

Cloud Hosting Cost Intricacies For IoT Applications

by Satya K Vivek | July 25, 2017



If you're struggling to get a handle on cloud costs for IoT applications, you're not alone. IT executives are struggling to keep up with constantly shifting cost structures in the midst of a prolonged price war between Amazon Web Services (AWS), Google Cloud Platform, Microsoft Azure, etc. In addition, defining your application's unique usage requirements, and managing your cloud migration and deployment costs both pose significant challenges.

Your initial IoT Application architecture discussions should include IT experts who will consider the scope of the following major subsystem usages on the cloud as part of the App design and implementation:

- 1) Device Gateway – the interface between connected devices (sensors/actuators) and the cloud server.
- 2) API Gateway – the interface between Mobile/Web applications and the cloud server.
- 3) Back-end server/services – handle application business logic and data processing/ management.
- 4) Notification Services - the interface to notify various events to the users via different mechanisms like Push notifications, Text message, eMail, etc.
- 5) Database – storing the data collected from connected devices, applications, etc.

Device Gateway

The purpose of the device gateway is to enable secure and reliable interaction with cloud applications and connected devices. IoT projects have additional complexities when compared to other cloud-centric technology applications, such as:

- Diverse hardware

- Diverse operating systems and software on the devices
- Different network gateway requirements

Commonly used protocols to communicate to the cloud server are MQ Telemetry Transport (MQTT), Advanced Message Queuing Protocol (AMQP) and HTTP. The device gateway can be implemented using various components like MQTT/ AMQP brokers (ActiveMQ, HiveMQ, RabbitMQ etc.), REST API servers etc., and deployed on private/ public cloud virtual machines or standalone servers. The advantages of using services for the Device Gateway are listed below:

- **Security & Authentication:** The services provide strong authentication modes, including: mutual authentication, token based authentication, as well as supporting encryption at all points of connection. Data is never exchanged between devices and device gateway without proven identity.
- **Scalability:** The services are auto-scalable and support millions of devices and messages.
- **Availability:** The services are replicated over multiple data centers to provide highly available services.
- **Ease of Development:** The device gateway service provides additional features for device management. Most of the services also provide SDKs in popular programming languages for both device management and messaging.
- **Support for Full Stack Development:** Services offer support on full stack development of IoT data, including: data storage, rule based processing, analytics, etc.
- **Logging and Monitoring:** All cloud service providers are offering good support on logging and monitoring.

Selection of the device gateway and its properties will vary based on the use cases and other aspects including bidirectional connectivity, distributed computing, availability, scalability, security and reliability. The cost of the device gateway will vary based upon these factors. The table below provides a comparison on the device gateway offered by major cloud providers.



| Cloud Service Provider | Device Gateway |
|------------------------|---|
| Microsoft | Service: Azure IoT Hub Protocol supported: HTTP, AMQP, WebSocket, MQTT. Azure IoT hub also allows to push custom protocols Price: Unlimited devices, 400,000 messages per day: \$50/month Unlimited devices, 60,00,000 messages per day: \$500/month Unlimited devices, 3000,00,000 messages per day: \$5000/month Message Block size for billing: 4KB Max Message size: 256 KB Security: SSL support, Per-device with SAS token |
| Amazon | Service: AWS IoT Protocol supported: HTTP, MQTT, Websockets Price (US region): \$5 per 1M messages Message Block size for billing: 512 bytes Max Message size: 128 KB Security and Authentication: SSL support, X.509 certificate based authentication, IAM role based security, or 3rd party authentication via AWS Cognito |
| Google | Service: Cloud IoT Core (Private Beta version) Protocol supported: MQTT Price: Not available Security and Authentication: SSL support, Certificate based authentication, IAM role based security |
| IBM | Service: IBM Watson IoT Protocol supported: HTTP, MQTT Price: Tier 1: 1 MB - 450GB \$0.0013 USD per MB Exchanged Tier 2: 450GB - 7TB \$0.00091 USD per MB Exchanged Tier 3: 7TB and above \$0.000182 USD per MB Exchanged Max Message size: 128 KB Security: SSL support, Per-device with token based authentication |

The major factors affecting the costing of a device gateway are the selection of the messaging protocol and its properties.

- MQTT and AMQP are lightweight publish/subscribe messaging protocols that support bidirectional communication between the devices and cloud server. These protocols require keep-alive (PING Messages) functionality to ensure the connection is still open and both broker and client are connected. HTTP messages work on request - response mode and connection is initiated from a client. The cost of the device gateway is directly proportional to the number of messages and the selection of a messaging protocol for a use case is an important aspect for the cloud host pricing.
- The selection of various properties in the messaging protocol impacts the cost of a device gateway. For example, the number of PING requests will vary based on the keep-alive interval. Setting smaller keep-alive intervals ensures better connectivity between the devices and cloud server and this will increase the number of messages, which will increase the cloud cost.
- Message size is another factor having a cost impact. Many players have a defined message block size for billing purposes. If you assume the message block size is 4KB, and you send a message of 4.1KB, this will be treated as 2 messages.
- For most cloud providers costs vary based on region. For example, AWS IoT service costs \$5 per million messages in US East (N. Virginia) region, but \$6 per million messages in Asia Pacific (Sydney) and \$8 per million messages in Asia Pacific (Singapore) regions.

API Gateway



APIs allow software applications to exchange data and play an important role in modular architecture designs. An API gateway optimizes communication between clients and the application and it also extend the benefits of an API management platform for hybrid cloud computing deployments. API Gateway benefits include:

- **API Lifecycle Management:** API Gateways helps centralize management of all the APIs. It supports staging and versioning that can be aligned with the rest of the application providing flexible deployment options.
- **Elastic:** API Gateways are elastic and can scale-out and scale-in dynamically without manual configuration.
- **Flexible deployment:** API backend integration and deployment can be easily done. Most providers will support the integration of APIs with their backend services.
- **API Logging, Caching, Throttling and Monitoring:** API Gateways provide inbuilt mechanism for logging, caching, throttling and monitoring.
- **Native Code Generation:** Many providers offer options to generate native SDKs. Once the API is deployed developers can generate SDKs for Android, JavaScript and iOS.
- **Security and Authentication:** The services provide various security solutions including SSL, API authentication, API Keys management, API quota, rate limit, etc.

The table below provides a comparison of the API Gateways offered by major cloud providers.

| Cloud Service Provider | API Gateway |
|------------------------|---|
| Microsoft | Service: Azure API Management Price for Premium tier: \$91.91 per day (supports 32 M API calls/ day and 161 GB data transfer / day) Price for Standard tier: \$22.55 per day (supports 7 M API calls/ day and 32 GB data transfer / day) Price for Standard tier: \$1.59 per day (supports 32 K API calls/ day and 161 MB data transfer / day) Security: Secure APIs with key, JSON Web Token (JWT) validation, and IP filtering |
| Amazon | Service: AWS API Gateway Price: \$3.50 per 1M messages and \$0.09/GB for the first 10 TB Security: Support IAM and Amazon Cognito, support custom authorizers written as AWS Lambda functions |
| Google | Service: Google Cloud End Points Price: 0-2M API calls per month- \$0.00 per million messages 2M+ API calls per month- \$3.00 per million messages Security: JSON Web Token validation and a streamlined developer experience for Firebase Auth, Google Auth and Auth0. |
| IBM | Service: IBM API Connect Price: Enterprise Tier - \$100.00 USD/100K API calls per month (25 million calls per month) Professional Tier - \$80.00 USD/100K API calls per month (5 million calls per month) Free Tier - \$0.00 USD/100K API calls per month (50K calls per month) |

- The number of API calls and the data transferred are the factors affecting the cost of the API gateway.
- For some of the cloud providers, the cost will vary based on the region selected.

Backend Server/Services

Backend server/services will handle application business logic and data processing/management. The server applications can be deployed on a cloud Virtual Machine or deployed as server-less functions on the cloud. The server-less functions have the following benefits:

- **Lower operational and development costs:** Pay only for the compute time and the resources used.
- **Easily scalable:** Administrators do not need to add or upgrade existing servers.



- Simplifies packaging and deployment and requires no system administration.
- Reduces software complexity.
- Logging and monitoring: Server-less platform providers offer good logging and monitoring options

Server-less architectures have the following drawbacks:

- Server-less is not efficient for long-running applications: Usually there will be limitation for maximum duration of execution of a function. In certain cases using long tasks can be more expensive than running a workload on a dedicated server or virtual machine.
- Vendor lock-in: Your application is completely dependent on a third-party provider.
- Server-less applications take some time to handle the first request: This problem is known as “cold start” —a platform needs to initialize internal resources.
- Most of the providers restrict the maximum function deployment package size as well as maximum execution duration per request.

The server-less architecture is more suitable for Internet of Things applications due to:

- Server-less based computing models are well-suited to event based IoT applications.
- Server-less architectures dynamically manage IoT applications, which need to be highly scalable.

The table below provides a comparison of the server-less functions offered by major cloud providers.

| Cloud Service Provider | Backend Server/ Service |
|------------------------|---|
| Microsoft | Service: FUNCTIONS Price: \$0.20 per 1 million requests and \$0.000016 for every GB-s Note: 1M free requests per month and 400,000 GB-seconds of compute time per month is free. There will be additional cost for Data Transfer. Gigabyte Seconds (GB-s): Observed resource consumption is calculated by multiplying average memory size in Gigabytes by the time in seconds it takes to execute the function. |
| Amazon | Service: AWS Lambda Price: \$0.20 per 1 million requests and \$0.00001667 for every GB-second Note: 1M free requests per month and 400,000 GB-seconds of compute time per month is free. There will be additional cost for Data Transfer. GB-second is calculated using "number of seconds the Lambda function executed" and "the allocated memory for the function to execute" |
| Google | Service: CLOUD FUNCTIONS (BETA VERSION) Price: \$0.40 per 1 million invocations, \$0.0000025 per GB-second and \$0.0000100 per GHz-Second Note: 2M free requests per month and 400,000 GB-seconds and 200,000 GHz seconds per month is free. Monthly Data Transfer of 5GB is free. The additional data transfer cost is \$0.12 per GB. 1 GB-second is 1 second of wallclock time with 1GB of memory provisioned 1 GHz-second is 1 second of wallclock time with a 1GHz CPU provisioned |
| IBM | Service: OpenWhisk Price: \$0.000017 per second of execution, per GB of memory allocated Note: 400,000 GB-seconds of compute time per month is free. There will be additional cost for Data Transfer. Gigabyte Seconds (GB-s) is calculated using "memory allocated" and "time to execute the function". |

The major factors affecting the costing of a server-less function are given below:

- The cost of server-less function directly depends on the time to execute the function. Writing efficient code with good response times will help to reduce the costing.



- Computing resources is another important cost parameter. We should choose the computing resource depending on the memory intensity of the function. Server-less functions use 100mSec for billing. e.g., if a function will get executed in 10 mSec on 1024MB compute engine, the GB second is calculated as **0.1 GB-Sec** ($0.1 \times 1024 / 1024$). If the same function will get executed in 50 mSec on 128MB, the GB-second is calculated as **0.0125 GB-Sec** ($0.1 \times 128 / 1024$). Selection of proper computing resources will reduce costs.

Notification Services

Faster communication is an intrinsic benefit and driver for the Internet of Things. The efficient and faster modes of notification servers are needed to establish real-time monitoring or controlling systems in IoT. The various methods of notifying end users are Push notifications, Text message, eMail, etc. Notification Services are the interface to end users via these different methods. Push notifications are used for send notifications from backend server/application to various user interfaces like mobile applications, desktop applications, etc. eMail and Text messages are used to send notifications to users via their eMail or Phone numbers.

The table below provides a comparison of notification services offered by major cloud service providers.

| Cloud Service Provider | Notification Service |
|------------------------|---|
| Microsoft | Service: Notification Hub Push Notification: Basic Tier: \$10/month (it includes 10 Million free push notification) Supports iOS, Android, Windows, or Kindle devices |
| Amazon | Service: AWS SNS (Simple Notification Service) Push Notification: \$0.50 USD / Million Messages Supports push notifications to iOS, Android, Fire OS, Windows, Baidu-based devices. eMail: \$2.00 USD / 100000 Messages Text Messages: Price will vary depending on countries. Transactional SMS costs \$0.00645 in US Supports 200+ countries. Service: AWS SES (Simple Email Service) Price: \$0.10 per 1,000 eMails |
| Google | Service: Firebase Cloud Messaging Supports Push Notification to iOS, Android, or Web (JavaScript) client app with no cost |
| IBM | Service: Push Notification Push Notification: Price: \$1.00 USD / Million Digital Message |

The decision on the mode of notification plays a vital role in the cost of Notification services.

Push notifications are the cheapest option when compared to other methods - eMail and SMS. Push notifications can be implemented with minimal to no cost, however, Text messages (SMS) are very expensive, and performance and price vary. Designing an IoT system with SMS notifications only for critical events, eMail for medium events and Push notification for all type of events will help to reduce overall costs.

Database

A database is used for storing the data collected from connected devices, applications, etc. A NoSQL Database is more suitable in IoT than Relational databases due to these factors:

- NoSQL databases are designed to support huge amount of data and with the IoT constantly expanding (50 billion objects predicted to be connected to the network by 2020) NoSQL databases are the right choice for IoT applications.



- More and more heterogeneous data is generated by exponentially growing number of diverse sensors and devices. This leads to the need of highly flexible data models. NoSQL databases offer more flexibility in data models.
- Analytics is becoming the cornerstone of the Internet of Things, requiring enhanced and multiple analytical approaches to address the requirements of the applications. NoSQL databases support multiple analytical approaches.

The table below provides a comparison of the NoSQL Database services offered by major cloud providers.

| Cloud Service Provider | Database |
|------------------------|--|
| Microsoft | <p>NoSQL DB: Azure Cosmos DB Price: \$0.25 per GB / month \$0.008/hr (throughput reserved in units of 100 RUs/second)</p> <p>Note: The standard datatransfer rate is also applicable. A Request Unit (RU) is the measure of throughput in Azure Cosmos DB. 1 RU corresponds to the throughput of the GET of a 1KB item.</p> |
| Amazon | <p>NoSQL DB: Dynamo DB Storage Price: 0.25 per GB-month Write Throughput: \$0.0065 per hour for every 10 units of Write Capacity (enough capacity to do up to 36,000 writes per hour)* Read Throughput: \$0.0065 per hour for every 50 units of Read Capacity (enough capacity to do up to 180,000 strongly consistent reads, or 360,000 eventually consistent reads, per hour) Note: The standard data transfer rate is also applicable.</p> |
| Google | <p>NoSQL: Google Cloud Data Store Price: \$0.18 / GB / month \$0.06/100k read operations \$0.18/100k write operations \$0.02/100k delete operations</p> |
| IBM | <p>NoSQL: IBM Cloudant \$1.00 per GB of data storage / month \$0.25 USD/Lookup per second \$0.50 USD/Write per second \$5.00 USD/Query per second</p> |

Storage utilization and the number of concurrent operations are major factors effecting the cost of the database service. The IoT application with optimal database operation (read/ write/ delete) will help to reduce the cost of Database services.

For most cloud providers, the cost will vary based on the selected region. For example, Azure Cosmos DB will have the storage cost of \$0.25 GB/month in Central US region, but Azure Cosmos DB will have the storage cost of \$0.2825 GB/month in Japan East region.

About Gadgeon Smart Systems

Gadgeon is a Product Engineering Services company helping clients in realizing their ideas in the Internet of Things, Communication Service Provider, Healthcare, Industrial Automation, and Networking spaces. At Gadgeon, we have been delivering end to end solutions from requirements to deployment, from hardware to application.



With a team of 200+ experts with required breadth and depth of skills, domain knowledge, Program management and Process expertise we have been able to give significant ROI for both startups and established players in terms of quality and time to market.

