adjust the machine roles to adapt to changing load.

Cost Savings

In a supported cloud environment, BigML can deploy instances that are the right size for the task at hand, and can power down the compute layer when no jobs are running, which reduces costs. In an un-supported cloud environment or On-Premises deployment, BigML will assume that all machines are always running. For the BigML Managed VPC running in AWS, there is the possibility to leverage even more advanced cost savings by deloying spot instances when the market price falls below a threshold.

These advantages of choosing a BigML supported cloud provider are summarized in [Table 2.3](#).

Feature	Supported Cloud	Un-Supported Cloud / On-Premises
Storage	Cloud Managed: S3, GCS, Blobs	Locally Managed
Topology	Flexible	Fixed
Cost	Matches Demand	Exceeds Demand

Table 2.3: Supported Cloud Provider Advantages

BigML Architecture

BigML has been designed for scalability, flexibility, and reliability. BigML has a distributed logical architecture that preserves the distributed structure at the physical level. This allows BigML Private Deployments to run in a single server using one or multiple services for every logical component, or in multiple instances using one or multiple servers to replicate each logical component.

BigML's architecture follows a multi-layered approach where the components are grouped in four main layers:

User Interface

The web-based front end provides users with an intuitive dashboard. Each user can create, inspect or interact with resources such as sources, datasets, models, clusters, ensembles, anomaly detectors, evaluations, association discovery, or predictions in just a few clicks.

REST API

The middle end (REST API) can programmatically and asynchronously create, read, update, or delete resources using REST.

Machine Learning Back End

A Distributed Machine Learning back end provides a clean API to scalable Machine Learning algorithms which have been designed to handle real world data.

Infrastructure Management

A system that automatically adapts a BigML deployment for reliability and optimum performance.

Each layer uses its own, possibly physically distributed storage, with a fully independent redundancy scheme.

3.1 Web-Based Front End




Figure 3.1: BigML Front-End Layer

BigML's web-based front end is composed of:

- A reverse proxy for HTTPS

- A pre-fork worker model speedy HTTP server
- A high-performance Web Application
- A visualization library and static content that is served to Web browsers

BigML's web-based front end uses the exact same REST API that is available to access BigML programmatically. This is an important feature because the web-based front end can be used to quickly and easily prototype a Machine Learning application, and once complete, it can be fully automated using the REST API.

3.2 Middle End (REST API)

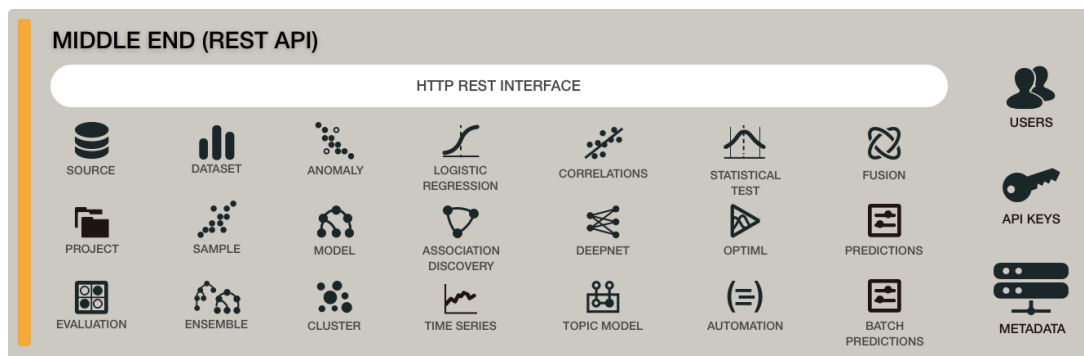


Figure 3.2: BigML Middle-End Layer

BigML's middle end offers a REST API that is used by the front end and can be used independently to programmatically utilize BigML. It is characterized by:

- A reverse proxy, HTTP server, and Web application similar to the one used by the front end
- A REST resource for each high-level Machine Learning component that BigML exposes
- Each REST resource can be created, read, updated, or deleted using basic HTTP requests
- Resources are always returned in JSON format

User authentication and authorization takes place at this level, simplifying the Machine Learning back end. The front end has its own authentication system that is used to store the keys that grant access to the middle end.

3.3 Distributed Machine Learning Back End




Figure 3.3: BigML Back-End Layer

The Distributed Machine Learning back end contains a collection of patent-pending services and libraries. It is organized as a set of services accessible to internal clients via an HTTP interface, using a JSON-based, REST-like API.

Currently, the back end offers seven service access points:

Source

Handles registration of data sources (“raw data”). It also knows how to peek at their contents to infer the number and types of fields they contain.

Dataset

Handles creation and transformation of datasets. A dataset is a structured version of the data source where a series of statistical summaries have been computed for each field.

Model

Responsible for creating models. Models can be single decision trees, ensembles, clusters, anomaly detectors, associations, and logistic regression, etc.

Predict

Performs predictions, batch predictions, computes centroids, anomaly scores, batch centroids, and batch anomaly scores, etc.

Evaluation

Performs and registers evaluations.

Sample

Performs in-memory samples of datasets for visualizations.

WhizzML

Compiles and executes libraries and scripts to create workflows in BigML.

Each access point is reachable at a specific port, multiplexed to one or more “proxy” services that attend and create job requests that are then queued as asynchronous messages using a broker. There is at least one queue per access point, which is read by specialized “server” processes which perform the actual computation. In other words, proxies act as publishers and servers act as subscribers, mediated by a message broker.

This decoupled architecture enables a scalable and reliable service, robust failover and job recovery mechanisms, and the maximum flexibility in terms of scalability policies. For example, new servers are connected to the system on-demand according to the queue size or the number of waiting requests. Also, incoming jobs can be segregated by expected resource needs to allow the creation of sub-clusters with the appropriate CPU and RAM allocations.

The separation between proxies and servers also permits the latter to be liberated of any storage duties, which are centralized at the proxy sub-layer, making resource construction and recovery even more efficient and robust.

3.4 Storage Types

Each BigML layer uses independent object stores according to the type of data and the type of deployment, as shown in [Table 3.1](#). This independence allows the storage systems to evolve according to the diverging needs of each layer. If required, BigML can work with other types of data stores, but this will involve additional time for customization and integration.

Layer	Managed VPC	Self-Managed VPC	On-Premises
Web-Based Front End	MongoDB	MongoDB	MongoDB
REST API	MongoDB	MongoDB	MongoDB
ML Back End	MongoDB	MongoDB	MongoDB
Object Storage	S3	S3 / GCS / Blobs	GridFS

Table 3.1: Storage Types

3.5 Authentication

BigML's On-Premises deployment can be integrated with external LDAP or CAS services for authentication. By default, BigML runs in a self-contained authentication mode with user information stored in the front-end layer and API Keys stored in the middle-end layer. LDAP or CAS is not required.

Deployment Planning

4.1 Deployment Sizing

BigML Private Deployments are licensed in four basic sizes:

Single Instance

The entire BigML architecture including all four layers is deployed on a single machine.

Multiple Instances - Max 12 Servers / 96 Cores

The back end computation layer is distributed among up to 12 servers, with a maximum total number of 96 cores.

Multiple Instances - Max 36 Servers / 288 Cores

The back end computation layer is distributed among up to 36 servers, with a maximum total number of 288 cores.

Custom

For deployments requiring more than the max 36 servers.

In order to determine the required deployment size, resources, storage, etc., it is necessary to consider the following questions about the data processing and modeling needs:

Volume of Data

What amount of data will be processed on a regular basis? Does the data vary frequently or just from time to time?

Data Complexity

What is the complexity of the data? How many columns? How many text fields?

Number of Concurrent Tasks

How many resources (datasets, models, evaluations, etc.) need to be created at the same time?

Modelling Speed

How frequently do new resources need to be created? Every hour or every day? Once a month?

Complexity of Modelling Tasks

Will large ensembles be required? Or automated feature selection tasks?

Number of Concurrent Users

What is the maximum number of users expected to be using the system concurrently? Will all of them be creating new models, or will some just be browsing or inspecting models created by others?

Level of Automation

Will machine learning tasks be fully automated? Or will BigML be used primarily through its web interface?

Prediction Speed

Are real-time predictions needed, or will offline batch predictions be sufficient?

Once the data processing and modeling needs are defined, BigML's Support Team can help determine if the single instance deployment or a distributed, multi-instance deployment is required.

Note: While it is important to choose a deployment size sufficient to handle the expected workload, it is not hard to change between sizes after the deployment is installed. For more information about deployment sizing including pricing, see [Chapter 9](#).

4.2 Deployment Options

There are several options to decide on before the BigML deployment can be configured and implemented:

Base Number of Instances

Based on the questions in [Section 4.1](#), choose a single or multi-instance deployment.

Secured Access

Choose secured connections using restricted source IP access, or using a gateway instance that constructs an actual VPN between a BigML VPC and your internal network.

Look & Feel

Choose customized colors and company logo for the BigML web interface.

Authentication and User Management

Choose how BigML authenticates users within your organization.

Gallery and Resource Sharing

Choose to enable BigML's gallery and resource sharing for internal collaboration.

BigML's Support Team will create a functional preview of the system that you can try before it is deployed to ensure everything you need is included.

4.3 Pre-Deployment Steps

BigML's Private Deployment setup can usually be completed in 1-2 business days, but this requires some preparation as detailed below. To guarantee a successful deployment, BigML offers assistance for both Self-Managed VPCs and On-Premises Deployments.

Note this list is an overview of the information that will be required for each deployment type. Once a specific deployment option has been chosen, BigML will provide a Pre-Deployment Checklist suitable for a technical audience detailing all of the information needed.

4.3.1 Managed VPCs

For Managed VPCs:

1. Provide the source IPs that should have access to your service, or accept non-firewalled.
2. Provide the domain names authorized to receive emails.
3. Prepare your systems to accept emails from @vpc.bigml.com.
4. Add records to your DNS, if the deployment will use a custom domain.
5. Provide BigML with SSL certificates if using a custom domain.

4.3.2 Self-Managed VPCs

Before BigML can start the deployment process, you will need to:

- Create an appropriate storage space and access policy if using a supported cloud storage (S3/Blobs/GCS)
- Deploy a supported instance to host “Sauron”, the BigML deployment management tool. If the deployment type is a single instance, this can be the same as the instance that is running the BigML application.
- Create a user in the cloud environment that has permissions to stop/start instances and can change the instance metadata.
- Set up a VPC subnet or two depending on your specific configuration.
- Configure DNS records for BigML’s Front End, API and static content.
- Provide BigML with SSL certificates for the domain names.

As mentioned previously, it is possible to setup a Self-Managed VPC as if it was an On-Premises Deployment. This allows for the possibility of running BigML in an unsupported cloud provider, with the limitations as described in [Section 2.3](#) on page 8.

Note: It is important to think of the machines assigned to BigML in a Self-Managed VPC as “appliances”. The BigML management software, “Sauron”, expects to have full control of the machines and full access to all system resources. Attempting to share the machines in a BigML deployment with other services will cause problems and is not supported.

Note: If your environment requires a “hardened” server configuration for maximum security, this should be implemented after BigML’s deployment is complete.

4.3.3 On-Premises Deployment

For an On-Premises Deployment, you must allocate an adequate number servers within your infrastructure. BigML’s Support Team will work together with your IT Staff on a very specific definition of the system requirements to deploy BigML.

Note: It is important to think of the machines assigned to BigML in an On-Premises Deployment as “appliances”. The BigML management software, “Sauron”, expects to have full control of the machines and full access to all system resources. Attempting to share the machines in a BigML deployment with other services will cause problems and is not supported.

Note: If your environment requires a “hardened” server configuration for maximum security, this should be implemented after BigML’s deployment is complete.

Deployment Steps

The steps shown here are only an overview. Once you have chosen a deployment option, BigML will provide you with a document which details the exact deployment steps.

5.1 Managed VPCs

BigML's Support Team will provide a Pre-deployment checklist which covers the details outlined in [Subsection 4.3.1](#). Once the technical requirements have been established, BigML will take care of all steps required to deploy and maintain the application.

5.2 Self-Managed VPCs

The deployment steps for a self-managed VPCs are as follows:

1. BigML will provide a pre-deployment checklist which covers the details outlined in [Subsection 4.3.2](#)
2. Once the pre-deployment checklist has been completed, BigML will deploy a functional preview of the customized application
3. BigML will create a package distributed via an AWS S3 bucket that contains the software to be installed
4. The first machine in the deployment, which will run the BigML management software "Sauron", must be brought up manually and installed
5. If the installation is a multi-node deployment, Sauron can deploy the remaining hosts after configuration
6. BigML will provide a manual for maintenance of the deployment

BigML will maintain an internal repository with your specific configuration along with manifest files that record how Sauron should take an instance from one version to another. When BigML releases a new "version", we will give you access to new files and notify you. Then you can instruct Sauron to automatically upgrade for you.

5.3 On-Premises Deployment

The deployment steps for an On-Premises Deployment are the same as for the Self-Managed VPC ([Section 5.2](#)) except:

- If the On-Premises Deployment is offline, then the software package will be physically delivered

- The nodes in the deployment must be individually configured to allow ssh access from the Sauron node with root access

Post Deployment Steps

6.1 Application Testing

After the deployment steps have been completed, BigML's Support Team will:

1. Test the deployment to ensure that the application is ready to:
 - Create new data sources using local or distributed data.
 - Create new resources such as datasets, models, clusters etc., that are similar in size and complexity to your expected usage.
 - Create new predictions.
2. Review security settings to ensure that they are aligned with your security policy and make any needed adjustments.
3. Assist with the setup of user accounts.

6.2 Documentation

If your deployment is a Self-Managed VPC or On-Premise Deployment, BigML's Support Team will provide two documents:

1. An "Installation" document describing:
 - The deployment process in enough detail to be repeatable
 - Available configuration settings
 - A detailed description of the delivered components
 - How to test functionality of the BigML application
 - Uninstallation instructions
2. A "Maintenance" document describing how to:
 - Manage user accounts
 - Access support
 - Upgrade the BigML application
 - Monitor functionality
 - Backup and restore data

6.3 Maintenance Training

The MLaaS BigML.com platform is a massively distributed system comprised of more than 30 distinct components that work together in a horizontally and vertically scalable way. BigML's Private Deployment is not a monolithic piece of software; rather it is a part of this same distributed system to work in a local environment. This means that much like the MLaaS version, BigML's on-site deployment can be easily deployed to an on-site cluster and scaled to match customer requirements.

While this is very powerful, it means that the system is somewhat complex and training to manage the deployment is recommended. Training will ensure that your IT staff:

- know how to create and manage user accounts.
- know which parts of the deployment must be monitored and the regular monitoring and maintenance tasks that should be under their control.
- understand the location of log entries and the way to use them for monitoring.
- can follow the upgrade process.

6.4 Remote Access

In addition to being available to answer questions and to schedule support training, BigML's Support Team is happy to directly help upgrade, maintain, or troubleshoot both Self-Managed and On-Premises BigML Private Deployments. To facilitate this assistance, it is necessary to have a method of remote access for BigML's Support Team to manage the private deployment.

Remote access can be provided thru whatever VPN or remote desktop control software is standard for your company, and is always at your discretion.

Support & Training

7.1 Support

After your Private Deployment has been implemented, BigML's Support Team is always available to quickly respond to queries about product features and functionality. BigML's Support Team works collaboratively with other functional areas within BigML to answer customer questions in a timely fashion.

BigML's support offerings include:

1. A range of support options that give you the flexibility to select the right amount of service for any size organization.
2. A dedicated email address (e.g., `yourcompany-support@bigml.com`) with guaranteed response time.
3. Telephone and/or chat-based support that provides rapid response and up-to-the-minute information.
4. Detailed user manuals that explain the use and configuration of each feature in the BigML application.

In addition, BigML also offers special assistance and training plans that can help you quickly develop predictive applications on top of your BigML Private Deployment.

7.2 Training

BigML offers both on-line training and in-person training to ensure that your team gets the most from the BigML application. Free easy-to-use education materials are provided as well.

7.3 Critical Fixes

Minor fixes will be packaged together into future general updates. Critical fixes or security improvements will be provided as soon as they are available. How the fixes are delivered depends on your deployment type:

Managed VPC

Critical fixes may be applied without prior notification if the expected downtime is minimal, for example only a few seconds. BigML will always notify you of any changes that are made.

Self-Managed VPC & On-Premises

BigML's Support Team will prepare the deployment package and contact you to apply the update.

7.4 Updates and Upgrades

The BigML platform is constantly evolving with minor improvements and major improvements such as new algorithms or visualizations. These improvements will be regularly rolled up into an upgrade package that can be applied to your Private Deployment:

Managed VPC

BigML's Support Team will apply upgrades during periods when your deployment is not being used. If you prefer, BigML can contact you to arrange a convenient downtime for each deployment.

Self-Managed VPC & On-Premises

BigML's Support Team will prepare the deployment package and contact you to apply the update.

After each update, you will receive a document describing the changes included in the deployment, detailing the *New*, *Improved*, and *Fixed* items.

License

Below is BigML's standard license for BigML Private Deployments.

8.1 Definitions

“BigML Products” shall include all forms of the BigML software programs, along with documentation, upgrades, and updates, which may be provided by BigML.

“Proprietary Rights” include any and all rights, whether registered or unregistered, in and with respect to patents, copyrights, confidential information, know-how, trade secrets, moral rights, contract or licensing rights, confidential and proprietary information protected under contract or otherwise under law, trade names, domain names, trade dress, logos, animated characters, trademarks, service marks, and other similar rights or interests in intellectual or industrial property.

“Server” means a physical or virtual machine capable of running BigML Products. For purposes of this definition, each virtual machine, hardware partition, or blade is considered to be a separate physical hardware system, and, therefore, a separate server.

“Core” means the processor or execution core contained in the same integrated circuit within a server's central processing unit, whether such cores are virtual or physical.

8.2 Subscription License Grant

Subject to the terms and conditions of your license agreement, BigML will grant your company (the “licensee”) a limited, revocable, non-exclusive, non-transferable, non-assignable, temporary right under the Proprietary Rights of BigML to use BigML Products.

The license will permit the licensee to install BigML Products on up to a maximum number of Servers simultaneously, or up to the total number of Cores where the BigML Products are installed. BigML's License does not restrict the volume of data analyzed or number of concurrent users for either modeling or predictive tasks provided that the total number of computers where the BigML Products are installed does not exceed the Servers or Cores permitted.

8.3 License Term

The License granted hereunder shall continue for one-year period (“Initial Term”) unless and until terminated pursuant to the restrictions hereof and subject to Licensee's proper performance of its obligations hereunder.

At the end of the Initial Term, the License will automatically renew for successive three (3) year periods (“Renewal Term”) unless Licensee provides BigML sixty (60) days advance written notice prior to the commencement of the Renewal Term.

8.4 License Fee

Licensee shall pay, upon delivery of the BigML Products, the customization and configuration Fee, the License Fee, the maintenance Fee, and the support Fee set out in next section “Pricing”.

The License Fee, maintenance Fee, and support Fee are payable in full at the beginning of the Initial Term or successive Renewal Terms. All payments shall be made within thirty (30) days of the delivery of the BigML Products to the Licensee.

Unless timely provided with a valid certificate of exemption or other evidence that items are not taxable, BigML will invoice Licensee for all applicable taxes including, but not limited to, VAT, GST, sales tax, consumption tax and service tax. Licensee will make all payments free and clear of, and without reduction for, any withholding or other taxes; any such taxes imposed on payments by Licensee hereunder will be Licensee sole responsibility.

BigML’s License can be upgraded at any time and the current Fees will be applied towards the new License Fees.

8.5 Ownership

BigML owns the BigML Products and all intellectual property rights embodied therein, including copyrights and valuable trade secrets embodied in the product’s design and coding methodology. United States copyright laws and international treaty provisions protect BigML Products. BigML will provide Licensee only a limited use license, and no ownership of any intellectual property.

8.6 Restrictions

BigML Products are not designed or licensed for use in hazardous environments requiring fail-safe controls, or other mission critical application where human life or property may be at stake including, without limitation: operation of nuclear facilities, aircraft navigation/ communication systems, air traffic control, and/or life support or weapons systems. The licensee understands that the BigML Products are not designed for such purposes and that its failure in such cases could lead to death, personal injury, or severe property or environmental damage for which BigML is not responsible.

The licensee will neither use nor permit others to use their BigML Products (a) for any unlawful, invasive, infringing, defamatory, fraudulent, or obscene purpose; (b) to create any virus, worm, Trojan horse, or harmful code; (c) for any illegal or unauthorized purpose; (d) to violate any laws in their jurisdiction (including but not limited to patent, copyright, trademark, or other intellectual property laws); or (e) to alter, steal, corrupt, disable, destroy, trespass, or violate any security or encryption of any computer file, database, or network. If the licensee, or a third party through the BigML Products, violates any of the foregoing prohibitions, BigML may immediately suspend the license without further liability or obligation to the licensee. Furthermore, the licensee shall be responsible for the cost of any services/actions BigML takes to mitigate any damage caused by such impermissible behavior.

The licensee will make no attempt to, and will not permit them or any other third party to make any attempt to:

1. Alter, modify, improve, reverse engineer, disassemble, or decompile the BigML Products.
2. Interfere in any manner with the hosting of the Applications or BigML Products associated with them.

3. Reproduce, duplicate, or copy any portion of the BigML Products, use of the BigML Products, or access to the BigML Products without the express permission by BigML.

8.7 General Conditions

BigML will use reasonable efforts to ensure that the BigML Products are available twenty-four hours a day, seven days a week. However, there will be occasions when BigML Deployments will be interrupted for maintenance, upgrades, emergency repairs, failure of telecommunications systems, equipment failures, acts of god, or other situations. BigML will make reasonable efforts to minimize such disruption where it is within BigML's reasonable control. Licensee agrees that BigML will not be liable in any event for any suspension, modification, discontinuance, lack of availability, or other termination of the BigML Products.

Licensee understands and agrees that the BigML Products are provided on an "AS IS" and "AS AVAILABLE" basis. Licensee agrees that access to and use of the BigML Products is at their own risk. Licensee acknowledges and accepts that there may be technical downtime with BigML Deployments. Such downtime or errors may prevent their datasets, models or predictions being created, viewed or accessible. BigML shall not be liable if the creation or access to any dataset, model or prediction is delayed, prevented from access or displayed incorrectly by any cause beyond BigML's control including server downtime.

BigML will not be responsible or liable for any harm to their computer system, loss of data, or other harm of any kind to Licensee or any other third party that results from their access to or use of the BigML Products, or any data. Licensee agrees that BigML has no responsibility or liability to them or any other third party for the deletion of, or the failure to store or to transmit any data and/or any other communications maintained by the BigML Products, or any partial or inaccurate transmission of data.

Licensee agrees to, upon request from BigML and within thirty (30) days of such request, fully document and certify the compliance of the License regarding the usage and number of copies and installations of the BigML Products.

Upon the expiration or termination of this License, all rights granted to the Licensee under this agreement shall forthwith terminate and immediately revert to BigML and Licensee shall discontinue all use of the BigML Products and the like. Licensee agrees to return to BigML or to destroy all copies of the BigML Products upon termination of the License.

This License is the entire and exclusive agreement between BigML and Licensee regarding the BigML Products. This License replaces and supersedes all prior negotiations, dealings, and agreements between BigML and Licensee regarding the BigML Products.

This License shall be governed by the law of the state of Oregon.

8.8 Confidentiality

Licensee acknowledges that the BigML Products contain trade secrets and other valuable and confidential information of BigML. Licensee shall not act, or fail to act, in any way or manner to intentionally or negligently harm BigML or BigML's rights in its intellectual property in this License. This License, together with any information learned in connection therewith that should reasonably be considered confidential under the circumstances, are "Confidential Information". Licensee shall only disclose Confidential Information on a need-to-know basis to Licensee's employees. Licensee may not disclose any Confidential Information to a third party and shall use all reasonable care to keep the Confidential Information confidential and consistent with the grant of rights to Licensee under this agreement.

Exceptions. The confidential obligation shall not apply to any information or materials which: (a) was in Licensee's possession before being received from BigML, (b) are or become publicly available through no fault of Licensee; (c) are independently developed without reliance on Confidential Information; (d) are received from a third party with no duty of confidentiality.

8.9 Warranty Disclaimers

BIGML MAKES NO WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, AND BIGML DISCLAIMS THE IMPLIED WARRANTIES OF TITLE, NON-INFRINGEMENT, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, SYSTEM INTEGRATION, AND DATA ACCURACY. LICENSEE ACKNOWLEDGES THAT NO REPRESENTATIONS OTHER THAN THOSE CONTAINED IN THIS AGREEMENT HAVE BEEN MADE RESPECTING THE LICENSED PRODUCTS OR SERVICES TO BE PROVIDED HEREUNDER, AND THAT LICENSEE HAS NOT RELIED ON ANY REPRESENTATION NOT EXPRESSLY SET OUT IN THIS AGREEMENT. FURTHER, FOR THE SOFTWARE WHICH THIS MAY APPLY, LICENSEE ACKNOWLEDGES AND AGREES THAT NEITHER THE INTERNET OR THE CLOUD PROVIDER ARE NOT ESTABLISHED OR MAINTAINED BY BIGML, THAT BIGML HAS NO CONTROL OVER THE INTERNET OR CLOUD PROVIDER, AND THAT BIGML IS NOT LIABLE FOR THE DISCONTINUANCE OF OPERATION OF ANY PORTION OF THE INTERNET OR CLOUD PROVIDER OR POSSIBLE REGULATION OF THE INTERNET WHICH MIGHT RESTRICT OR PROHIBIT THE OPERATION OF THE LICENSED PRODUCTS.

Pricing

9.1 BigML Lite

BigML Lite provides the power of BigML at a fixed, accessible price point. As outlined in [Table 2.1](#), BigML Lite includes access for 5 users within 1 organization, with the standard BigML theme and customer support via email and chat. This deployment option gives access to 1 server (8 cores), which is a single instance running the entire BigML application stack: front end, API middle end, Machine Learning back end, and storage. Your company can choose to pay for BigML Lite at either a monthly rate of \$1,000 or a yearly rate of \$10,000 (discounted).

9.2 BigML Enterprise

BigML's standard pricing for Private Deployments is an up-front, yearly subscription based on the maximum number of back end **servers** (Dataset Server, Model Servers, Evaluation Servers, etc.) and the maximum number of **cores** needed by your deployment. BigML's license does not otherwise impose limits on the number of integrated applications, the number of users, or the volume of data processed by your deployment. The number of **servers** or **cores** needed for BigML's front end, middle end and storage do *not* count against the maximum number of **servers** or **cores** allowed by your license, *except* for the single instance deployment which must always be deployed on one **server** with no more than the maximum number of **cores**. BigML's license includes regular updates and upgrades.

BigML's standard pricing for Private Deployments is independent of the type of deployment (Managed, Self-Managed or On-Premises) that you choose for your company. That is, the license cost is exactly the same independent of whether BigML runs on-premises, in your cloud, or fully-managed by BigML.

BigML also offers special pricing for 3-year, 5-year, 7-year contracts or global licenses.

Kindly note these additional considerations:

- BigML does not rigorously enforce an activation key to deploy BigML back end servers. BigML trusts that your company will respect the licensing agreement and will not surpass the licensed number of **servers** or **cores** in your deployment.
- Maintenance and support fees are based on the maximum number of back end servers allowed by your license.
- On-Premises deployments require an additional set-up fee that will vary depending on your location, level of integration, and complexity of your IT infrastructure.
- For Managed VPCs the corresponding cloud costs will be billed back to you monthly.

9.2.1 Single Instance - Max 1 Server / 8 Cores

This is a single instance running the entire BigML application stack: front end, API middle end, machine learning back end, and storage, all in one server.

Single Instance	Amount
<i>Customization and Configuration</i>	\$10,000
<i>License</i>	\$30,000/year
<i>Maintenance Fee</i>	\$6,000/year
<i>Support Fee</i>	\$9,000/year

Table 9.1: BigML Private Deployment - Single Instance

For a Managed VPC, the BigML single instance is deployed in the default subnet of an AWS VPC. Services will be available only to localhost except of course for the HTTPS endpoints for the web-based front end and API middle end access. These endpoints will be accessible only from the list of IPs that you provide and by BigML. BigML does not offer a VPN connection to single instance VPCs.

See [Table 9.6](#) for the estimated cloud costs of a 100% utilized Single Instance deployment on AWS.

9.2.2 Multiple Instance - Max 12 Servers / 96 Cores

A 12-instance BigML VPC deployment removes the limitation of running the BigML application on a single instances and distributes the computational workload across as many as twelve separate back end servers. This has several advantages over a single instance deployment:

- Highly available, redundant instances
- Option to connect with VPN
- Ability to handle more concurrent users
- Ability to handle more concurrent and larger tasks
- Ability to auto-scale, within limits
- Server/core limitation applies only to back end servers

Multiple Instances (Max 12 Servers / 96 Cores)	Amount
BigML Private Deployment - Customization and Configuration	\$20,000
BigML Private Deployment - License	\$70,000/year
BigML Private Deployment - Maintenance Fee	\$14,000/year
BigML Private Deployment - Support Fee	\$21,000/year

Table 9.2: BigML Private Deployment - Max 12 Servers / 96 Cores

See [Table 9.7](#) for estimated costs of a 12-instance (fully-utilized) Multiple Instance deployment on AWS. Other configurations or cloud providers may result in different pricing.

9.2.3 Multiple Instance - Max 36 Servers / 288 Cores

For even more computational capacity, consider the 36-node BigML VPC deployment which has several advantages over the single instance or 12-instance BigML VPCs:

- Additional capacity, bigger storage, and more RAM

- Storage completely separated from other services
- More multi-user ability
- Higher availability as machines can be deployed across availability zones, ensuring that the VPC continues to run even if AWS loses one availability zone
- Server/core limitation applies only to back end servers

Multiple Instances (Max 36 Servers / 288 Cores)	Amount
BigML Private Deployment - Customization and Configuration	\$30,000
BigML Private Deployment - License	\$120,000/year
BigML Private Deployment - Maintenance Fee	\$24,000/year
BigML Private Deployment - Support Fee	\$36,000/year

Table 9.3: BigML Private Deployment - Max 12 Servers / 288 Cores

See [Table 9.8](#) for estimated costs of a 36-node (fully utilized) Multiple Instance deployment on AWS. Other configurations or cloud providers may result in different pricing.

9.2.4 Unlimited Server License

The Unlimited Server License option allows for unlimited nodes and cores running in a single deployment supporting an unlimited number of users. This option removes the need to be concerned with the deployment size and allows scaling the BigML Private Deployment to handle any computational workload.

- Unlimited capacity, storage and RAM
- Storage completely separated from other services
- Scales to any number of users in a single country
- Higher availability as machines can be deployed across availability zones, ensuring that the VPC continues to run even if AWS loses one availability zone

Unlimited Servers / Cores in a Single Site	Amount
BigML Private Deployment - Customization and Configuration	\$50,000
BigML Private Deployment - License	\$500,000/year
BigML Private Deployment - Maintenance Fee	\$100,000/year
BigML Private Deployment - Support Fee	\$150,000/year

Table 9.4: BigML Private Deployment - Unlimited Servers / Cores

9.2.5 Unlimited Deployment License

The Unlimited Deployment License option allows for unlimited nodes and cores running in an unlimited number of deployments supporting users within a single company. This option removes the need to be concerned with the deployment size and allows scaling the BigML Private Deployment to handle any computational workload, including isolated deployments for separate departments.

- Unlimited capacity, storage and RAM
- Storage completely separated from other services
- Scales to any number of users across entire company worldwide

- Higher availability as machines can be deployed across availability zones, ensuring that the VPC continues to run even if AWS loses one availability zone

Unlimited Deployments / Cores	Amount
BigML Private Deployment - Customization and Configuration	\$100,000
BigML Private Deployment - License	\$1,500,000/year
BigML Private Deployment - Maintenance Fee	\$300,000/year
BigML Private Deployment - Support Fee	\$450,000/year

Table 9.5: BigML Private Deployment - Unlimited Deployments / Cores

9.2.6 Custom Deployment

If the data load of your private deployment might require more than 36 servers or 288 cores and you are not quite ready for a [Unlimited Server License](#) or [Unlimited Deployment License](#), please contact us at <mailto:info@bigml.com> to get a personalized quote.

9.2.7 Estimated AWS Cost

Below are the estimated AWS costs for various deployment sizes. A few things to keep in mind about these examples:

- Example pricing assumes 100% utilization, that is all instances running 24x7. However, in the multiple instance deployments the autoscaler can power down the back end servers when they are not needed.
- Instance sizes are only suggestions. Actual required instances will depend on your compute requirements, as laid out in [Section 4.1](#)
- These estimates do not include possible additional costs for Elastic IPs, Data Transfer, Elastic Load Balancing, VPN gateways, or other related components if utilized.
- Better prices for AWS EC2 instances can be achieved with a 1-year or 3-year contract by using EC2 reserved instances.

Single Instance

Component	Num	Specification	Cost/mo	Total/mo
<i>Instance below counts against 1-Instance limit</i>				
EC2 Instance	1	8 CPU, 30GB, 160GB volatile SSD storage (m3.2xlarge)	\$390	\$390
<i>Storage</i>				
EBS RAID Metadata Storage	4	General Purpose SSD, 375 IOPS, 125GB	\$12.50	\$50
EBS Backup Storage	2	Magnetic, 100 IOPS, 1TB	\$64	\$128
S3 Resource Storage	1TB	Sources, Datasets, Models, etc, 1TB, PUTS/GETS++	\$30/TB++	\$32
Estimated AWS Costs for on-demand Single Instance VPC				\$600.00

Table 9.6: Estimated monthly cost of a 100% Utilized Single Instance VPC in the AWS us-east-1 region

Max 12 Instance / 96 Cores

Component	Num	Specification	Cost/mo	Total/mo
<i>Instances below count against 12-Instance limit</i>				
Half-stack Front End Middle End and Proxy Services	2	8 CPU, 30GB, 160GB volatile SSD storage (m3.2xlarge)	\$390	\$780
Datasource Server 2 workers	2	2 CPU, 3.75GB RAM, 410GB volatile MAG storage (m1.medium)	\$64	\$128
Dataset Server 1 worker	2	8 CPU, 15GB RAM, 160GB volatile SSD storage (c3.2xlarge)	\$307	\$614
Model Server 1 worker	2	2 CPU, 17.1GB RAM, 420GB volatile MAG storage (m2.xlarge)	\$179	\$358
Evaluation/Predict Server 1 worker	2	8 CPU, 15GB RAM, 160GB volatile SSD storage (c3.2xlarge)	\$307	\$614
WhizzML Server 1 worker	2	2 CPU, 3.75GB RAM, 410GB volatile MAG storage (m1.medium)	\$64	\$128
<i>Storage</i>				
EBS RAID Metadata Storage	8	General Purpose SSD, 750 IOPS, 250GB	\$25	\$200
EBS Backup Storage	4	Magnetic, 100 IOPS, 1TB	\$64	\$256
S3 Resource Storage	4TB	Sources, Datasets, Models, etc, 4TB	\$30/TB++	\$128
Estimated AWS Costs for on-demand Multiple Instance VPC				\$3206.00

Table 9.7: Estimated monthly cost of a 100% Utilized 12-Instance VPC in the AWS us-east-1 region

Max 36 Instance / 288 Cores

Component	Num	Specification	Cost/mo	Total/mo
Front End Stack	2	2 CPU, 3.75GB RAM, 32GB, volatile SSD storage (c3.large)	\$77	\$154
Middle End Stack	2	4 CPU, 7.5GB RAM, 80GB volatile SSD storage (c3.xlarge)	\$154	\$308
MongoDB	2	4 CPU, 30.5 RAM, 1 x 80 volatile SSD storage (r3.xlarge)	\$243	\$486
<i>Instances below count against 36-Instance limit</i>				
Proxy Services	2	4 CPU, 30.5GB RAM, 80SSD volatile SSD storage (r3.xlarge)	\$243	\$486
Datasource Server 2 workers	2	2 CPU, 3.75GB RAM, 410GB volatile MAG storage (m1.medium)	\$64	\$128
Dataset Server 1 worker	4	8 CPU, 15GB RAM, 160GB volatile SSD storage (c3.2xlarge)	\$307	\$1228
Model Server 1 worker	16	2 CPU, 17.1GB RAM, 420GB volatile MAG storage (m2.xlarge)	\$179	\$2864
Evaluation Server 1 worker	4	8 CPU, 15GB RAM, 160GB volatile SSD storage (c3.2xlarge)	\$307	\$1228
Predict Server 1 worker	2	8 CPU, 15GB RAM, 160GB volatile SSD storage (c3.2xlarge)	\$307	\$614
Predict Proxy	4	2 CPU, 3.75GB RAM (m1.medium)	\$64	\$256
WhizzML Server	2	2 CPU, 3.75GB RAM (m1.medium)	\$64	\$128
<i>Storage</i>				
EBS RAID Metadata Storage	8	General Purpose SSD, 1500 IOPS, 500GB	\$50	\$400
EBS Backup Storage	4	Magnetic, 100 IOPS, 1TB	\$64	\$256
S3 Resource Storage	4TB	Sources, Datasets, Models, etc, 4TB	\$30/TB	\$128
Estimated AWS Costs for on-demand Multiple Instance VPC				\$8664

Table 9.8: Estimated monthly cost of a 100% Utilized 36-Instance VPC in the us-east-1 region

List of Figures

- 2.1 Managed VPC 4
- 2.2 Self-Managed VPC 6
- 2.3 On-Premises 7

- 3.1 BigML Front-End Layer 9
- 3.2 BigML Middle-End Layer 10
- 3.3 BigML Back-End Layer 10

List of Tables

2.1	Private Deployment Types Feature Comparison	3
2.2	Private Deployment Feature Comparison	4
2.3	Supported Cloud Provider Advantages	8
3.1	Storage Types	11
9.1	BigML Private Deployment - Single Instance	27
9.2	BigML Private Deployment - Max 12 Servers / 96 Cores	27
9.3	BigML Private Deployment - Max 12 Servers / 288 Cores	28
9.4	BigML Private Deployment - Unlimited Servers / Cores	28
9.5	BigML Private Deployment - Unlimited Deployments / Cores	29
9.6	Estimated monthly cost of a 100% Utilized Single Instance VPC in the AWS us-east-1 region	29
9.7	Estimated monthly cost of a 100% Utilized 12-Instance VPC in the AWS us-east-1 region	30
9.8	Estimated monthly cost of a 100% Utilized 36-Instance VPC in the us-east-1 region	31

bigml[®]