The 9 verbs of IoT Cloud Platforms

A white paper by Joakim Fernstad

Summary

This white paper is aimed at CTOs, VPs of Engineering and Product Development Managers for companies developing electronics products and plan to create products with connectivity to a cloud.

There is a market expectation that new electronics products are connected to the internet. Regardless of the product, whether it's a simple sensor node or a complex autonomous vehicle, the decision to add connectivity to a product opens a company up to many new situations, and while they have a clear view of what value the connectivity add to their customers, it is less clear how to facilitate that value on the cloud side. Should they create their own IoT cloud infrastructure or rely on existing IoT cloud platforms for their solution? The answer is not obvious since the upside of using an existing IoT cloud platform provider can be great in the short run, it can however be expensive in the long run.

In June 2017 there were 450 global IoT cloud platform providers[1], and the number has certainly increased since then. The amount of research required to understand IoT cloud platform providers offerings is massive and delaying an decision for too long can lead to loss of market share.

By following the advice of "The 9 verbs of IoT Cloud Platforms" companies can save time and effort by limiting the number of suitable IoT cloud providers, while be confident that the chosen provider is delivering a solid base for their IoT cloud solution. This white paper is also a valuable resource for companies considering creating their own IoT cloud solution from scratch.

[1] https://iot-analytics.com/iot-platforms-company-list-2017-update/

Glossary

Cloud	Server solution in a remote data center where physical hardware is abstracted into virtual hardware, allowing multiple services execute separately but on the same hardware.
Cluster	Type of server architecture that shares storage and/or CPU load of a service.
Sharding	Type of database architecture that divides stored data amongst several servers, where no server contains the full set of data.
IoT Cloud Solution	The final solution built on top of the IoT Cloud Platform.
IoT Cloud Platform	The platform enabling the final IoT Cloud Solution.
Time Series data	Numerical values that are informative only together with a timestamp.
Replica	Identical copy of a cloud component, often part of a horizontal scaling solution.
Cloud Orchestration	Software responsible for directing traffic within a cloud and ensure that components are up and running. It is also responsible for scaling services.
Container	A modern type of virtual machine.
Serverless Computing	Next revolution in cloud computing, where a service does not exist when it is not in use.
Provisioning	The act of giving a device an identity and rights within the IoT Cloud Solution.
X509 certificate	A file which contain authenticated and verifiable information. Commonly used to authenticate web sites domain names.
TLS	Transport Layer Security, the de facto encryption protocol for internet communication.
ACL	Access Control Lists, defines access rights to specific resources within a cloud.
FOTA	Firmware Over The Air update, firmware update after a device is deployed
IaC	Infrastructure as Code. A way to describe cloud infrastructure and connections between components by writing code rather than manually configuring them.

Background

When a company decides to create a connected device, they will have have to think about more than the product's primary functionality. Aspects like security, upgradability, and cloud architecture are crucial to consider early on and can take up a considerable part of the total project time.

There are many cloud providers that offer a plethora of services to help companies focus on product development and not other aspects, thus reducing the time to market. At the same time, most IoT cloud solutions offer similar functionality and expand them with specific applications on top.

The question remains: "Should we use an existing IoT cloud platform or build our own?". Regardless of the answer, "The 9 verbs of IoT Cloud Platforms" will point you in the right direction to finding an IoT cloud platform that works for your company and help you select components for your specific IoT cloud solution.

Premise

At Endian, we are convinced the best IoT cloud platform is secure, based on standard protocols, in your full control, and ready for application development with minimal effort. This conviction is the base for our research.

Standard protocols
Relying on proprietary or obscure protocols can jeopardise the future of your solution by not being able to switch to cloud components with similar functionality, or by having to develop a solution to get around bugs in protocols which are not mature enough for production purposes. By using standard and open protocols, the chances of being able to swap out a cloud component, without negatively affecting the rest of the solution, increases greatly. Prepare for future changes by using well tested and widely used protocols for all communication in your solution.

In your full control
Using cloud provider services for your solution gets you going fast and it will most likely work well for a long time. However, it can be a challenge to understand how cloud providers charge for their services, and when you reach a scale where the cost is outweighing the benefit, it can be difficult and expensive to transfer to another provider's platform. This is the vendor lock-in conundrum. Again, by making sure the IoT cloud platform use standard and open protocols you reduce your reliance of specific providers and open up for future migration of your services.

Ready for application development
A great IoT cloud platform has all the supporting components you need, configured and ready to go, and enable your developers focus on product development. This includes a strategy to secure communication between your devices and your cloud solution.

The 9 verbs of IoT Cloud Platforms

The essential functionality of an IoT cloud platform can be described with 9 verbs. They should provide a guide to which components to look for when creating an IoT cloud solution and will help in selecting a suitable IoT cloud platform provider.

SCALE

One of the strongest cases for having your IoT solution in a cloud, rather than on-premise, is because of a cloud's scaling capabilities. Adding more resources to a cloud does not require new hardware purchase, instead you lease another virtual server and your solution can take full advantage of these added resources.

Modern IoT cloud solutions are based on Cloud Orchestration software and can be mixed with Serverless Computing for extra flexibility. Modern solutions also allow the cloud infrastructure to be described using computer code. This is called Infrastructure as Code (IaC) and rather than manually having to manage the infrastructure, it is described using written logic and managed by the Cloud Orchestration software.

Vertical Scaling
Before cloud computing, additional computing resources where added by purchasing more powerful hardware and moving the solution to this new hardware. This meant ordering servers, migrating services over to the new hardware, and decommissioning the old servers. The term for this is Vertical Scaling and it is cumbersome and time consuming.

Horizontal Scaling
The opposite of Vertical Scaling is Horizontal Scaling where a service is replicated on several, often less powerful, machines, and the load on the service is distributed evenly over these machines. Scaling happens by cloning, or replicating, a service onto new hardware.

Virtual servers
Hardware can nowadays be virtualised and several virtual servers can occupy the same hardware, sharing computing resources and improving hardware utilisation. Virtualisation enables easy replication of services by simply copying the virtual server onto more machines, and thereby scale horizontally. This is the definition of the cloud; virtual servers in data centers which can easily be moved and/or replicated.

Containers
A more modern type of virtual machine is a Container[2], which differs from virtual machines in that it shares the host machine's kernel for execution. Containers are popular these days because of their short startup time, small sizes and ease of management.

Cloud Orchestration
With the rise of Hyperscale data centers[3], a solution was needed to manage these virtual servers and containers effectively. This solution is called Cloud Orchestration and by basing your IoT solution on Cloud Orchestration software, you can easily manage and scale computing resources. In order to fully be able to utilise the scaling functionality of a cloud, the components chosen for your IoT cloud solution must have the ability to scale horizontally. This is often achieved with replicating, clustering or sharding technologies.

Serverless Computing
Cloud Providers now offer the next revolution of cloud computing, called Serverless Computing[4], where virtual servers are no longer used and instead cloud resources are allocated for specific functions and only when needed. E.g if a data analysis component executes every hour it will only exist and cost money while it is executing. Serverless Computing is still in its infancy and is not suitable for all types of services, e.g when persistent connections are required between devices and the IoT cloud solution.

SECURE

Your devices will be deployed on networks outside your control, and there is no guarantee that device data only travels through friendly networks. To ensure that device data is not manipulated or captured on its way to your cloud solution, the communication link must be encrypted. The encryption solution will depend on what protocols are used for data transfer and the device's hardware capabilities. When possible, protect communication links using TLS. Securing the communication links are only a part, although a vital part, of your full security solution.

---

[2] https://en.wikipedia.org/w/index.php?title=Operating-system-level_virtualization&oldid=832533207

[3] http://www.areadevelopment.com/data-centers/Data-Centers-Q1-2018/the-surge-of-hyperscale-data-centers.shtml

[4] https://en.wikipedia.org/w/index.php?title=Serverless_computing&oldid=835406190

PROVISION

Provisioning is all about Device Management and it is a big piece of your security solution. A device should have a unique identity used to authenticate with your IoT cloud solution and be given a minimum amount of access rights. Access Control Lists (ACL) is a common way to provide device access to specific cloud resources. Proper provisioning of devices protects your IoT cloud solution from imposters or cloned devices and stops compromised devices from accessing data owned by other devices. IoT cloud platform providers should have provisioning solutions available.

Refrain from using the same username / password pair for all devices. It only takes one compromised device to compromise them all, and put the rest of us at risk[5]. Using username / password pair as credentials for authentication in a secure way requires encryption of the communication links or they can be captured by bad actors while being transmitted to the IoT cloud solution. A stronger authentication scheme is TLS Mutual Authentication, where the credentials are stored in a x509 certificate unique to each device. This authentication scheme is based on asymmetric public/private key-pair cryptography and is extremely hard to forge. Modern Linux-based embedded devices are powerful enough to support this type of authentication.

COLLECT

There are many reasons to create a connected product. Centralised storage of device and sensor data is often high up on the list. By having data from all devices in a central location one can apply big data algorithms, teach artificial intelligence about normal behaviour, and get an instant overview of the device fleet status. It can save your company money, e.g by not having to dispatch maintenance crew unnecessarily; add efficiency to your customers, e.g by understanding device behaviour on a large scale and adjust the software to utilise resources better; and enable new business models, e.g by having customers subscribe to data streams and usage reports.

In order to facility these upsides, there are 3 choices to be made before you can collect device data; selection of communication protocols, communication server components, and device data storage components.

Communication protocol
There are several standard and open protocols to help devices communicate with an IoT cloud solution. Common ones are HTTP, MQTT, and CoAP, where MQTT and CoAP are specifically designed for communication with embedded devices. Evaluate the protocols from the view of your products' needs.

Considering the SECURE verb, the selected protocol should support encryption.

Communication server
The server should implement the protocol selected in the previous step and, considering the SCALE verb, the server should also support horizontal scaling. This often means that clustering capabilities are needed.

Considering the PROVISION verb, the server must support the authentication method defined by the provision component.

Device data storage component
Device data is often sensor or event data; numerical values that are informative only together with a timestamp. This is called Time Series data and it can be stored in relational databases or in files, but is more effectively stored in databases specifically designed for this type of data. Choose a database with Time Series data support.

Considering the SCALE verb, the storage component should also scale horizontally with clustering or sharding functionality.

[5] https://en.wikipedia.org/wiki/Mirai_(malware)

MONITOR

Your full IoT cloud solution will consist of many components, some of which will have several replicas. By having a monitoring solution that collect metrics and components status regularly, unexpected service downtime can be reduced by catching failures earlier. An effective monitoring solution will have a rules engine which act on collected metrics and device data. It will notify you of anomalies or automatically act on cloud metrics, such as high CPU load or excessive traffic. A common use of a rules engine is to manage the available computing resources in the cloud; automatically increase or reduce capacity when needed.

VISUALISE

A lot of data will flow through your IoT cloud solution and a great way to get an snapshot of the status of your devices and cloud components is to visualise them using graphs. Time Series data is especially suited to be displayed as graphs. By having a visualisation component in your IoT cloud solution from the start, your developers will have greater insight into how their software affect the data flow. When your IoT solution is up and running in production mode, the visualisation component will help you stay on top of developing issues early on.

The visualisation component is often a web application or part of your own application. Find a visualisation component that work well with your device data storage component.

UPGRADE

Some time could pass between manufacturing your devices and your customers receiving them. During this period many things could happen which might render your devices vulnerable or obsolete. Make sure you retain the ability to update your devices' firmware after they are deployed. By having a Firmware Over The Air (FOTA) component, additional features can be added later and security issues can be patched with little effort. If the FOTA solution is available early on in application development, it can be part of the development process to streamline quality assurance.

EXPAND

No IoT cloud solution is complete without customisation. That is what differentiates an IoT cloud platform from an IoT cloud solution; your application specific components. Here is where you should focus your development and the IoT cloud platform is there for support.

Application

Common components include web services that feed data to web applications and/or mobile apps. A web service is the glue between your application databases and the applications themselves and it enforces correct business logic for your applications.

Web services and web applications are types of components which are easy to scale horizontally and your Cloud Orchestration component should handle this.

Considering the SECURE verb, web services and web applications are exposed to the internet and thus all communication with them should be encrypted.

Storage

An application solution normally requires storage for different data types; relational data, key/value data, and unstructured data, and there are database components optimised for each type. Depending on your application, you will have a mix of these data types and should consider suitable storage components for each.

Considering the SCALE verb, the storage components should scale horizontally using clustering or sharding functionality.

Additional

Your full application solution could include additional components such as data import or export components that transfer data to and from other cloud services. Also, data analytics components

such as big data algorithms are common. There could be Open Source components matching your needs and it's worth the time finding a component, rather than developing a proprietary one. Realistically, your IoT cloud solution will be expanded with a mix of open source and proprietary components.

BACKUP

The device data is often the most valuable asset in your IoT solution but customer information is also a big valuable asset. Choose a backup solution early on to make sure you never loose device and customer data. Backup solutions are hard to generalise since they depend on how the data is stored and distributed. Most cloud providers offer backup solutions and those are most likely the best choice for your IoT cloud solution.

Focus on your business

IoT cloud solutions have many moving parts that are not directly related to your devices' primary function but are necessary to support a successful and scalable solution.

Unless cloud infrastructure is your core business, chose an IoT platform provider that allow you to get started with application development right away and does not lock you into their solution without also providing a way for you to retain control. They should present solutions for every one of "9 verbs of IoT Cloud Platforms" to effectively be a base for your IoT cloud solution.

Endian Technologies provides an Open Source based IoT cloud platform called Endian SCORE which allows you to delay the decision to purchase or build your own platform. You take full control over the Endian SCORE based IoT cloud solution when you are ready. Let us focus on your IoT cloud architecture while you focus on building your business.

Reach out to score@endian.se to find out what your IoT cloud solution will look like.