

Designing MQTT Topics for AWS IoT Core

October 2018



© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Contents

Introduction	1
MQTT Communication Patterns	1
Point-to-Point	1
Broadcast	2
Fan-In	3
Communication Workflows	4
MQTT Design Best Practices	5
General Best Practices	5
Best Practices for Telemetry	7
Best Practices for Commands	9
Applications on AWS	15
AWS IoT Shadow Example	15
MQTT Command Topics Example	17
MQTT Telemetry Topics Example	21
Best Practices for Using MQTT Topics in the AWS IoT Rules Engine	23
Conclusion	27
Contributors	28
Further Reading	28
Appendix: MQTT Topic Best Practices Checklist	29

Abstract

This whitepaper focuses on best practices for designing MQTT topics in Amazon Web Services (AWS) Internet of Things (IoT). It covers how developing an optimal MQTT topic schema can improve the overall architecture of your IoT applications by providing greater visibility into cloud to device communication, providing more fine-grained security permissions, and enhancing integration options with other AWS IoT services, such as the AWS IoT Rules Engine, AWS IoT Device Shadow, AWS IoT Device Management, and AWS IoT Analytics. This whitepaper is intended for technical architects, IoT cloud engineers, and embedded engineers. This paper assumes the reader has an understanding of fundamental MQTT concepts and terminology.

Introduction

AWS IoT Core supports MQTT, a widely adopted lightweight messaging protocol designed for constrained devices. MQTT participants receive information organized through MQTT topics. An MQTT topic acts as a matching mechanism between publishers and subscribers. Conceptually, an MQTT topic behaves like an ephemeral notification channel.

For AWS IoT, one of the first considerations when using MQTT is the design strategy of your MQTT topics. MQTT topics must balance current device communications, cloud side operations, and future device capabilities. Therefore, it can be challenging to design an ideal MQTT topic structure that creates enough of a schema to enforce least privilege communication but does not create a rigid structure that makes it challenging to support future device deployments.

This document provides you with MQTT topic design best practices and guidance. It outlines a set of commonly used MQTT topic structures that can be implemented to solve various device message patterns, then applies several example design patterns using different AWS IoT services.

MQTT Communication Patterns

IoT applications support multiple communication scenarios, such as device to device, device to cloud, cloud to device, and device to or from users. Although the range of patterns can significantly vary, a majority of MQTT communication models derive from three MQTT patterns: point-to-point, broadcast, and fan-in.

Point-to-Point

A point-to-point communication pattern is one of the basic building blocks of how devices commonly send and receive messages in MQTT. Two things use a single MQTT topic as the communication channel. The thing that receives the event subscribes to an MQTT topic. The thing that sends the message publishes to the same known MQTT topic. This approach is common in smart home scenarios where an end user receives updates about the thing in the home. In the following example, the robotic arm publishes a message on a topic that the mobile application subscribes to.




Figure 1: One to One Messaging in Point-to-point Communication

Point-to-point communication is not limited to one to one communication between devices. Point-to-point is also used in one-to-many communication where a single publisher can publish to individual devices using a different MQTT topic per device. This approach is common in notification scenarios where an administrator sends distinct updates to specific devices. In the following example, the repair service uses a set of point-to-point communications to programmatically loop through a list of appliances and publish a message.

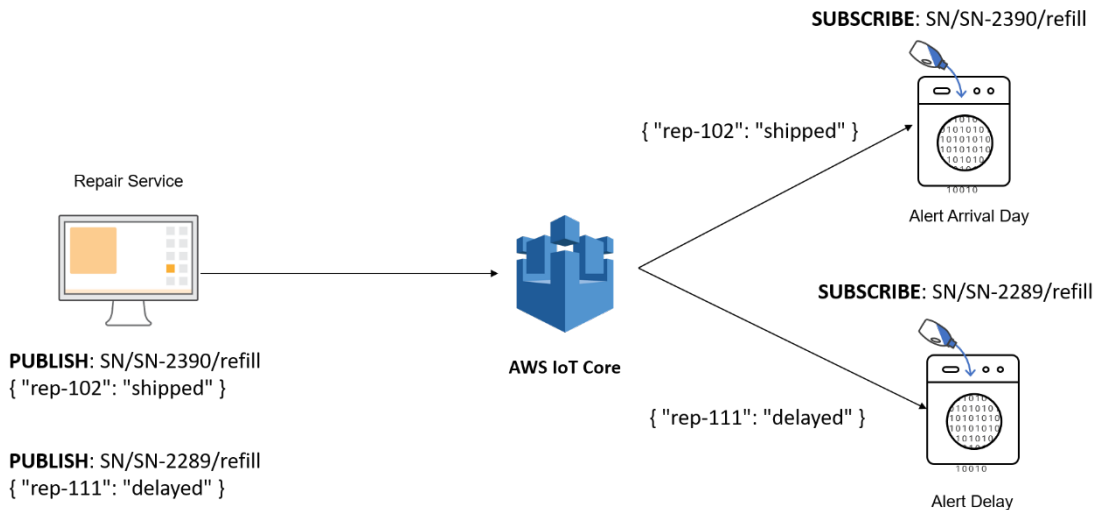


Figure 2: One-to-many messaging in point-to-point communication

Broadcast

Broadcast patterns are used for one-to-many messaging. The broadcast pattern sends the same message to a large fleet of devices. In a broadcast, multiple devices subscribe to the same MQTT topic, and the sender publishes a message to that shared topic. A typical use of a broadcast pattern is to send a notification to devices based on the category or group of the device. For example, a weather

station transmits a broadcast message based on a device's current geolocation, where devices are grouped by location.

The following illustration depicts an example where a broadcast pattern sends a message on a weather topic that all buses in the state subscribe to. Based on the current location of the bus, it can ignore or respond to the message.

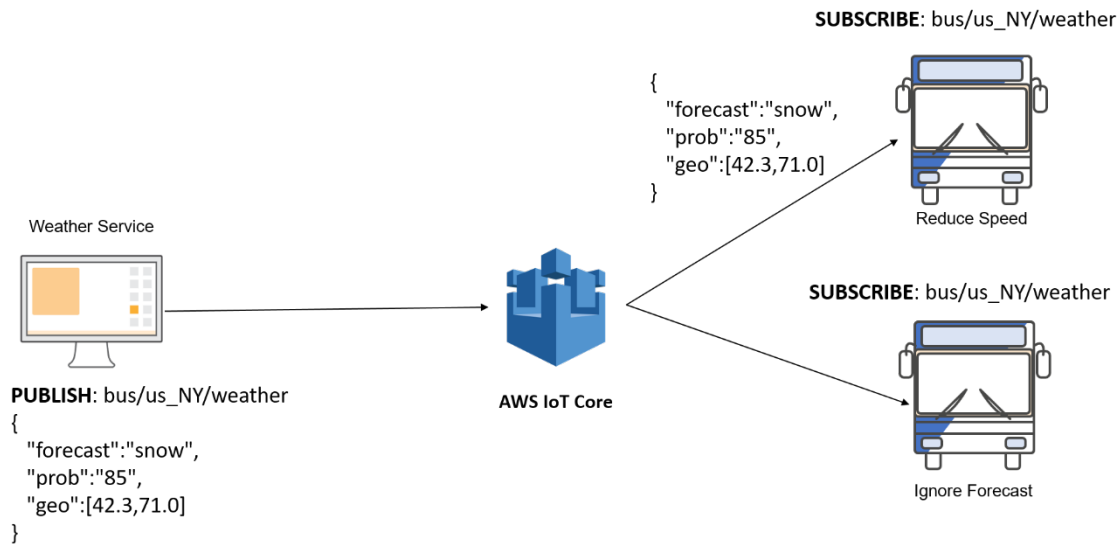


Figure 3: One-to-many messaging in broadcast communication

Fan-In

The fan-in pattern is a many-to-one communication pattern and can be thought of as the reverse of the broadcast pattern. Multiple devices publish on a shared topic with a single subscriber to that topic. With the fan-in pattern, the subscriber has the added ability to leverage wildcards as the publishers all use a similar but unique MQTT topic. The fan-in pattern is commonly used to facilitate aggregation of telemetry to an IoT application.

In the following example, each end device publishes to an MQTT topic containing a known group identifier. The AWS IoT Rules Engine uses a wildcard subscription to receive the messages and route them to an [Amazon Kinesis](#) stream. Specifically, the robotic arms publish on a fan-in topic associated with a specific building (building1234). The administrative system receives all updates for the building using an MQTT wildcard (+).

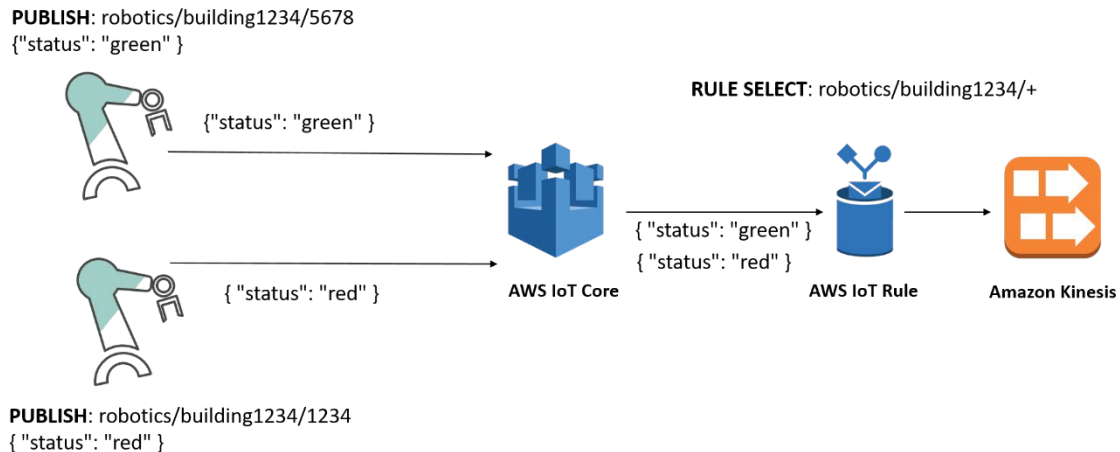


Figure 4: Many to one communication in fan-in pattern

When devices communicate via the cloud using MQTT, avoid using the fan-in pattern to a single subscribing end device, because this routing may hit a non-adjustable limit on a single device MQTT connection. Instead, use the fan-in pattern to route a large fleet of messages to your IoT application via the AWS IoT Rules Engine. For large scale fan-in scenarios, combine the Rules Engine with a wildcard subscription pattern and a Rules Engine action to route to Amazon Kinesis Data Streams, Amazon Kinesis Data Firehose, or [Amazon Simple Queue Service \(Amazon SQS\)](#).

Communication Workflows

The three common communication workflows are device to device, device to cloud, and cloud to device. Each workflow determines the order of your MQTT topics. In the case of device to device, MQTT topics should contain identifiers for either the sender or receiver of a message. For device to cloud, MQTT messages should include information about the target application. The target application is responsible for augmenting any MQTT messages with internal metadata about the device. Last, for cloud to device communication, MQTT messages should contain session information for tracking acknowledgment of any critical messages.

MQTT Design Best Practices

General Best Practices

Although there are numerous combinations of IoT communication patterns that share common approaches, there are several best practices that apply to any message pattern irrespective of how a device is publishing or receiving a message. This section articulates several overall best practices for you to review and implement as you design your topic structures.

Ensure MQTT topic levels only use lowercase letters, numbers, and dashes. Since MQTT topics are case sensitive, it is important to use a standard set of naming conventions when designing MQTT topics. For this reason, customers should only use lowercase letters, numbers, and dashes when creating each topic level. Customers should avoid camel casing and using hard to debug characters such as spaces.

Ensure MQTT topic levels structure follows a general to specific pattern. As topic scheme flows left to right, the topic levels flow general to specific. For example, an HVAC system is associated with an IoT platform named **hv100**, is located in the **basement** of building **bld1518**, and has a Thing Name of **hvac719**. The topic structure begins with the general group, in this case, the name of the IoT platform, and ends with the most specific identity, the Thing Name. This example creates the following topic level structure:



Include any relevant routing information in the MQTT topic. Relevant routing information includes, but is not limited to, the IoT application identifier, any groups the device may be a part of, such as installed location, and the unique identity of your IoT device. To continue with the previous HVAC system example, the MQTT topic `hv100/bld1518/basement/hvac719` includes all relevant routing information. Based on this MQTT topic, you can design a system that captures any data related to the entire application using the identifier, `hv100`, but also can target different areas of interest for subscribing to messages, such as the building location.

Prefix your MQTT topics to distinguish data topics from command topics. Make sure that your MQTT topics do not overlap between commands and data messages. By reserving the first topic level to denote data and command topics, you are more easily able to create fine-grained permissions using IoT policies, and monitor the status of commands and command responses separately from passive telemetry commands. For example, use the AWS IoT Device Shadow for tracking reported and desired states, and use a separate data topic for passive, real-time telemetry data.

Document proposed MQTT topic structures as part of your operations practice. The document should include all topics available for publishing, subscribing, or receiving data, along with the intended producers and consumers of the data. Review the document to ensure it adheres to any AWS IoT limits, internal security requirements, and any application use cases.

Use the AWS IoT Thing Name as the MQTT client ID for connecting as a device over MQTT. MQTT requires a unique client ID when connecting to AWS IoT. Use the unique Thing Name per device as the MQTT client ID. By matching the client ID to the Thing Name, you are more easily able to correlate information, such as IoT logs, to the associated AWS IoT devices. Additionally, if the clientId matches the Thing Name, you can use additional features such as [Thing Policy Variables](#). Thing Policy Variables enables AWS IoT policies to include thing properties, such as thing name and thing attributes, and can help reduce the total number of policies you need to provision across your IoT devices.

Include the Thing Name of the device in any MQTT topic the device uses for publishing or subscribing to its data. To track messages destined for a particular device, include the Thing Name as part of any MQTT message that is published by the device or sent to a specific device. The Thing Name should appear near or at the end of the MQTT topic after any routing topic information.

Review the AWS IoT Core default service limits. Design your communication pattern so that it aligns with any adjustable IoT [service limit](#). AWS IoT has several adjustable and non-adjustable limits associated with using the AWS IoT core service. As part of your topic review, review the AWS IoT limits, and ensure your MQTT topic and device communication do not conflict with any non-adjustable limits.

Include additional contextual information about a specific message in the payload of the MQTT message. This contextual information includes, but is not limited to, a session identifier, the requestor identifier, logging information, or the return topic on which a device is expecting to receive a response. Although the [MQTT 3.1.1 specification](#) does not require specific payload attributes, we recommend you include relevant tracking information inside of the MQTT payload. By creating a standard structure including fields such as session identifier and success or error codes, you can more easily analyze trends in device behavior. Standardizing the communication schema also strengthens a shared vernacular of device use cases across IoT teams.

Avoid MQTT communication patterns that result in a sizeable fan-in scenario to a single device. Some AWS IoT limits cannot be raised as part of a limit increase and frequently correlate to per device actions, such as maximum publish-in on a single MQTT connection. Do not allow a single device to subscribe to a shared topic that is being published to by a large number of other devices. By avoiding this pattern, you are more likely to avoid hitting a single connection device limit, particularly a throughput per connection per second limit.

Never allow a device to subscribe to all topics using #, and only use multi-level wildcard subscriptions in IoT rules. By using multi-level wildcards, you can create unintended consequences when you inadvertently add new topics to the hierarchy that may not be intended for that particular device. Instead, reserve use of multi-level wildcards as part of the IoT rules engine, and use single level wildcards (+) for device subscriptions.

Best Practices for Telemetry

Telemetry is read-only data that is transmitted by the device and aggregated in the cloud. Telemetry follows the device to cloud pattern along with the fan-in pattern for communication. Telemetry does not require an acknowledge message back from the MQTT broker beyond optionally setting a higher quality of service (QoS) level. Since telemetry is a passive activity, the MQTT topic for telemetry cannot overlap with any MQTT topics for active workflows, such as commands. A telemetry topic supports more complex devices that publish telemetry on behalf of other devices, such as an edge gateway or a mesh network with a single coordinator.

In AWS IoT, you have the ability to use different AWS IoT services to support telemetry communication patterns. We recommend that you use standard MQTT topics to support your telemetry use cases.

Using the MQTT Topics for Telemetry

MQTT Telemetry Topic Syntax

The following example and sections provide the MQTT topic structure for telemetry:

```
dt/<application>/<context>/<thing-name>/<dt-type>
```

dt: Set prefix that refers to the type of message.

For a telemetry topic, we use `dt`, short for data. All telemetry topics use this top-level prefix for an application. By reusing the same value for telemetry, you can identify the intent of a message by referring to the initial prefixed value. In this case, any `dt` topic is a telemetry topic.

application: Identifies the overall IoT application associated with the device.

Commonly used application attributes include device hardware version or an internal identifier for a cloud application that is the primary ingestion point for a message. The IoT application is associated with an internal name for your overarching IoT product or relates specifically to the type of hardware of your device. Because the application topic portion correlates to a group of device messages and is immutable, the application prefix portion of the telemetry MQTT topic is placed immediately after the `dt` message type.

context: Single or multiple levels of additional contextual data about the message a device is publishing.

Contextual information is related to information that is set during device provisioning. For example, contextual information in a factory setting could include the current physical location of a device in the facility. Another example of contextual information is a `group-id` in the MQTT topic. The `group-id` denotes when multiple devices have an inherent relationship based on specific attributes, such as buying a package of smart light bulbs to control lighting a

room. The group-id enables numerous devices to coordinate activities as a single unit.

thing-name: Identifies which device is transmitting a telemetry message.

dt-type (optional): Associates a message with a particular subcomponent of a device, or for edge gateways, any downstream devices.

A complex device often has multiple subcomponents with specific tasks, such as sensors, actuators, or separate system on chips (SOCs). The `dt-type` allows you to associate each subcomponent of a particular device to an individual MQTT topic. For example, a subcomponent that measures geolocation and direction of a vehicle. That subcomponent would have a `dt-type` value of `geo` to distinguish its geolocation messages from other components of the car, such as the accelerometer.

Best Practices for Commands

In IoT applications, command topics are used to control a device remotely and to acknowledge successful command executions. Unlike telemetry, command topics are not read-only. Commands are a back and forth workflow that can occur between two devices or between the cloud and devices. Because commands are actionable messages, isolate the MQTT topic for command messages from telemetry topics.

Several services are available for you to implement command and control operations on AWS IoT. With the capability to store the desired and reported states in the cloud, the AWS IoT Shadow is the preferred AWS IoT Service for implementing individual device commands. AWS IoT Device Jobs should be used for fleet-wide operations as it provides extra benefits, such as Amazon CloudWatch metrics for Job tracking, and the ability to track multiple in-transit Jobs for a single device. You can use a combination of the AWS IoT Shadow, AWS IoT Job documents, and standard MQTT topics to support your command use cases.

Using the AWS IoT Shadow for Commands

[The Device Shadow](#) service acts as a state intermediary, allowing devices and applications to retrieve and update a device's shadow state. You can use the shadow to get and set the state of a device over MQTT or HTTP. The shadow

includes the following individual state properties that support command and control:

desired state. Applications that have permissions to send commands to a device can write the requested state changes to the desired portion of the shadow document. By updating the desired state, the AWS IoT Shadow service stores the desired state change in the AWS cloud and then sends an MQTT message to the device using a reserved shadow topic. When a device receives a shadow request, it can execute the changes required from the desired state.

reported state. The reported state of the thing stores the current attributes last published by a device. Things write to this portion of the document to record their new state while applications read this portion of the Shadow to determine the state of a specific device. For implementing command topics, the AWS IoT Shadow is the recommended feature for storing, retrieving, and changing attributes of a device.

Use the AWS IoT Shadow in situations where a command persists for later use, even if the device is currently offline. For example, if a GPS system is sent a new destination through the Shadow desired state but is not immediately reachable, the new coordinates remain in the GPS IoT Shadow. Once the GPS system regains connectivity, it can actively request its last shadow state and retrieve the new coordinates. The shadow is also ideal for storing the last reported state for attributes of the device.

Best Practices for Using the AWS IoT Shadow

The AWS IoT Shadow is an excellent mechanism for command and control along with storing the reported state of a specific device. The following list provides best practices for maximizing the efficiency of commands through the shadow.

IoT devices should not share shadows. To separate commands for each device, make sure that each device has permissions to its own shadow and that devices do not share a single shadow. For complex scenarios, like edge gateways or large device assets with multiple subcomponents, the primary asset should use multiple IoT Things and shadows individually associated with the downstream devices.

Use the shadow for state or commands that have a medium to low transaction per second (TPS). The shadow is an ideal fit for infrequent updates that occur in minutes, hours, or days. Because the shadow publishes on additional topics to acknowledge an action was successful, use a standard MQTT command topic for high throughput telemetry, such as requests that are sending updates multiple times per second or every second.

Use the shadow for storing status metrics of a device. Store informational data about the current health of the device including, but not limited to, connectivity, the status of device sensors and control units, and any error information about those subcomponents. If you know the current status of the device, you can make actionable decisions during command requests.

Use the device shadow to catalog the current firmware version. The shadow is an ideal location for a device to report the installed firmware version on the hardware. The firmware should be a simple attribute, such as a field that highlights `major.minor.patch` version of the service.

Use the optional `clientToken` field with device shadow updates to track the sender of a shadow message. The `clientToken` is a field in the Shadow that enables a subscriber to associate the responses with requests in your MQTT application. If a device sets the `clientToken` during a shadow update request, the AWS IoT Shadow service includes that same `clientToken` in the associated Shadow output events.

Using the AWS IoT Jobs for Commands

AWS IoT Jobs is a service that allows you to define a set of remote operations that are sent to and executed on one or more things connected to AWS IoT. For command use cases, Jobs allows applications to run tasks that require executing multiple steps. An AWS IoT Job contains instructions that the thing must run to complete its transaction. AWS IoT Jobs are the recommended feature for fleet-wide operational tasks, such as software updates, that are only executed by trusted administrators of the entire IoT application.

Best Practices for Using the AWS IoT Jobs

Use thing groups to organize devices for AWS IoT Jobs. Create multiple thing groups organized by common device attributes, such as the current firmware version, hardware version, or deployment environments (for example, staging or production). Thing groups should also have common hierarchical

structures, such as business units or locations. During deployments, you can use thing groups as the deployment target for a specific IoT Job.

Use staged rollouts to deploy commands using Device Jobs. Device Jobs are the ideal solution for delivering fleet-wide operations to devices. Create multiple smaller deployments first to subsets of the fleet, letting the devices apply your changes, and then rolling out the commands to a greater number of devices over time. By allowing changes to progress over weeks and months, you can have more confidence that there are fewer unforeseen issues and you can more quickly react if there is an issue during an earlier rollout.

Using the MQTT Topics for Commands

MQTT Command Topic Syntax

In some scenarios, you may want to design your command communication using the standard MQTT publish and subscribe model. These types of situations may occur when a device must execute a command that is temporal (that is, can only be processed at this current type and should fail if the device is unavailable) or run a single command across multiple devices simultaneously. It is also possible in a brownfield environment where a device may be incapable of leveraging higher level AWS IoT services. You may also require the flexibility to choose your own set of MQTT topics to define commands to and responses from devices.

In cases where you are using a separate set of command topics, follow similar best practices for MQTT command topics as described for telemetry. A command topic should have flexibility for complex devices that publish or relay commands to other devices. Command topics should also provide visibility into essential attributes. MQTT command topics should be designed in a way that can answer operational questions based on the MQTT topic and payload:

- Who is the originator of the command?
- Who is the intended receiver of the command?
- Was the command processed successfully?
- What is the current status of the command?
- If the command was not processed successfully, what is the error?

In addition to these questions, you may also want to determine when a command was requested, when a device responded, and to monitor the state of any single request among the fleet in the cloud.

When you design MQTT topics for command requests, follow this structure:

```
cmd/<application>/<context>/<destination-id>/<req-type>
```

Since commands are two-way communication patterns, design a similar MQTT topic structure for responding to commands, such as the following:

```
cmd/<application>/<context>/<destination-id>/<res-type>
```

Because telemetry topic design is similar to command topic design, this section provides only the portions of the IoT topic for command requests and responses that differ.

cmd: Prefix that refers to the type of message.

Command topics use `cmd`, short for command. By prefixing all commands with `cmd` and all telemetry with `dt`, telemetry and commands are isolated on separate MQTT topics.

req-type: Classifies the command.

For simple request and response patterns, the `req-type` attribute should be a single command request static value such as `req`. In cases of limited command types, the MQTT message includes the additional data in the payload.

In more complex systems, where a device is orchestrating multiple devices, actuators, or subcomponents, the `req-type` attribute relates to each subcomponent available to receive commands. For example, if a device is mobile, you may want to steer the device remotely or receive navigational information about the device's surroundings. This type of subcomponent would have a `req-type` of `nav` where commands are sent steering in single or multiple planes.

destination-id: Identifies the destination device or application for this message.

By including the `destination-id`, the target device can subscribe to its own set of command topics and receive any command requests.

res-type: Denotes command responses and identifies responses that are related to a previously sent command.

The `res-type` enables a single device to use one single-level wildcard subscription for all incoming command acknowledgments. If device has limited commands, the response topic can use a static field, such as `res`.

MQTT Command Payload Syntax

In addition to creating a clear MQTT topic structure for commands, make sure you generate a schema for message payloads. MQTT payload information is parsed by the receiving device or IoT application to inform it of any additional logic it may need to complete its operation. For MQTT commands, include the following fields with the command message payload:

session-id: Identifies a unique session.

The requestor generates the `session-id` for the command and includes it in the request payload. The response topic uses the `session-id` upon command completion. By using a `session-id`, the AWS IoT Rules Engine can store and track the status of commands and determine if a request is still in transit, successful, or in error. Devices can also keep track of in-transit requests when communicating with multiple devices.

response-topic: In a command, there is a request for an action to happen and then a response that indicates the status of the command (successful or error). To avoid hard-coding response topics, we recommend that for any MQTT command, the command request payload includes a field that has a response topic. The device publishes its response payload using the response topic. For example, consider the following command topic:

```
cmd/security/device-1/cert-rotation
```

In the payload of this request, the IoT application includes a field that denotes where the device (`device-1`) should send its response and a session identifier for tracking. See the following example for this command's payload structure:

```
{
  "sessionId":"session-820923084792",
  "resTopic":"cmd/security/app1/res"
}
```

Applications on AWS

The following sections provide use cases for implementing MQTT topic best practices using AWS IoT.

AWS IoT Shadow Example

In this example, an administrator is sending a command to a wind turbine to change the current angle of the turbine blade to accommodate an approaching wind change. To do so, the administrator first leverages the AWS IoT Shadow as the primary mechanism for issuing commands to a device and as a storage location for essential device readings related to the device.

This scenario assumes the following details:

- The wind turbine as a thing name of `turbine8000`
- The administrator is leveraging an internal application which he is authorized to use with temporary IAM permissions scoped to `turbine8000`
- The administrator has a session identifier of `session-admin100`

Command Request to Wind Turbine Device Shadow

The administrator publishes a message to the IoT Shadow using the reserved MQTT Shadow topic for `turbine8000`. The command payload includes setting the desired state of the shadow. For tracking purposes, the administrator also embeds a session identifier as the `clientToken` in the IoT Shadow request.

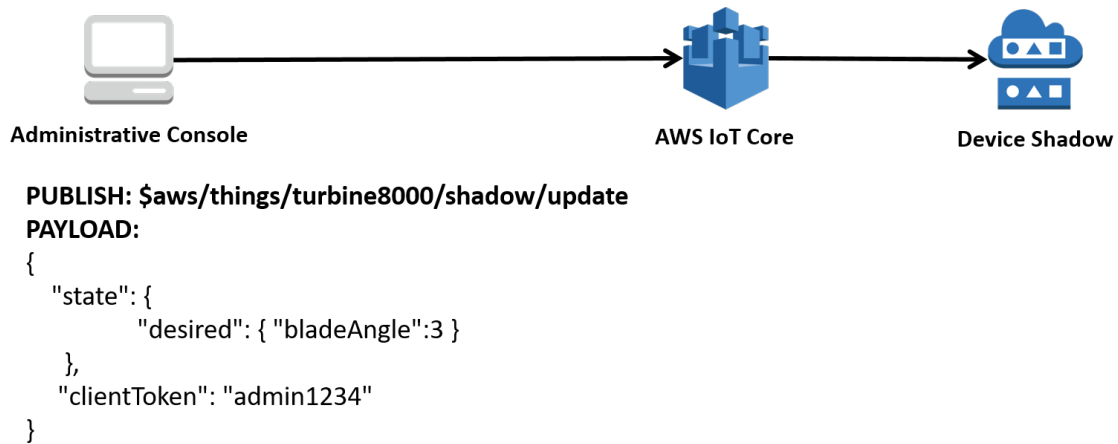


Figure 5: Command request to wind turbine device shadow

Command Processing On Wind Turbine

When using the shadow, the end device is responsible for subscribing to several different MQTT topics to receive the appropriate notifications from the AWS IoT Shadow Service. At a minimum, make sure devices subscribe to the `update/delta`, `update/rejected`, `get/accepted`, and `get/rejected` reserved Shadow Topics to receive successful and unsuccessful shadow requests.

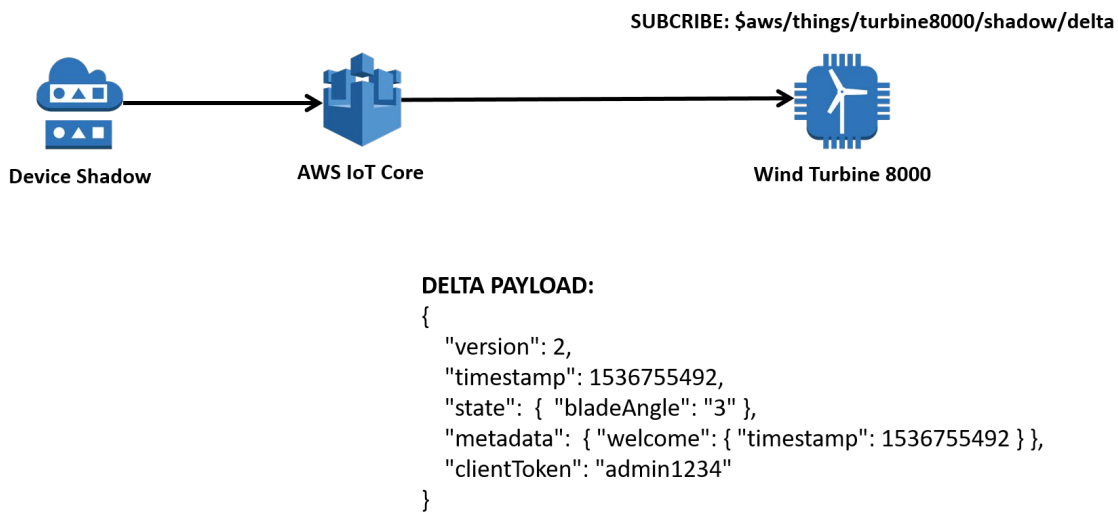


Figure 6: Command processing on wind turbine

In the example, the wind turbine receives a message on `update/delta` that the reported state is different than the desired state. Before processing the command, the wind turbine checks the `clientToken` field to determine the

sender of the request. Then, it inspects the timestamp field in the metadata. (For some use cases, commands should only be deemed valid for a given period of time.) Once confirmed, the wind turbine adjusts its blade angle and publishes a response with a new reported state using the `update` topic.

Command Response Delivered to Administrator

For the administrator to receive updates on reported state changes, the administrative UI subscribes to the `update/accepted` and `update/rejected` topics. Once the turbine has published its updated state to the IoT Shadow, the administrator receives an acknowledgment of the updated messages on the `update/accepted` topic.

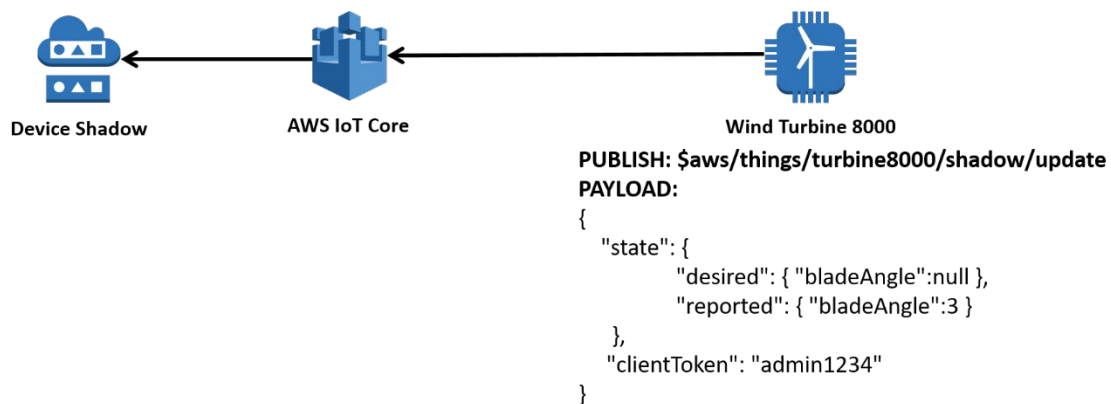


Figure 7: Wind turbine publishes update to device shadow

MQTT Command Topics Example

For a smart door lock application, a user must be able to submit a command to the lock that triggers a temporary key to be issued for an approved visitor. The temporary key consists of a TTL, code, and information about the authorized user. The ability to create a temporary key allows another individual to open the lock for a specified period. This use case would apply in scenarios such as visiting family member arriving at the home while the primary owner is at work.

This scenario assumes the following details:

- The homeowner has a mobile device id of `mobile-1`
- The approved visitor has a mobile device id of `mobile-2`

- The smart lock is installed as part of a group of other locks in the home. The set of locks has a context for `groupId` where `groupId` equals `group-3`
- The smart lock for the front door has a `lockId` of `lock-1`
- The smart lock hardware has a series number of `series100`. The series is a unique identifier for this product version.

Command Request to Generate a Smart Lock Code

The primary owner first publishes a command to the lock to request a temporary access code created for an approved visitor. The command payload includes the identification of the homeowner's mobile device, a randomly generated session identifier for tracking the current request, an action field that contains the type of command, and a topic field. The smart lock uses the topic in the topic field for publishing its response back to the homeowner.

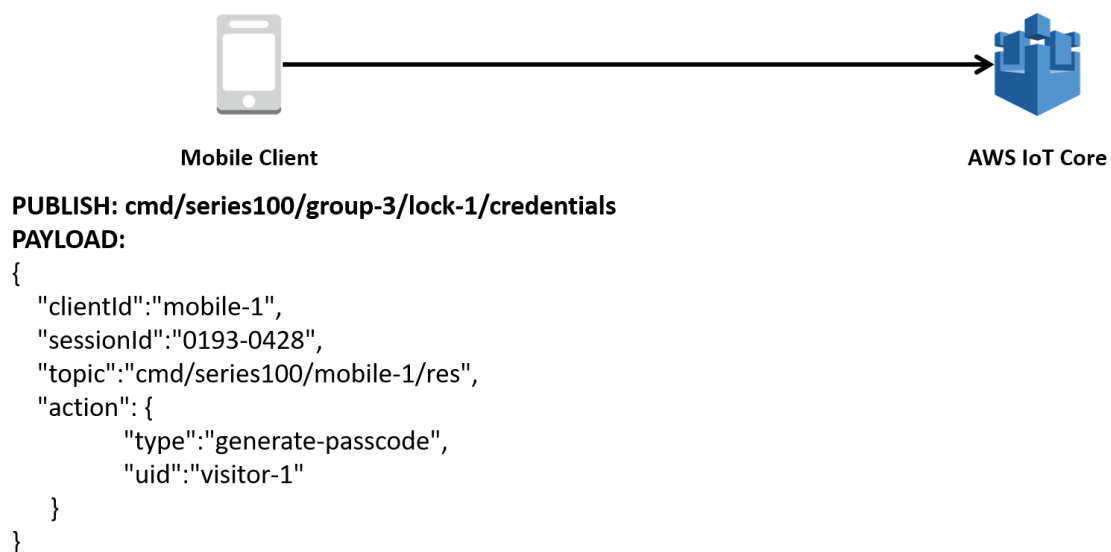


Figure 8: Mobile user requests temporary credentials for the front door

Augment MQTT messages using any internal standards for requests, responses, and telemetry. For example, by adding in the requestor of the command in the payload, applications can specify different response topics based on the use case. If the homeowner requests a temporary lock but needs the response to reach all door locks in the home, the topic field could be changed to send to a group of devices.

Command Processing on the Smart Lock

The smart lock receives the command message from the MQTT topic. By leveraging a consistent naming schema for commands in this application, the smart lock can ensure that it only receives commands on its specific command topic. This topic design also makes it possible for the lock to subscribe using a single level MQTT wildcard after its identifier. The single level wildcard command is backward compatible as the IoT application adds new command types.

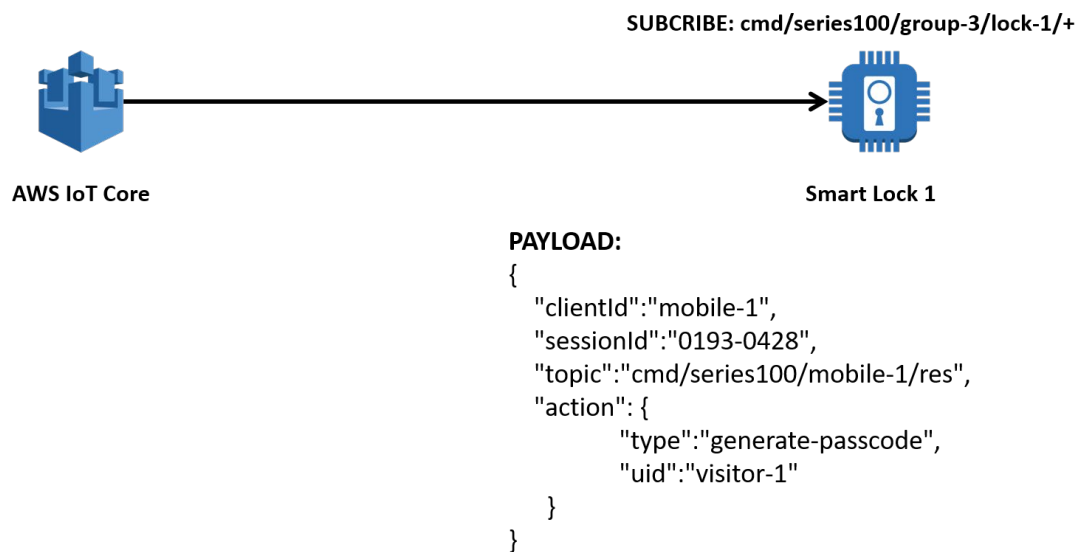


Figure 9: Command sent from the AWS IoT Device Gateway to the smart lock

After receiving the MQTT payload, the smart lock parses the command request to determine what type of action it runs. In this case, the command is related to credentials. The device also extracts the client ID, the session ID, and the response topic from the command payload. Because a home can have multiple authorized homeowners, the client ID determines which homeowner has requested this change. In this example, the action field includes the type of credentials request, generate-password, and the associated user for the temporary key. Last, the device obtains the response topic field in the MQTT message and uses this information to publish its response.

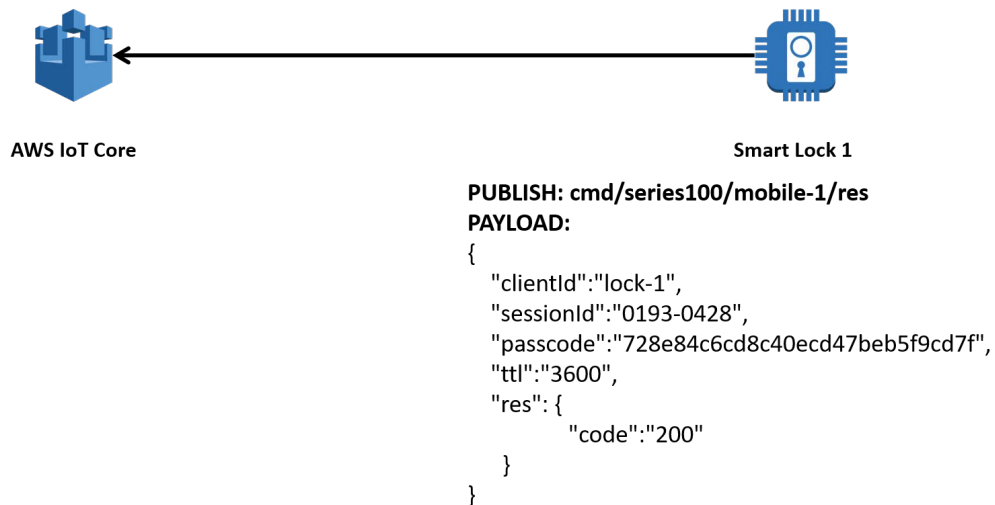


Figure 10: Smart lock response published to AWS IoT Core

Command Response Delivered to the Mobile Client

The homeowner's mobile client subscribes to command responses for any smart locks in the group. Whenever the mobile client receives a successful command response, the temporary code along with any additional authorization permissions can be processed on the device and simultaneously stored in the AWS cloud. At a later time, the authorized visitor can use the temporary code along with additional security credentials, such as exchanging for OAuth credentials, proving local presence to the door, and so on, to apply the temporary code to the smart lock.

In this workflow, the application used a command topic for a single device. However, a similar workflow may be used to request commands for multiple devices in a group, such as locking all of the doors in the home.

SUBSCRIBE: cmd/series100/mobile-1/+



Mobile Client



AWS IoT Core

PAYLOAD:

```
{
  "clientId":"lock-1",
  "sessionId":"0193-0428",
  "passcode":"728e84c6cd8c40ecd47beb5f9cd7f",
  "ttl":"3600",
  "res": {
    "code":"200"
  }
}
```

Figure 11: Command response sent from AWS IoT Core to mobile client

MQTT Telemetry Topics Example

This section is an example of aggregating telemetry from a set of occupancy sensors that are placed throughout a building to monitor room usage. The occupancy sensors communicate to a local gateway that is running [AWS Greengrass](#). AWS Greengrass then delivers all sensor metrics on an MQTT topic to AWS IoT Core. Because this use case is focused on telemetry, a response topic is not needed between Greengrass and AWS IoT Core, either locally or upstream.

This scenario assumes the following details:

- The occupancy sensors have ids of `occupancy-1` and `occupancy-2`
- The building is called `building-fresco`
- Each sensor is placed on a specific floor and room name in `building-fresco`.
- The Greengrass local gateway has a unique identifier of `gateway-1`
- The current building automation system correlates to an internal project called `acme`

Local Telemetry from Occupancy Sensor to Greengrass

Each occupancy sensor publishes an occupancy reading once per minute and whenever a person enters or leaves the room. Because the occupancy sensors do not correlate precisely to the state of the device and instead refer to the state of the room, the sensor publishes the room status on a telemetry topic. The payload includes a timestamp, occupancy count, and any efficiency countdown for turning off lights if a room is vacant. The occupancy sensor uses an MQTT topic that includes the contextual information about the position of the sensor within the building and its associated project. A Greengrass core receives all occupancy sensor data locally.




Figure 12: Local occupancy sensor publishes sensor reading to AWS Greengrass

Greengrass Telemetry from Edge to Cloud

In this example, the primary role of Greengrass is to aggregate the data from multiple occupancy sensors then send the data to AWS IoT Core. Because Greengrass is the local bridge to the cloud, Greengrass adds metadata to each sensor reading. Greengrass adds building information to each message to show the overall usage of the building in 5 minute increments. Greengrass also augments the MQTT topic by including the appropriate application identifier, acme.

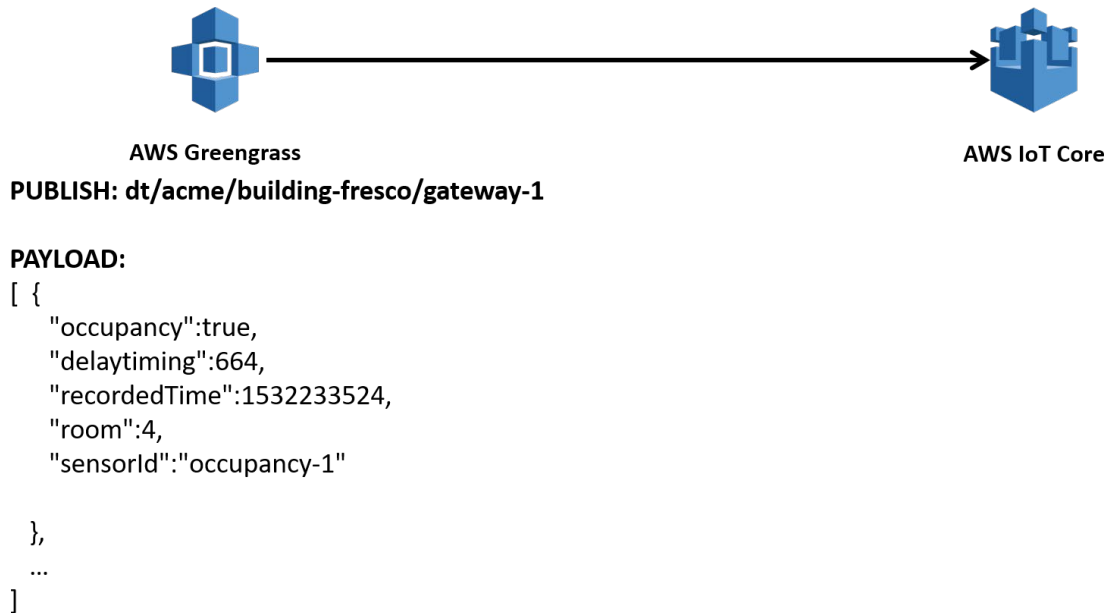


Figure 13: Greengrass aggregates and augments telemetry then forwards messages to AWS IoT Core

Best Practices for Using MQTT Topics in the AWS IoT Rules Engine

The [AWS IoT Rules Engine](#) enables you to define how messages sent to AWS IoT Core can interact with AWS services. An AWS IoT rule consists of a SQL SELECT statement, a topic filter, and a rule action. The SQL SELECT statement can extract data from incoming MQTT messages. The topic filter of an AWS IoT rule specifies which MQTT topics trigger an IoT Rule Action.

The rules engine plays a pivotal role in intelligently directing messages to other AWS services or republishing to devices. IoT rules support use cases, such as gathering operational metrics, data enrichment, data aggregation of device telemetry for analytics purpose, and for troubleshooting errors.

Rules Engine Integration with Telemetry Topics

We recommend using a topic structure for telemetry similar to the following:

```
dt/<application-prefix>/<context>/<thing-name>/<dt-type>
```

The second field in the MQTT topic telemetry pattern defined as `application-prefix` represents an immutable, natural bifurcation between your devices in a fleet. A common attribute for the application is the device hardware version or the name of the IoT application. Using the telemetry MQTT structure, you can create an IoT rule to capture all telemetry associated with a specific application version:

```
{
  "sql": "SELECT *, topic(2) as applicationVersion, topic(3) as
contextIdentifier FROM 'dt/#'",
  "awsIoTSqlVersion": "2016-03-23",
  "ruleDisabled": false,
  "actions": [{
    ...
  }]
}
```

Because the MQTT topic structure mirrors a hierarchy, this rule can select different parts of the MQTT topic hierarchy and inject it into the new payload. These attributes provide further context as messages are processed and stored in other AWS services.

Rules Engine Integration with Command Topics

The Rules Engine is an ideal service to capture insight into the success or failure of commands, regardless of whether the commands are sent using the AWS IoT Device Shadow, IoT Job, or by using an MQTT command topic.

Tracking Success of Commands

The AWS IoT Rules Engine can be used to track the success rates of individual commands. The IoT rule extrapolates payload information, such as session identifier; generates additional metadata in the rule select statement, such as creating a time to live; and temporarily stores the new message payload into a data store, such as [Amazon DynamoDB](#). The rule that follows mirrors a common implementation of this use case for an AWS customer. The IoT rule stores each session as an individual DynamoDB record and because the WHERE clause identifies this as an incoming command, the rule adds a literal value named `status` that marks the command as `In progress`.

```
{
  "sql": "SELECT session ID AS token, timestamp()/1000 as ttl,
  topic AS responseTopic, client ID AS requestorId, action.type AS
  commandType, 'In Progress' AS status FROM 'cmd/series100/# WHERE
  topic(5) == 'credentials' ",
  "awsIoTSqlVersion": "2016-03-23",
  "ruleDisabled": false,
  "actions": [{
    "dynamoDBv2": {
      "roleArn": "arn:aws:sns:us-east-1:12345678:sns_role",
      "putItem": {
        "tableName": "command_sessions_table"
      }
    }
  ]
}
```

As the IoT ecosystem receives command requests, the preceding rule maintains a record of all in-transit commands.

By following the MQTT topic best practices, the response topic includes overlapping information as the command itself (for example, the original session ID and the response topic used by the smart lock). As an added capability, the cloud application may have a second IoT rule that uses the session ID to update the status of a specific command using information from the response metadata.

See the following example:

```
{
  "sql": "SELECT sessionId AS token, timestamp()/1000 as ttl,
  topic() AS responseTopic, client ID AS requestorId, res.code AS
  responseCode, 'Complete' AS status FROM 'cmd/series100/# WHERE
  topic(5) == 'res' ",
  "awsIoTSqlVersion": "2016-03-23",
  "ruleDisabled": false,
  "actions": [{
    "dynamoDBv2": {
      "roleArn": "arn:aws:sns:us-east-1:12345678:sns_role",
      "putItem": {
        "tableName": "command_sessions_table"
      }
    }
  ]
}
```

Aligning Rules Engine Capabilities with MQTT Topics

As you define your use of the IoT rules, review the following recommendations as you relate to MQTT topics and the AWS IoT Rules Engine:

- Use the [topic\(Decimal\)](#) rule function to augment your MQTT messages with contextual information contained in your MQTT topics.
- Use the [timestamp\(\)](#) rule function to include a timestamp that correlates the time that a message reached AWS IoT Core.
- If your commands are in JSON, reference any contextual payload metadata, such as session ID, in the [SELECT](#) and [WHERE](#) clause of the rules engine. The additional payload information can be used to determine if and when a rule should trigger.
- Use [Substitution templates](#) in your AWS IoT actions to express variables as part of the AWS IoT Rule action that is triggered. Substitution expressions make it easier to scale and dynamically route to downstream IoT Rule actions.
- To track the completion of a command or request, use the AWS IoT Rules Engine to store the data and status in a service, such as DynamoDB. As messages are processed, data can be automatically

archived from DynamoDB using a TTL field. In cases where commands are sent at a high throughput rate, you can leverage the IoT Rules Engine with Amazon Kinesis to buffer data prior to DynamoDB storage.

- Use the AWS IoT Rules WHERE clause to filter messages that do not apply to an AWS IoT Action. The WHERE clause can be used with the JSON payload or Rules Engine Functions, such as [get_thing_shadow\(thingName, roleARN\)](#) or [aws_lambda\(functionArn, inputJson\)](#).
- After AWS IoT Core receives a message, use AWS services like Amazon Kinesis or Amazon SQS to buffer the message payload along with the MQTT topic the message was published to. Once messages are buffered, you can run your own logic on [AWS Lambda](#) or [Amazon Elastic Compute Cloud \(Amazon EC2\)](#) to map fields from the payload or the topic and enrich the payload with additional metadata related to the individual devices, the type of device, or the device group.

Conclusion

MQTT is a simple, secure, flexible, and robust IoT protocol. It allows you to define communication networks between devices and the cloud that can tailor fit an increasingly large number of customer use cases. To support you on your initial steps in using MQTT on AWS IoT, this whitepaper has presented several best practices, guidelines, and considerations that can be used when reviewing how to scale IoT device communications.

AWS IoT enables the definition of several MQTT communication patterns—point-to-point, broadcast, and fan-in—that relate to different use cases. In addition, AWS IoT Services provide you with additional managed services including, but not limited to, AWS IoT Jobs, AWS IoT Shadow, and the AWS IoT Rules Engine.

When scaling IoT solutions, use the best practices for defining telemetry and commands while ensuring that the overall MQTT design does not conflict with the overall best practices for MQTT outlined in this paper. Last, consider how your IoT topic structure impacts operational and business visibility in the cloud. To achieve strong business value, you must plan, develop, design, and use AWS IoT rules to fully leverage the IoT application's MQTT topic schema.

Contributors

The following individuals and organizations contributed to this document:

- Olawale Oladehin, Solutions Architect, AWS IoT

Further Reading

For additional information, see the following:

- [AWS Whitepapers](#)
- [AWS IoT Core Documentation](#)

Appendix: MQTT Topic Best Practices Checklist

Use the following list to quickly confirm that you have considered all of the AWS IoT recommended best practices for designing your MQTT topics.

- Ensure MQTT topics from left to right go from most broad to most specific routing information
- Include any relevant routing information in the MQTT topic such as the IoT application identifier, static device groups, and the unique identity of the device
- Prefix your MQTT topics to distinguish data topics from command topics
- Use only lowercase letters, numbers, and dashes “-“ in your MQTT topics
- Use the AWS IoT Shadow for command and control to devices
- Document topic structures and device communication patterns as part of your operations practice
- Use the AWS IoT Thing Name as the MQTT Client ID
- Include the Thing Name of the device in any MQTT topic or payload the device uses for publishing its data
- Review the AWS IoT Core default service limits. Design device communication patterns that align with adjustable IoT service limits
- Include additional contextual information about a specific message in the payload of the MQTT message
- Avoid MQTT communication patterns that result in a large fan-in to a single device
- Never allow a device to subscribe to all MQTT topics using #
- Only use multi-level wildcard subscriptions, #, in IoT rules
- Ensure devices do not share a single Device Shadow for commands

- Use the Device Shadow for device state or commands that have a medium to low transaction per second (TPS)
- Use the Device Shadow for storing status metrics of a device
- Store the current firmware version as an attribute in a device's shadow
- Use the optional clientToken field with the shadow when you need to track to sender of a shadow update
- Use Thing Groups to organize your IoT devices by common device attributes
- Implement Device Jobs with a staged rollout in order to gradually deploy fleet-wide commands across your devices
- For large-scale fan-in to the AWS IoT Rules Engine use a buffer such as Amazon Kinesis or Amazon SQS to separate your IoT application code from ingestion
- Create AWS IoT Rules responsible for tracking the progress of important MQTT messages in your application