

Using Cloudmesh GAS for Speedy Generation of Hybrid Multi-Cloud Auto Generated AI Services

Gregor von Laszewski*, Anthony Orlowski*, Richard H. Otten*, Reilly Markowitz*,
Sunny Gandhi*, Adam Chai*, Geoffrey C. Fox* and Wo L. Chang†

* *Dept. of Intelligent Systems Engineering, Indiana University, Bloomington, IN 47408, USA*
Email: laszewski@gmail.com

† *Information Access Division, National Institute of Standards and Technology*
Gaithersburg, MD 20899, USA

Abstract—Today’s problems require a plethora of analytics tasks to be conducted to tackle state-of-the-art computational challenges posed in society impacting many areas including health care, automotive, banking, natural language processing, image detection, and many more data analytics-related tasks. Sharing existing analytics functions allows reuse and reduces overall effort. However, integrating deployment frameworks in the age of cloud computing are often out of reach for domain experts. Simple frameworks are needed that allow even non-experts to deploy and host services in the cloud. To avoid vendor lock-in, we require a generalized composable analytics service framework that allows users to integrate their services and those offered in clouds, not only by one, but by many cloud compute and service providers.

We report on work that we conducted to provide a service integration framework for composing generalized analytics frameworks on multi-cloud providers that we call our Generalized AI Service (GAS) Generator. We demonstrate the framework’s usability by showcasing useful analytics workflows on various cloud providers, including AWS, Azure, and Google, and edge computing IoT devices. The examples are based on Scikit learn so they can be used in educational settings, replicated, and expanded upon. Benchmarks are used to compare the different services and showcase general replicability.

Index Terms—Software/Software Engineering, Tools, Programming Environments/Construction Tools, Code generation, Distributed systems, Artificial Intelligence

Keywords— code generation, AI services, multi-cloud, hybrid cloud, OpenAPI, REST, cloudmesh

I. INTRODUCTION

Today’s application scientists want to share their services with many colleagues while not only offering the services as bare metal programs but exposing the functionality as a Software as a Service (SaaS). This has the advantage that the services can be readily reused by other applications and hosted in the cloud, allowing access to state-of-the-art services or collections of resources that otherwise would not be accessible to individual domain experts. Through the increased availability, resource constraints can be reduced, and scientists can offer their analytics workflows as services to the community. This may include long-lasting services envisioned by cloud computing as part of its Software as a Service (SaaS) paradigm or for smaller analytics functions as microservices.

Furthermore, a subset of analytics functions can be offered as part of a serverless computing model, elevating the penetration from a pure bare metal solution to a multi-pronged cloud-based service offering.

While working with many professionals, researchers, and students, we found that the entry barriers to accomplishing this goal remain very high, and would elude many domain experts as they have neither the expertise nor the time to learn the foundation to conduct the infrastructure-related tasks integrating DevOps and analytics. Although recent developments, especially on the serverless computing side, have made progress, we ought to leverage the existing expertise of the domain scientists while automating the creation of various services to leverage them in SaaS, microservices, and serverless computing.

Having worked with this community, we found that the educational steps involved for a beginner take about two to three months to get up to a the level where the development of cloud-based services is possible. We set the goal to explore if it is possible to drastically reduce the the time needed to create such services.

For this reason, we developed a sophisticated but easy to use framework that takes a *regular* python function and converts it automatically into a secure REST service and OpenAPI specifications [1] that can be reused in the ecosystem of cloud services. We used this framework to create many AI-based REST services to showcase the approach’s validity. We used examples from SciKit-learn [2] and benchmark the execution of the resulting REST services on various clouds and an IoT device.

The paper is structured as follows. In Sec. II we will give a very brief background overview to allow domain experts to catch up with the terminology and concepts used in our architecture. The background analysis leads us to our requirements presented in Sec. III and our architectural design shown in Sec. IV. Our benchmarks are collected in Sec. VI. We present our conclusion in Sec. VIII.

II. BACKGROUND

In this background section, we provide a small summary of activities related to this research so that domain experts can get an introduction to concepts that we use to implement

our architecture. It is beyond the scope of this paper to give more detailed introductions in topics such as IaaS, SaaS, microservices, serverless computing, OpenAPI, and REST services. The sections will, however, be useful as a starting point for further research to the reader.

A. The Big Data Reference Architecture

NIST has developed a Big Data Reference Architecture as part of the NIST Big Data Interoperability Framework (NBDIF) [3] and identified several use cases that motivate it [4]. The reference architecture is depicted in Fig. 1. It includes the following components: Data Provider, Big Data Application Provider, Big Data Framework Provider, Data Consumer and System Orchestrator as well as two overarching fabrics: security and privacy and system management. There are three types of linkages, namely *Big Data Information flow*, *Service Use* and *Software Tools*, and *algorithms transfer*. The architecture presents a level of abstraction to define Big Data applications. Components that implement sophisticated functionality work in concert to address the challenging creation of instantiating architectures beyond the conceptual stage. As such, the components interact with each other that are expressed through the linkages within the NBDIF. The next logical step is to explore how it can benefit and be used for analytics services.

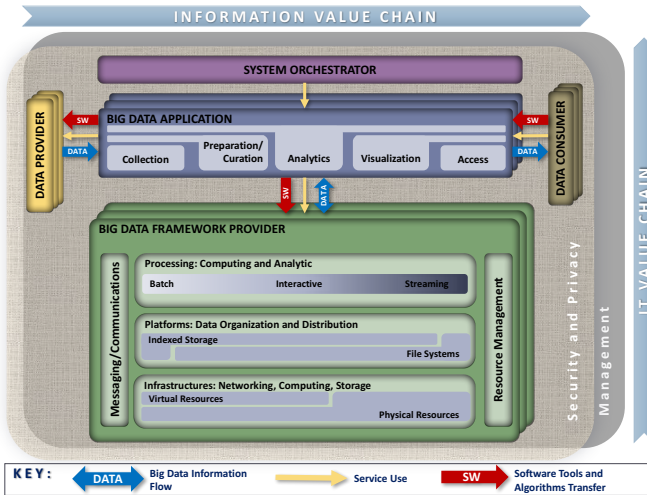


Fig. 1. NIST Big Data Reference Architecture [5]

NIST has developed through NBDIF open working group participation the following documents [3]–[11]. Within these activities, Volume 8 is of especial importance as it allows a set of Big Data Architectural needs [12] [13]. This effort builds the basis of our activities reported here while expanding it to cloud providers and services focusing on *Analytics Services*, which are not covered by the current volumes.

In a previous effort, we have developed a reference implementation that follows the architecture laid out in NBDIF and is easy to use by scientists. However, it focused mostly on multi-cloud provider access via REST services and command-line tools. The reference implementation is done as part

of the Cloudmeshproject [14], which was one of the first hybrid multi-cloud provider interfaces, even including cloud technologies that are no longer active such as Eucalyptus [15], OpenCirrus [16], FutureGrid [17], and Comet Cloud [18], demonstrating its longevity. Today, it supports clouds such as AWS [19], Azure [20], Google Cloud Platform [21], Oracle [22], and OpenStack [23]. It could offer further value by exploring the integration of MapReduce frameworks such as Hadoop [24] and Spark [25], as well as container-based frameworks such as Docker [26], and Kubernetes [27].

However, the work presented here focuses on creating analytics services that can be automatically created and hosted on any of the clouds supported by Cloudmesh. This is a non-trivial effort due to the large number of technologies involved and is outside of the expertise of domain scientists. However, the use of Cloudmesh makes it possible for the domain scientist to easily access these services and leverage our more than ten years of experience in this field.

The previous work provides us with a blueprint on how to proceed. We list the following main findings of our earlier work that we leverage as part of this work.

Software Defined Analytics Services and Applications.

Just as in the NBDIF working group, the utilization of *DevOps* to deliver Software-Defined (SD) Big Data applications is of utmost importance for the design of reusable services and components [14].

Multi-cloud Provider Interfaces. NIST BDRA Volume 8 [5] was through community input shaped in such a form that it allows multi-cloud interfaces. Such interfaces have been in practical use in our software and showcase the validity of the NIST-BDRA approach. It is clear that we need to introduce such multi-cloud and multi-service interfaces for analytics-related tasks whenever possible as motivated in our introduction.

Use Case Collection. NIST has provided as part of the NIST BDRA document Vol. 6 [3] several use cases that can be analyzed and from which common big data services can be detected. These use cases were sufficient to drive the NIST BDRA document [3] and allowed the community to investigate initial implementations. These use cases also motivate the work conducted in this effort.

Independent API Specification Leveraging OpenAPI.

Although the use of OpenAPI [1], [28] is not required as part of the NIST specification, it can be used to formulate services in a language-independent fashion. Hence it allows *creating, evolving and promoting a vendor-neutral description format*. This is important also for our analytics services approach to promote a vendor-neutral and independent effort.

API's and Tools Targeting A Multi-Layered Architecture.

In our previous effort, we learned that we need to provide support for tools, services, and APIs on multiple levels in a multi-layered architecture. While some users expect a generalized specification other users may require access on the command line, deployed services, or even a Jupyter notebook. We observe that the entry-level

to define API specification is too high for many. This is often the case for domain experts in the analytics community that may lack the necessary expertise for general service integration and deployment.

Hence, our previous experiences and work provides us with a blueprint on how to proceed, which we summarize as follows:

Create an easy to use framework that allows the scientists (a) to develop shareable analytics components (b) allows for the deployment of them, and (c) allows for the easy reuse of the services by community members leveraging the deployments.

B. OpenAPI

One of the important aspects of generating REST services is a language-independent. For this reason, the “OpenAPI Specification (OAS) defines a standard, language-agnostic interface to RESTful APIs which allows both humans and computers to discover and understand the capabilities of the service without access to source code, documentation, or through network traffic inspection. When properly defined, a consumer can understand and interact with the remote service with minimal implementation logic [1].”

Hence the purpose of the specification is not only to display the documentation but also to use it to generate the clients and server stubs from it automatically. OpenAPI can be formulated as a YAML Ain’t Markup Language (YAML) [29] file.

An OpenAPI definition can then be used by documentation generation tools to display the API, code generation tools to generate servers and clients in various programming languages, testing tools, and many other use cases. One of the issues with using the OpenAPI during the design of a project is that it takes considerable effort to understand the specification. Based on our experience of integrating it into university courses, it is a formidable effort to learn and use it. The lessons from this educational effort that includes researchers, professionals, graduate, and undergraduate students motivated this work.

C. Hybrid Multi-Cloud Computing with Cloudmesh

Cloud computing providers offer their customers on-demand self-service computing resources that are rapidly elastic and accessible via broad network access [30]. They accomplish this through the economies of scale achieved by resource pooling (serving multiple customers on the same hardware) and using measured services for fine-grained customer billing [30]. Cloud providers offer these resources in multiple service models including infrastructure as a service, platform as a service, software as a service, and, recently, function as a service [30]. These providers are rapidly offering new platforms and services ranging from bare-metal machines to AI development platforms like Google’s TensorFlow Enterprise platform [31], and AI services such as Amazon’s text-to-speech service [32].

Customers can take advantage of cloud computing to reduce overhead expenses, increase their speed and scale of service deployment, and reduce development requirements by using cloud providers’ platforms or services. For example, customers’ developing AI systems can utilize clouds to handle big

data inputs for which private infrastructure would be too costly or slow to implement. However, having multiple competing cloud providers leads to situations where service availability, performance, and cost may vary. Customers must navigate these heterogeneous solutions to meet their business needs while avoiding vendor lock-in and managing organizational risk. This may require comparing or using multiple cloud providers to meet various objectives.

Today’s infrastructure deployments can benefit from a *hybrid multi-cloud* strategy in which a mix of cloud-enabled services such as computing, storage, and other services are integrated from on-premises infrastructure, private cloud services, and a public cloud.

As pointed out earlier, Cloudmesh [14] is a framework and toolkit that enables users to easily access hybrid multi-cloud environments. Cloudmesh is an evolution of previous tools that have been used by many users. Cloudmesh makes interacting with clouds easy by creating a service mashup to access common cloud services across numerous cloud platforms. Cloudmesh contains a sophisticated command shell, a database to store JSON objects representing virtual machines, storage, and a registry of REST services [12]. Cloudmesh has a sophisticated plugin concept that is easy to use and leverages python namespaces while integrating plugins from different source code directories [33]. Installation of Cloudmesh is available for macOS, Linux, Windows, and RaspberryOS [14].

Cloudmesh works with a variety of cloud providers, including Amazon Web Services, Microsoft Azure, Google Cloud Platform, and Oracle’s OpenStack based providers such as the academic research serving Chameleon Cloud [34].

Recently we have also explored containers and microservices. The work presented here summarizes some of this effort. With the help of a plugin *cloudmesh-openapi* we can generate REST services, including microservices and containers, to organize its functions and code. In addition, Cloudmesh can be distributed as a container and used in a containerized environment. Through this ability, Cloudmesh services generated with *cloudmesh-openapi* can also be deployed on Kubernetes.

As pointed out earlier, Cloudmesh [14] is a framework and toolkit that enables users to easily access hybrid multi-cloud environments. Cloudmesh is an evolution of previous tools that have been used by many users. Cloudmesh makes interacting with clouds easy by creating a service mashup to access common cloud services across numerous cloud platforms. Cloudmesh contains a sophisticated command shell, a database to store JSON objects representing virtual machines, storage, and a registry of REST services [12]. Cloudmesh has a sophisticated plugin concept that is easy to use and leverages python namespaces while integrating plugins from different source code directories [33]. Installation of Cloudmesh is available for macOS, Linux, Windows, and RaspberryOS, and containers [14].

Cloudmesh works with a variety of cloud providers, including Amazon Web Services, Microsoft Azure, Google Cloud Platform, and Oracle’s OpenStack-based providers such as the academic research clouds.

Recently we have also explored containers and microservices. The work presented here summarizes some of this effort. With the help of a plugin *cloudmesh-openapi* we can generate REST services, including microservices and containers, to organize its functions and code. In addition, cloudmesh can be distributed as a container and used in a containerized environment. Through this ability, Cloudmeshservices generated with *cloudmesh-openapi* can also be deployed on Kubernetes.

III. REQUIREMENTS

Next, we present the most critical requirements that motivated our architecture and design. We start with a set of general requirements.

Leveraging new Python features. Python is a very popular choice with many data scientists. Our framework needs to leverage the newest Python 3 features such as *Typing Interface* [35] in order to increase robustness and future-proofing of our codebase.

Ability to be used within Jupyter Notebooks.

The framework must be able to integrate with Jupyter notebooks as they are very popular with today's data scientists. The functionality must be easily accessible not only as part of Python programs but also within Jupyter notebooks. This is of special importance also for using cloud services such as Google Colab [36] which for example offers cloud-based Notebooks.

Easy of use is a critical aspect of the framework that it addresses from the start ease of generation, deploying, and using the user-defined analytics services. This is accompanied by easy-to-use command-line tools.

Next, we list some more specific requirements that motivate our architectural design.

Multi-Cloud Service Integration. The framework must allow us to integrate multiple cloud services, including IaaS, PaaS, and SaaS. This also includes the ability to access AI-based services offered by the various cloud providers.

Hybrid-Cloud Service Integration. The framework must allow integrating on-premise, private, and public clouds.

Generalized Analytics Service Generator. We need a generalized analytics service generator. The first step in the activity to generate an analytics service is to provide an OpenAPI Service generator. Our generator will allow us to define essential analytics functions such as (a) uploading and downloading files to an analytics service; (b) specifying the functionality through typing enhanced Python functions; and (c) generating the code for the service.

Generalized Analytics Service Deployment. After the service is generated, it needs to be deployed. For this step, we will be reusing the Cloudmeshdeployment mechanism to instantiate services on-demand on specified cloud providers such as AWS, Azure, and Google.

Generalized Analytics Service Invocation. The next step includes the invocation of the deployed services. While

analyzing some use cases, we identified that users often need to invoke the same service many times to tune service parameters in a quasi-realtime fashion while using parameters that can not be included in the URL. Hence we will need to upload input parameters through files, if the simple typing data types provided by our proposed framework is not sufficient.

REST Services Architecture. As REST has become the most prominent architectural design principle, our Generalized service architecture needs to be able to produce REST services.

Automated REST Service Generation for other Languages.

Our framework must have provisions included that allow the integration into other programming languages and, on the other hand, allows the integration of services and functions developed in other languages.

Generalized Analytics Service Registry. As users and communities may develop many different services, we must provide the ability to (a) find specifications of generalized analytics services (b) find use-cases of generalized analytics services (c) find infrastructure on which such services can be deployed, and (d) find deployed analytics services. For this, we need a registry that can be queried by the community.

Generalized Composable Analytics Services. Services must be allowed to reuse other services to allow for easy integration. Thus we need to make our services composable. This also includes the choreography of the execution of such composable services.

IV. ARCHITECTURE

To satisfy our requirements, we have designed a multilayered architecture delivering a framework and toolkit to allow the easy generation and deployment of generalized AI-based REST services.

It contains two main layers. The first layer is concerned with generating the REST services, while the second focuses on easy deployment in a multi-cloud environment. However, as we also deal with hybrid infrastructure, we allow in the second layer access to HPC and local on-premise resources. In addition, our architecture addresses the creation of containers and their deployment in docker and Kubernetes. This way our framework is not only capable of deploying into cloud virtual machines, but also into other infrastructure services, either offered locally or in the cloud. Both layers can be accessed and control via a convenient command shell and client through Cloudmeshplugins. As we utilize REST services, the deployed services can easily be accessed from other resources via the REST API. Next, we describe some of the important features within each of the layers while starting with the infrastructure deployment.

A. IaaS Access Layer

This layer allows us to deploy different infrastructure services on demand while introducing an abstraction layer for compute resources that allow IaaS access across the

different platform offerings. Here, we can leverage from the Cloudmeshtoolkit that provides us with the basic interfaces to virtual machines to conveniently access in homogeneous fashion AWS, Azure, Google, Oracle, and OpenStack. Access to HPC and Bare metal can be integrated and has been showcased in previous versions of Cloudmesh. We also have prototyped Cloudmeshinterfaces for accessing compute resources via docker and Kubernetes.

All of the deployment can easily be managed through a simple client shell that can also be used as a command line executor. This system is one of the key components of Cloudmeshand allows easy integration of new commands and modules. This makes Cloudmeshextensible, while others can provide new functionality that can be accessed in the command shell and command line interpreter. We use the Cloudmeshcommand shell to integrate the functionality of the Generalized AI Service (GAS) Generator and describe its functionality in more detail next.

B. GAS Generator

The Generalized AI Service (GAS)¹ Generator creates the REST service from a simple function or class definition while utilizing the newest Python language features such as documented typing information integrated into the function or method signatures as part of the program specification. The GAS Generator provides us with the *fuel* that is needed as part of the service deployment. This is manifested in a number of artifacts for the deployment. The artifacts include a specification derived from the python program in OpenAPI format, the server code that is derived from the OpenAPI format, and an optional container specification file (e.g., Dockerfile). In addition, as we expect that the service is going to be reused, we use a GAS service registry in which we record the specification description of the service as well as deployment information on which the service ought to be deployed. This deployment specification can be derived from other prototype Cloudmeshcomponents such as *cloudmesh-frugal*, which can obtain resources based on minimal cost. We have not explicitly included this component in our architecture picture as we have not used it as part of our benchmarks or deployments that we describe later.

C. GAS Workflow

To showcase why the framework is so useful for data scientists, we are contrasting the service creation and use that a scientist undergoes while using OpenAPI without and with GAS Generator in Fig. 3 and Fig. 4.

The workflow in Fig. 3 showcases a typical workflow as promoted by the developers of Swagger codegen [37]. The user identifies from his use case an OpenAPI schema that becomes the basis to generate the code. However, this results in an unnecessary high entry barrier as the creation of OpenAPI schemas is complex. While using the swagger code generator,

¹The name GAS is derived from two different common usages. First, it refers to gasoline, referring to fuel that we need to generate the services; the second is an expandable material that fills the whole of a container.

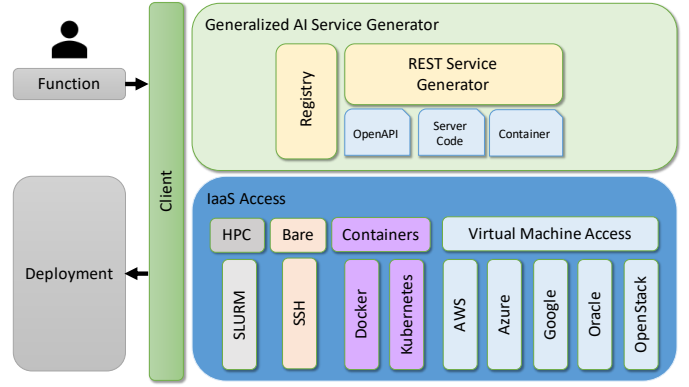


Fig. 2. Layered architecture of the Cloudmesh OpenAPI framework.

a variety of code stubs in many different languages can be created. The code generated requires an unnecessary high entry barrier as we next need to identify how and where we include an implementation of a function that we want to expose as a REST service. Once complete, the rest of the activity requires the remaining steps to be executed by hand, but scripts could be developed to automatize it.

Next, we like to contrast this with our much-simplified approach. As we know, the data scientists have the knowledge to write a Python function (or class); we simply leverage this expertise and take the function (or class) and provide it as input (fuel) to the GAS generator. This is done with a simple one-line command invocation that includes the name of the python program in which the function (or class) is defined. The scientist does not have to learn REST, the scientist does not have to look into a code stub that is generated for him, the scientist does not even have to know how to instantiate or run the service. Furthermore, the scientist does not have to know about any security as we have added features to the code to leverage the existing security mechanism as a simple flag to the GAS Generator command line instantiation. This simplification allows the scientist to develop REST services in minutes rather than a month, as the entry barrier is very low. Additionally, as we are integrated with Cloudmeshinfrastructure deployment capabilities, the instantiation of the services can also be done with a one line command under the assumption the scientist has accounts registered with Cloudmeshallowing him to authenticate and authorize the deployment of the services in the cloud.

Not only has our workflow fewer complex steps, they are also less interrupted, as the white colored boxes requiring user input in Figures Fig. 3 and Fig. 4 indicate.

D. Scripting as Fuel for the GAS Generator

Next, we demonstrate through two examples how simple it is to obtain Services from Python function and class specifications. In our examples, we will focus on functions and methods that return a result. We chose a simple add function and list the code related to it in Fig. 5. We also expand upon the example and use a class definition to showcase how to




Fig. 3. Schema-based component flow to specify an analytics service.




Fig. 4. Function-based component flow to specify an analytics service.

derive services using multiple paths instead of being deployed in different services as showcased in Fig. 7. As we see from the example, other than using the new typing feature provided in Python, the example is just a regular Python program. It can be tested locally on the system to check its functionality before we generate the service, providing a big advantage for development.

Fig. 6 shows how to generate and deploy the service. As this process is the same for the class-based definition and only differs in the filename, we omitted it to include an explicit listing of the access method for it.

Once the service is deployed the *curl* calls in Fig. 6 and Fig. 8 showcase how to interact with the service from the command line. Naturally, one can use any programming language that has built-in libraries for HTTP requests to interface with the service (such as *requests* in Python [38]). Once we execute the following lines in a terminal, the result of the addition will be calculated in the REST service, and it is returned.

E. GAS Security

As we leverage OpenAPI and automatically generated OpenAPI services, it is possible to leverage security mechanisms from the underlying service implementation. To showcase this ability, we added basic authentication into our framework as an example configuration. However, it is certainly possible to extend this as the services we use also support OAuth, ApiKey Authentication, Bearer Authentication (JWT), and HTTPS support [39] [12].

```

1 def add(x: float, y: float) -> float:
2     """
3     adding x and y.
4     :param x: x value
5     :param y: y value
6     :return: result
7     """
8     result = x + y
9     return result
  
```

Fig. 5. Defining an analytics function that is used to generate a REST service.

```

10 $ cms openapi generate add --filename=./add.py
11 $ cms openapi server start ./add.yaml
12 $ curl \
13   -X GET "http://localhost:8080/cloudmesh/add?x=1&y=2" -H "accept:_text/plain"
14 # This command returns
15 > 3.0
  
```

Fig. 6. Generating, deploying, and invoking the REST service.

```

1 class Calculator:
2
3     @classmethod
4     def multiply(cls, x: int, y: int) -> int:
5         """
6         Multiply int by int and return an int.
7
8         :param x: the value of input #1
9         :param y: the value of input #2
10        :return: result of multiplying x by y
11        """
12        return x * y
13
14     @classmethod
15     def divide(cls, x: int, y: float) -> float:
16         """
17         Divide int by float and return a float.
18
19         :param x: the value of input #1
20         :param y: the value of input #2
21         :return: result of dividing x by y
22         """
23        return x / y
24
25 if __name__ == '__main__':
26     calc = Calculator()
27     print("multiply_1_2:", calc.multiply(1, 2))
28     print("divide_6_3.14:", calc.divide(6, 2.3))
  
```

Fig. 7. Defining an analytics function with the help of class methods to generate a REST service with multiple functions.

```

29 $ curl \
30   -X GET "http://localhost:8080/cloudmesh/multiply?x=1&y=2" \
31   -H "accept:_text/plain"
32
33 $ curl \
34   -X GET "http://localhost:8080/cloudmesh/divide?x=6&y=3.14" \
35   -H "accept:_text/plain"
  
```

Fig. 8. Defining an analytics function with the help of class methods to generate a REST service with multiple functions.

To demonstrate basic authentication, a Cloudmeshuser can create an OpenAPI server whose endpoints are only accessible as an authorized user. Currently, when basic auth is used, all endpoints are secured with this method. In future versions, we intend to allow securing selected methods. An example of the usage of basic auth is provided on our Web page at [40].

V. DEPLOYMENT SCENARIOS

Due to the versatility of REST and our ability to integrate with a variety of infrastructure services, a rich set of deployment scenarios is possible. Two important scenarios related to single and multiple service provider deployments are discussed next.

A. Single Cloud Provider Hosted AI Service

In this scenario, a user deploys CloudmeshOpenAPI on a virtual machine from a cloud provider and uses it to host auto-generated, RESTful, AI services. To start, the scientist constructs an AI service as a set of Python functions that implement a workflow, for example, downloading data from a remote server, training an AI model, uploading a new sample for prediction, and running a prediction on that sample. After the deployment, the service is accessible using standard HTTP request methods. In Fig. 9 we show a remote client that accesses such a typical AI service workflow. Here the service is just deployed on a virtual machine from a single cloud provider. Cloudmesh provides the choice on which infrastructure provider to place the service. Through our security mechanism, the service can either be exposed to the public or to authenticated and authorized users.

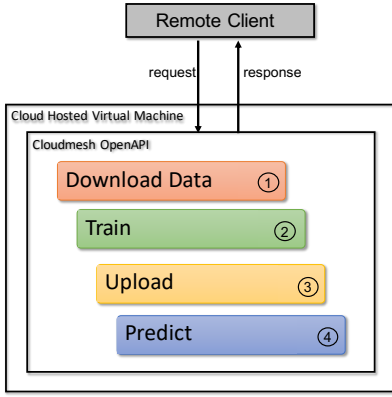


Fig. 9. Example AI Service Workflow to obtain data, train, upload data for prediction, as well as the interaction with it.

B. Multi-Cloud Hosted AI Service

In our next scenario, we like to focus on the possibility to deploy the same service on multiple clouds through the use of our sophisticated but easy-to-use command clients. Detailed information about the exact commands is provided in our manual [14]. Through this, we can, for example, evaluate suitable providers for our deployment through benchmarking the service execution on each provider and select those that meet our benchmark requirements. This is precisely what we will be showcasing in our benchmark section and demonstrate this approach's feasibility. Thus the scientist not only obtains the ability with GAS Generator to develop and deploy services, but also to evaluate their performance on a variety of infrastructures. An example is provided in Fig. 10 where a scientist deploys, for example, a service on AWS, Azure, and Google. As they are asynchronous services, the scientist can query the services simultaneously and gathers responses and benchmarks in parallel. Obviously, this can be used in a real scenario to integrate compute resources from multiple providers that can be accessed via our GAS services. It also allows specific adaptations such as the integration of cloud AI services with one provider that is not accessible in another.

Hence the framework can also be utilized to benchmark secondary services that are offered by a particular provider, or if they are offered by more than one provider, they can be comparatively evaluated.

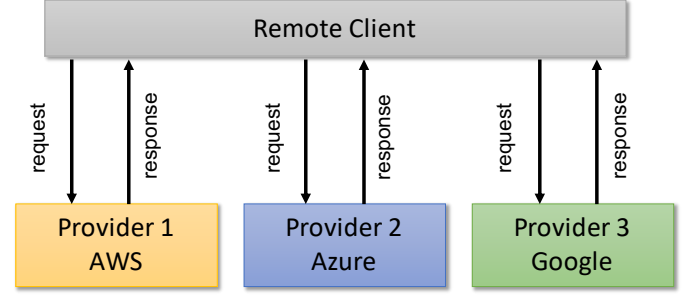


Fig. 10. Multi-Cloud AI Services: A client simultaneously accesses an AI service hosted on three separate cloud providers, AWS, Azure, and Google, to benchmark provider performance.

VI. BENCHMARK

In this section we describe our benchmark results.

A. Infrastructure

For a comparison of our services, we want to use service deployments on virtual machines that are hosted on various cloud providers. We have chosen to select similar virtual machines for conducting the benchmark. This includes AWS [19], Azure [20], and Google [21].

In addition, we are performing some bare metal experiments on two Raspberry PI clusters, one with all of them being Raspberry PI4's and the other with Raspberry PI with mostly 3b+'s. The latter has a management node, a PI 4, and worker nodes that are PI 3b+. The inclusion of the Raspberry platform was important to us as it demonstrates the capability of IoT and Edge computing devices that may become more prevalent in the future for delegating tasks to the edge. We further provide a docker container for a comparison of containerized services.

B. Application

We developed benchmark tests that utilize pytest to replicate Scikit-learn artificial intelligent algorithms on various infrastructures. These pytest are then run on different cloud services to benchmark statistics on how they perform.

The team obtained cloud service accounts from AWS, Azure, and Google. To deploy the pytest, the team used Cloudmesh and its OpenAPI based REST services to benchmark the performance on different cloud services.

Benchmarks include components like data transfer time, model train time, model prediction time, and more. Besides this report, scripts and other code are provided for others to replicate our tests.

We focus on two infrastructure benchmarks. The first deploys and measures the AI service on a single cloud provider at a time (see VI-E), and the second deploys a multi-cloud (see VI-F) AI service measuring the service across the clouds in parallel.

C. Algorithms and Datasets

This project uses a simple example algorithm and dataset. We have chosen to use an example included in Scikit-learn as they are widely known and can be used by others to replicate our benchmarks easily. Nevertheless, it will be possible to easily integrate other data sources, as well as algorithms, due to the generative nature of our codebase for creating REST services. Within Scikit-learn we have chosen the *Eigenfaces SVM Facial Recognition* example as it represents a very common data science usage pattern. This example conducts a facial recognition that first utilizes principle component analysis (PCA) to generate eigenfaces from the training image data, and then trains and tests an SVM model [41]. This example uses the real world *Labeled Faces in the Wild* dataset consisting of labeled images of famous individuals gathered from the internet [42].

D. VM Selection

When benchmarking cloud performance, it is important to identify and control VM deployment parameters. This allows one to analyze comparable service offerings, or identify opportunities for performance improvement by varying deployment features such as machine size, location, network, or storage hardware. These benchmark examples aimed to create similar machines across all three clouds, and measure their service performance. See Tab. I for a summary of the parameters controlled in these benchmark examples.

One key component is the virtual machine size, which determines the number of vCPUs, the amount of memory, attached storage types, and resource sharing policies. Resource sharing policies include shared core machine varieties—which providers offer at less expensive rates—that allow the virtual machines to burst over its base clock rate in exchange for credits or the machine’s inherent bursting factor [43], [44]. For this example, we chose three similar machine sizes that had comparable: vCPUs, underlying processors, memory, price, and were not a shared core variety. We installed the same Ubuntu 20.04 operating system on all three clouds.

Another factor that can affect performance, particularly in network latency, is the zone and region selected. We deploy all benchmark machines to zones on the east coast of the United States. This helps control variations caused by network routing latency and provides more insight into the inherent network performance of the individual cloud services.

Because cloud providers can observe varying loads during the day, the benchmark execution time is another parameter to control. In our single cloud provider benchmark for the Eigenfaces SVM example, clouds were tested at least twice and were run sequentially between the hours of approximately 19:45 EST and 03:30 EST starting with Google and ending with Azure. In the Eigenfaces SVM example, only 60 runs were conducted on Azure due to a failed VM deployment caused by factors outside of the benchmark script’s control. Compared to our single cloud provider benchmark, our multi-cloud benchmark benefits from all clouds being tested at the same time.

TABLE I
CONTROLLED VM PARAMETERS FOR CLOUD BENCHMARKS.

	AWS	Azure	Google
Size (flavor)	m4.large	Standard_D2s_v3	n1-standard-2
vCPU	2	2	2
Memory (GB)	8	8	7.5
Image	ami-0dba2cb6798deb6d8	Canonical:0001-com-ubuntu-server-focal:20_04-lts:20.04.202006100	ubuntu-2004-lts
OS	Ubuntu 20.04 LTS	Ubuntu 20.04 LTS	Ubuntu 20.04 LTS
Region	us-east-1	eastus	us-east1
Zone	N/A	N/A	us-east1-b
Price (\$/hr)	0.1	0.096	0.0949995

TABLE II
RASPBERRY PI AND DOCKER SPECIFICATIONS

	Pi 3B+	Pi 4	Docker (On MBP)	MacBook Pro i5 3.1GHz
Cores	4	4	2	2
Memory (GB)	1	8	2	8
OS	Raspberry OS 10	Raspberry OS 10	Ubuntu 20.04 LTS	macOS
Version	Kernel 5.4.51	Kernel 5.4.51	NA	Big Sur
Purchase Cost (\$)	51.99	109.99	NA	NA
Energy Cost (\$/year)	5.36	6.73	NA	NA
Price (\$/hr)	0.0065	0.0133	NA	NA

The Price is the purchase cost and 1 yr energy cost, amortized over a year and given for each hour of the year.

E. Single Cloud Provider AI Service Benchmark.

The benchmark script for the Eigenfaces SVM example uses Cloudmeshto create virtual machines and set up a CloudmeshOpenAPI environment sequentially across the three measured clouds, Amazon, Azure, and Google. After the script sets up the environment, it runs a series of pytests that generate and launch the Eigenfaces-SVM OpenAPI service, and then conducts runtime measurements of various service functions. Also, we run the same pytests on two Raspberry Pi models, a MacBook Pro running a Docker container, and a bare metal MacBook Pro to demonstrate CloudmeshOpenAPI’s flexibility for multi-platform use.

The benchmark runs the pytest in two configurations. After the benchmark the script sets up a virtual machine environment, it runs the first pytest locally on the OpenAPI server and measures five runtimes:

1. **download data:** Download and extraction of remote image data from ndownloader.figshare.com/files/5976015
2. **train:** The model training time when run as an OpenAPI service
3. **scikitlearn train:** The model training time when run as the Scikit-learn example without OpenAPI involvement
4. **upload local:** The time to upload an image from the server to itself
5. **predict local:** The time to predict and return the target label of the uploaded image

The benchmark runs the second pytest iteration as a remote client and interacts with the deployed OpenAPI service over the internet. It tests two runtimes:

1. **upload remote:** The time to upload an image to the remote OpenAPI server
2. **predict remote:** The time to run the predict function on the remote OpenAPI server, and return the target label of the uploaded image

In Fig. 11 we compare the download and extraction time of the labeled faces in the wild data set. This data set is approximately 233 MBs compressed, which allows us to measure a small but non-trivial data transfer. Lower transfer times imply the cloud has higher throughput from the data server, less latency to the data server, or that the cloud has a better performing internal network. The standard deviation is displayed to compare the variation in the download times. Because the difference between commercial and residential internet speeds dominates the function runtime, we do not compare the clouds to the Pi models, MacBook, or docker container.

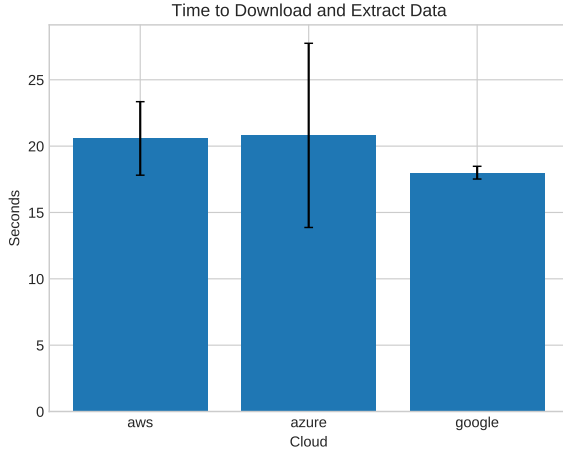


Fig. 11. Runtime for downloading the data used in the Eigenfaces SVM benchmark.

In Fig. 12 we measure the training time of the Eigenfaces-SVM model both as an OpenAPI service and as the basic Scikit-learn example. This allows us to measure the runtime overhead added by OpenAPI compared to the source example. Here, the two functions are identical except that the OpenAPI train function makes an additional function call to store the model to disk. This is necessary to share the model across the train and predict functions. In the Fig. 12 there are two bars per cloud provider. The blue bars are the training time of the model when hosted as a CloudmeshOpenAPI function. The orange bars are the training time of the Scikit-learn example code without CloudmeshOpenAPI involvement. The bars plot the mean runtimes and the error bar reflects the standard deviation of the runtimes. In Fig. 13 we show the same plot without the Pi models, MacBook, and docker results to allow a closer comparison of the three comparable clouds.

In Fig. 14 we measure the time to upload an image to the server both from itself and from a remote client. This allows us to compare the function runtime as experienced by the server, and as experienced by a remote client. The difference helps determine the network latency between the benchmark client and the cloud service. In the figure, there are two bars per cloud provider. The blue bars are the runtime of the upload function as experienced by the server, and the




Fig. 12. Runtime for training on the data used in the Eigenfaces SVM benchmark.

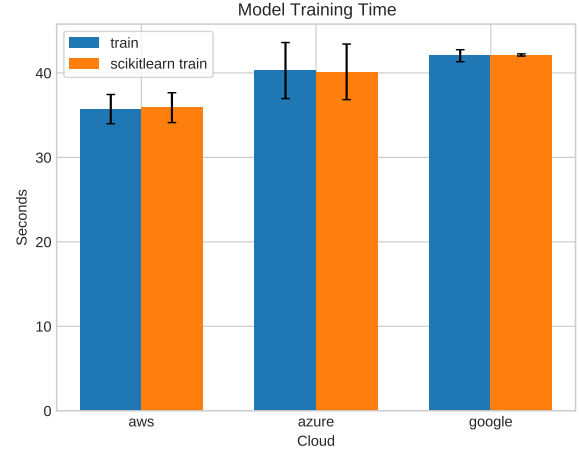


Fig. 13. Closeup of the Runtime for training on the data used in the Eigenfaces SVM benchmark without the data for the Pi.

orange as experienced by the remote client. The bars plot the mean runtimes and the error bar reflects the standard deviation of the runtimes. For the Pi models, MacBook, and docker container, we only measure the local function runtime.

In Fig. 15 we measure the time to call the predict function on the uploaded image. Again we run this once from the local server itself, and a second time from a remote client to determine client and server runtimes. In the figure, there are two bars per cloud provider. The blue bars are the run time of the predict function as experienced by the server, and the orange as experienced by the remote client. The bars plot mean runtimes and the error bar reflects the standard deviation of the runtimes. For the Pi models, MacBook, and docker container, we only measured the local function runtime.

Tab. III presents a full listing of test results. For the upload and predict tests, the *type* column denotes whether the test was run locally (server runtime) or remote (client runtime).

In Tab. V we present a cost analysis of the service functions.




Fig. 14. Runtime for uploading the data used in the Eigenfaces SVM benchmark.

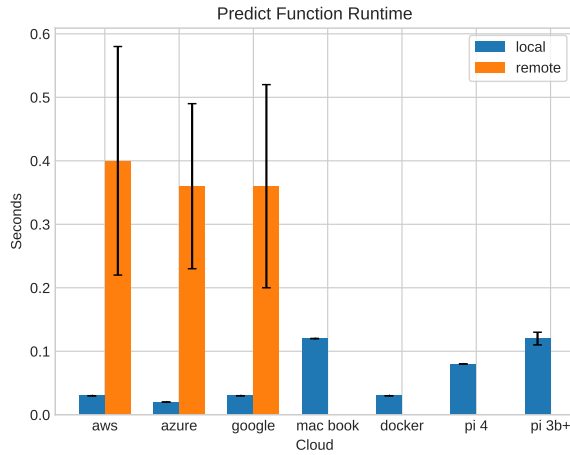


Fig. 15. Runtime for the prediction used in the Eigenfaces SVM benchmark.

The analysis uses the price from Tab. I and Tab. II. The price for the cloud virtual machines is based on provider advertised costs, while the price for the Pi models is based on the hardware cost and one year of energy cost amortized for one year. This does not include other costs such as cooling, networking, or real estate. For the Pi energy cost, we assume a full and constant load. We utilize power consumption benchmarks from [45] and Indiana residential kWh cost from [46] to calculate the expected Energy Cost per year. We calculate the cost to run each function and compare the clouds and Raspberry Pi 4 to the Raspberry Pi 3b+. We compare the percent runtime decrease from the Pi 3b+ to the clouds and Raspberry Pi4, and the percent cost increase from the Pi 3b+ to the clouds and Raspberry Pi 4.

F. Multi-Cloud AI Service Benchmark

In this benchmark, our script first acquires VMs, installs CloudmeshOpenAPI, and launches the Eigenfaces SVM AI

TABLE III
TEST RESULTS FOR THE EIGENFACES SVM SINGLE CLOUD PROVIDER BENCHMARK.

test	type	cloud	mean	min	max	std
download data	local	aws	20.58	17.23	31.80	2.77
download data	local	azure	20.81	13.56	42.70	6.94
download data	local	docker	820.98	820.98	820.98	0.00
download data	local	google	18.00	17.06	19.38	0.48
download data	local	pi 3b+	130.17	123.84	149.40	5.39
download data	local	pi 4	47.67	43.43	75.60	5.72
predict	local	aws	0.03	0.02	0.05	0.00
predict	local	azure	0.02	0.01	0.03	0.00
predict	local	docker	0.03	0.03	0.03	0.00
predict	local	google	0.03	0.01	0.06	0.00
predict	local	mac book	0.12	0.12	0.12	0.00
predict	local	pi 3b+	0.12	0.10	0.14	0.01
predict	local	pi 4	0.08	0.08	0.08	0.00
predict	remote	aws	0.40	0.26	0.80	0.18
predict	remote	azure	0.36	0.24	0.60	0.13
predict	remote	google	0.36	0.27	0.82	0.16
scikitlearn train	local	aws	35.89	35.11	46.45	1.77
scikitlearn train	local	azure	40.13	34.95	43.96	3.29
scikitlearn train	local	docker	53.76	53.76	53.76	0.00
scikitlearn train	local	google	42.13	41.77	42.49	0.13
scikitlearn train	local	mac book	32.53	32.53	32.53	0.00
scikitlearn train	local	pi 3b+	222.63	209.18	231.90	7.87
scikitlearn train	local	pi 4	88.32	87.78	89.14	0.33
train	local	aws	35.72	34.91	46.50	1.73
train	local	azure	40.28	35.30	47.50	3.32
train	local	docker	54.72	54.72	54.72	0.00
train	local	google	42.04	41.52	45.93	0.71
train	local	mac book	33.82	33.82	33.82	0.00
train	local	pi 3b+	222.61	208.56	233.48	8.40
train	local	pi 4	88.59	87.83	89.35	0.32
upload	local	aws	0.01	0.01	0.01	0.00
upload	local	azure	0.01	0.00	0.01	0.00
upload	local	docker	0.02	0.02	0.02	0.00
upload	local	google	0.01	0.01	0.01	0.00
upload	local	mac book	0.02	0.02	0.02	0.00
upload	local	pi 3b+	0.09	0.04	0.48	0.08
upload	local	pi 4	0.02	0.02	0.02	0.00
upload	remote	aws	0.43	0.16	1.13	0.21
upload	remote	azure	0.32	0.15	0.50	0.15
upload	remote	google	0.31	0.18	0.73	0.18

TABLE IV
TEST RESULTS FOR THE EIGENFACES SVM BENCHMARK DEPLOYED AS A MULTI-CLOUD SERVICE.

test	type	cloud	mean	min	max	std
download data	remote	aws	20.51	17.57	34.42	3.82
download data	remote	azure	18.60	13.49	32.65	4.53
download data	remote	google	17.90	17.13	21.86	0.85
predict	remote	aws	4.15	3.59	5.42	0.57
predict	remote	azure	3.93	3.40	6.65	0.74
predict	remote	google	4.13	3.74	6.37	0.60
train	remote	aws	35.61	35.24	39.53	0.73
train	remote	azure	35.89	35.08	40.00	0.95
train	remote	google	41.98	41.58	45.71	0.71
upload	remote	aws	10.08	4.89	16.52	4.38
upload	remote	azure	8.46	4.72	13.92	4.05
upload	remote	google	8.87	5.39	15.44	4.52

service on three separate cloud providers. To avoid side effects a script that serially deploys the VMs is used. After the services are running, we then run our tests in a parallel manner as depicted in Fig. 10. Testing in parallel provides faster benchmark results and better equalizes benchmark testing conditions. The benchmark conducts requests to each cloud in parallel, so they experience similar network conditions. For

TABLE V
COST ANALYSIS OF FUNCTION RUNTIMES WITH % COST INCREASE AND
% RUNTIME DECREASE RELATIVE TO THE RASPBERRY PI 3B+.

test	type	cloud	mean	cost	% runtime decrease	% cost increase
download data	local	aws	20.58	5.72e-04	NA	NA
download data	local	azure	20.81	5.55e-04	NA	NA
download data	local	google	18.00	4.75e-04	NA	NA
predict	local	aws	0.03	8.33e-07	75.00	281.87
predict	local	azure	0.02	5.33e-07	83.33	144.39
predict	local	google	0.03	7.92e-07	75.00	262.77
predict	local	mac book	0.12	NA	0.00	NA
predict	local	docker	0.03	NA	75.00	NA
predict	local	pi 4	0.08	2.96e-07	33.33	35.68
predict	local	pi 3b+	0.12	2.18e-07	0.00	0.00
predict	remote	aws	0.40	1.11e-05	NA	NA
predict	remote	azure	0.36	9.60e-06	NA	NA
predict	remote	google	0.36	9.50e-06	NA	NA
scikitlearn train	local	aws	35.89	9.97e-04	83.88	146.24
scikitlearn train	local	azure	40.13	1.07e-03	81.97	164.32
scikitlearn train	local	google	42.13	1.11e-03	81.08	174.60
scikitlearn train	local	mac book	32.53	NA	85.39	NA
scikitlearn train	local	docker	53.76	NA	75.85	NA
scikitlearn train	local	pi 4	88.32	3.27e-04	60.33	-19.26
scikitlearn train	local	pi 3b+	222.63	4.05e-04	0.00	0.00
train	local	aws	35.72	9.92e-04	83.95	145.10
train	local	azure	40.28	1.07e-03	81.91	165.33
train	local	google	42.04	1.11e-03	81.11	174.04
train	local	mac book	33.82	NA	84.81	NA
train	local	docker	54.72	NA	75.42	NA
train	local	pi 4	88.59	3.28e-04	60.20	-19.01
train	local	pi 3b+	222.61	4.05e-04	0.00	0.00
upload	local	aws	0.01	2.78e-07	88.89	69.72
upload	local	azure	0.01	2.67e-07	88.89	62.93
upload	local	google	0.01	2.64e-07	88.89	61.23
upload	local	mac book	0.02	NA	77.78	NA
upload	local	docker	0.02	NA	77.78	NA
upload	local	pi 4	0.02	7.40e-08	77.78	-54.77
upload	local	pi 3b+	0.09	1.64e-07	0.00	0.00
upload	remote	aws	0.43	1.19e-05	NA	NA
upload	remote	azure	0.32	8.53e-06	NA	NA
upload	remote	google	0.31	8.18e-06	NA	NA

example, in a serial testing model when downloading data from a remote server, the remote server may experience varying loads which will ultimately result in different throughputs for the various tests. Our parallel tests better equalize these conditions by having each cloud download the data under the same network conditions.

In the benchmark, we compute the means from 30 runs of a workflow that includes one download data invocation, one train invocation, 30 upload invocations and 30 predict invocations. We run the workflows in parallel on the separate clouds using multiprocessing on an eight-core machine.

In Fig. 16 we depict the combined runtime of our benchmark tests. This allows us to compare the complete execution time of an AI service workflow.

In Tab. IV we provide complete test results for the multi-cloud benchmark.

VII. LIMITATIONS

Azure has updated their libraries and discontinued the version 4.0 Azure libraries. We updated Cloudmeshto use the new library, but not all features, such as virtual machine delete, are implemented or verified.

VIII. CONCLUSION

This paper has introduced a framework and tool called GAS Generator that allows data scientists not experienced enough with REST and/or OpenAPI to generate REST services from python functions quickly. The overall time for deploying the resulting service was reduced from several months by




Fig. 16. Mean runtime of the Eigenfaces SVM workflow deployed as a multi-cloud service.

inexperienced data scientists to under a week. The service can be provisioned on public clouds and shared with other users. Authentication is built into our framework while leveraging common REST service practices. In a small benchmark executed on the various cloud providers as well as local hardware, including Raspberry PIs, we have seen that the cloud providers, when using similar resources and images, perform similarly. To compare the services with IoT devices such as Raspberry PI 3b+ and 4 we have chosen a small enough example that can be conducted on them and can be used as a reference to other IoT devices in the future. We found especially that in the case of the PI 4, the performance was quite good for our example. We also provided a cost-performance analysis to compare the IoT devices with the cost used on the cloud to conduct the task over a year's worth of activities. We find that the PI is surprisingly cost-effective.

However, our most significant gain from this project is the reduction in manpower and entry barrier it takes to create and deploy our AI services. Due to the generalized approach while using python function developers and data scientists can naturally integrate more complex tasks as well as tasks that leverage cloud-specific AI services that are uniquely offered by particular providers. GAS Generator is an open-source project, and we appreciate contributions to the project. Please contact the first author at laszewskigmail.com.

ACKNOWLEDGMENT

We like to thank Brian Kegerreis, Jonathan Beckford, Jagadeesh Kandimalla, Prateek Shaw, Ishan Mishra, Fugang Wang, and Andrew Goldfarb for developing the Python REST service generator on which this work is based on. We also like to thank the more than 70 contributors to the Cloudmeshtoolkit for developing the various cloud providers. We would like to thank the NBDIF working group for their contributions in discussions that lead to the development of this effort.

REFERENCES

- [1] OpenAPI, “Openapi specification,” Dec. 2020. [Online]. Available: <https://www.openapis.org/>
- [2] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [3] W. L. Chang, “Nist big data interoperability framework: Volume 6, reference architecture,” NIST, Tech. Rep., 2015. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1500-6.pdf>
- [4] —, “Nist big data interoperability framework: Volume 3, use cases and general requirements,” NIST, Tech. Rep., 2015. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1500-3.pdf>
- [5] G. von Laszewski and W. L. Chang, “Nist big data interoperability framework: Volume 8, reference architecture interface,” NIST, Tech. Rep., Oct. 2019. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1500-9r1.pdf>
- [6] W. L. Chang, “Nist big data interoperability framework: Volume 1, definitions,” NIST, Tech. Rep., 2015. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1500-1.pdf>
- [7] —, “Nist big data interoperability framework: Volume 2, big data taxonomies,” NIST, Tech. Rep., 2015. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1500-2.pdf>
- [8] —, “Nist big data interoperability framework: Volume 4, security and privacy,” NIST, Tech. Rep., 2015. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1500-4.pdf>
- [9] —, “Nist big data interoperability framework: Volume 5, architectures white paper survey,” NIST, Tech. Rep., 2015. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1500-5.pdf>
- [10] —, “Nist big data interoperability framework: Volume 7, standards roadmap,” NIST, Tech. Rep., 2015. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1500-7.pdf>
- [11] —, “Nist big data interoperability framework: Volume 9: Modernization and adoption,” NIST, Tech. Rep., Oct. 2019. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1500-10r1.pdf>
- [12] G. von Laszewski, “Cloudmesh openapi repository for automatically generated rest services from python functions,” Web Page, Dec. 2020. [Online]. Available: <https://github.com/cloudmesh/cloudmesh-openapi>
- [13] —, “Cloud computing,” Indiana University, Bloomington IN, U.S.A., Online Book, Sep. 2020. [Online]. Available: <https://cloudmesh-community.github.io/pub/vonLaszewski-cloud.pdf>
- [14] —, *Cloudmesh Manual*, cloudmesh.org, Dec. 2020. [Online]. Available: <https://cloudmesh.github.io/cloudmesh-manual/>
- [15] Eucalyptus Systems, Inc., “Eucalyptus,” Web page, Dec. 2020. [Online]. Available: <https://www.eucalyptus.cloud/>
- [16] R. Campbell, I. Gupta, M. Heath, S. Y. Ko, M. Kozuch, M. Kunze, T. Kwan, K. Lai, H. Y. Lee, M. Lyons, D. Milojevic, D. O’Hallaron, and Y. C. Soh, “Open cirrutmcloud computing testbed: Federated data centers for open source systems and services research,” in *Proceedings of the 2009 Conference on Hot Topics in Cloud Computing*, ser. HotCloud’09. USA: USENIX Association, 2009.
- [17] G. von Laszewski, J. Diaz, F. Wang, A. J. Younge, A. Kulshrestha, and G. Fox, “Towards generic futuregrid image management,” in *Proceedings of the 2011 TeraGrid Conference: Extreme Digital Discovery*, ser. TG ’11. New York, NY, USA: Association for Computing Machinery, 2011. [Online]. Available: <https://doi.org/10.1145/2016741.2016758>
- [18] R. Wagner, P. Papadopoulos, D. Mishin, T. Cooper, M. Tatineti, G. von Laszewski, F. Wang, and G. C. Fox, “User managed virtual clusters in comet,” in *Proceedings of the XSEDE16 Conference on Diversity, Big Data, and Science at Scale*, ser. XSEDE16. New York, NY, USA: Association for Computing Machinery, 2016. [Online]. Available: <https://doi.org/10.1145/2949550.2949555>
- [19] Amazon, “Amazon web services,” Web Page, 2017, accessed: 04/2018. [Online]. Available: <https://aws.amazon.com>
- [20] Microsoft, “Azure,” Dec. 2020. [Online]. Available: <https://azure.microsoft.com>
- [21] Google Staff, “Cloud computing services — google cloud,” Web Page, Dec. 2020. [Online]. Available: <https://cloud.google.com>
- [22] Oracle, “Oracle cloud infrastructure (oci),” Dec. 2020. [Online]. Available: <https://www.oracle.com/cloud>
- [23] N. Rackspace Hosting, “Openstack,” Web Page, Dec. 2020. [Online]. Available: <https://www.openstack.org/>
- [24] Apache, “Hadoop,” Dec. 2020. [Online]. Available: <https://hadoop.apache.org>
- [25] —, “Spark - Unified Analytics Engine for Big Data,” Sep. 2020. [Online]. Available: <https://spark.apache.org>
- [26] docker.com, “Docker: Empowering App Development for Developers,” Dec. 2020. [Online]. Available: <https://www.docker.com>
- [27] kubernetes.io, “Kubernetes,” Web Page, Dec. 2020. [Online]. Available: <https://kubernetes.io/>
- [28] “Openapi tools,” Dec. 2020. [Online]. Available: <https://github.com/OAI/OpenAPI-Specification/blob/master/IMPLEMENTATIONS.md>
- [29] Oren Ben-Kiki, Clark Evans, Ingy döt Net, “Yaml: Yaml ain’t markup language,” Web page, Oct. 2009. [Online]. Available: <https://yaml.org/>
- [30] P. Mell, T. Grance *et al.*, “The nist definition of cloud computing,” Computer Security Division, Information Technology Laboratory, National Institute of Standards and Technology, Tech. Rep., 2011. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>
- [31] Google, “Google cloud tensor flow enterprise,” Dec. 2020. [Online]. Available: <https://cloud.google.com/tensorflow-enterprise>
- [32] Amazon, “Amazon polly,” Dec. 2020. [Online]. Available: <https://aws.amazon.com/polly/>
- [33] G. von Laszewski, “Cloudmesh repositories,” Github, Dec. 2020. [Online]. Available: <https://github.com/cloudmesh/>
- [34] K. Keahey, J. Anderson, H. S. Gunawi, J. Mambretti, P. Ruth, and D. Stanzione, “Chameleon cloud,” Web Page, Dec. 2020. [Online]. Available: <http://chameleoncloud.org>
- [35] Python, “PEP 484 – Type Hints,” Dec. 2020. [Online]. Available: <https://www.python.org/dev/peps/pep-0484>
- [36] “Google colaboratory,” Web Page, Dec. 2020. [Online]. Available: <https://colab.research.google.com/notebooks/intro.ipynb>
- [37] Swagger, “Codegen,” Dec. 2020. [Online]. Available: <https://github.com/swagger-api/swagger-codegen>
- [38] Requests, “Requests: HTTP for Humans,” Dec. 2020. [Online]. Available: <https://requests.readthedocs.io/en/master>
- [39] Connexion, “Security Connexion 2.0 documentation,” Oct 2019. [Online]. Available: <https://connexion.readthedocs.io/en/latest/security.html>
- [40] cloudmesh.org, “Cloudmesh OpenAPI Security,” Dec. 2020. [Online]. Available: <https://github.com/cloudmesh/cloudmesh-openapi/blob/main/README-security.md>
- [41] SciKit Learn, “Faces recognition example using eigenfaces and svms,” Web Page, Dec. 2020. [Online]. Available: https://scikit-learn.org/stable/auto_examples/applications/plot_face_recognition.html
- [42] G. B. Huang, V. Jain, and E. Learned-Miller, “Unsupervised joint alignment of complex images,” in *2007 IEEE 11th International Conference on Computer Vision*. IEEE, 2007, pp. 1–8. [Online]. Available: <http://vis-www.cs.umass.edu/papers/iccv07alignment.pdf>
- [43] Amazon, “Amazon ec2 instance types - amazon web services,” Dec. 2020. [Online]. Available: <https://aws.amazon.com/ec2/instance-types/>
- [44] Google, “Google cloud machine types,” Dec. 2020. [Online]. Available: <https://cloud.google.com/compute/docs/machine-types>
- [45] J. Geerling, “Power Consumption Benchmarks: Raspberry Pi Dramble,” Web Page, Sep 2020. [Online]. Available: <https://www.pidramble.com/wiki/benchmarks/power-consumption>
- [46] I. E. Association, “Electricity Prices of Indiana Energy: Why are Electric Costs Increasing?” Jun 2019. [Online]. Available: <https://indianaenergy.org/why-are-electric-costs-increasing-electricity-prices>