

# Customizing Contentful

## A DEVELOPER'S GUIDE



# Table of Contents

2	Authors
3	Introduction, Kiersten Thamm
4	Content modeling process, Jim Sam
10	Configuring automatic backups, Edwin Maldonado
17	Data ownership with Contentful, Arjun Ganesh
21	Blue-Green deployment, Sia Rad
25	Smart web applications, Arjun Ganesh

# Authors

As a solutions architect with more than 16 years of experience, **Arjun Ganesh** designs cutting-edge solutions that help customers increase their speed and agility when bringing digital experiences to market.

In engineering roles and as a solutions architect, **Edwin Maldonado** prioritizes customers' needs and helps them learn the best practices of content modeling and software developer kits.

**Sia Rad** has worked as a technical architect and global IT director for more than 20 years, consulting, designing and delivering secure and distributed enterprise solutions across industries.

**Jim Sam** is a solution architect who is passionate about helping customers achieve their content goals. For nearly a decade, he has helped companies and nonprofits in their digital transformation by analyzing and designing systems that enable seamless access to content.

**Kiersten Thamm**, editorial manager of the Contentful blog and technical writer, is a Contentful power user and collaborates with engineers to write articles that help users build dream projects with Contentful.

# Introduction

The world's largest companies run on Contentful. Companies whose scale, scope and legal obligations often produce complicated challenges.

That's where Contentful's solution architects come in to help. Architects design solutions for the most complex situations, guiding customers to build secure and inventive digital experiences under any circumstance.

This book includes five solutions written by architects in response to specific customer requests; however, their solutions have wider applications. They cover advanced topics of content modeling, data backups and data ownership, blue-green deployment, and website personalization.

Some of these solutions depend on features only available to enterprise-level customers. If you aren't already working with Contentful, you can request a [free enterprise trial](#).

We hope this collection helps you use Contentful to build your most ambitious project.

# Content modeling process

Jim Sam



The content model is at the heart of a team's Contentful experience. Why? It's the technical way you structure content that determines business outcomes. Well-structured, strategic content enables digital teams to quickly ship their products across multiple channels.


How do you go about creating a successful content model? What does that process look like? We've hinted at this before with [content modeling basics](#), [composable entries](#), and [topics and assemblies](#), but let's discuss it on a more thorough level. Familiarity with creating databases or metadata schemas is helpful but not required, and familiarity with JavaScript is also helpful for some of the associated tools.

## Major themes

### TOPICS AND ASSEMBLIES

Content types and their entries can describe content (topics) or context (assemblies), which hard content flows into. Assemblies can aggregate topics or other assemblies. Any individual assembly can be medium-specific (web, mobile, etc.) or work across all media. The fields that go in a particular assembly have profound implications on content re-use, content security, personalization, workflows, and automation strategy.


A content type as a blueprint for content (a topic):

 **Product** The atomic description of the products the company offers. [Edit](#)

Fields (3)


JSON preview

Sidebar


 **Name** Short text

Entry title

Settings ...


 **Description** Rich text

Settings ...

 **Image** Media

Settings ...


A content type as a blueprint for context, a grouping of featured products (an assembly). Implicit in this assembly is that it will have a parent assembly:


 **Featured Products** An assembly of Products. [Edit](#)

Fields (2)

JSON preview

Sidebar

 **Title** Short text Entry title [Settings](#) ...

 **Products** References, many [Settings](#) ...

Note that this assembly does not imply a specific medium and could be used for any.

## FUNCTIONALITY CHOICES

Your content model depends on what functionality you want to prioritize. As with all software architecture decisions, there are trade-offs. You need to have in mind what [quality attributes](#) you are optimizing for: content reusability, authoring/editing usability, content security, automation, modifiability, app speed, etc. that can be in conflict with each other. There is no singular, definitively best content model.

## Who should define a content model?

We strongly suggest that a cross-functional team is involved early in the process. A successful content model drives successful business outcomes, and the best content model reflects the needs and expectations of everyone involved.

What may seem obvious to a developer may not be so evident to a content editor, and what an editor wants might be opposed to the needs of content strategy. For example, WYSIWYG editors in other content management systems do not allow you to enforce a content strategy between content creators, which leads to unique pieces of content that do not conform to a larger strategy.

A team can iterate on its content model more quickly when diverse roles discuss their challenges, and everyone maintains appropriate expectations since they understand the project as a whole.

## Optimizing for specific outcomes

The mechanics of defining what fields belong in what content types gets to the heart of the functionality question.


In general, each content type should be a conceptual thing unto itself (i.e. product, row of cards, button, etc.). But it's a key question every time, if an attribute or group of attributes should be a field on that content type or abstracted to a referenced content type. Different optimizations will show different, often unforeseen problems.

To better illustrate this, let's walk through a few common scenarios and outline the pros and cons. This is by no means an exhaustive list of trade-offs or content modeling tactics that you might face, but more of a greatest-hits collection. Where your team falls in each is dependent on your objectives.

## Authoring usability vs. content reusability

The hardest choice we often see with content modeling is the question of how flat a model should be, flat meaning there are many fields per content type with very few references.








Flat forms are easy to understand for an editor, but this flat pattern does not lend itself to content reuse through well-structured content. The same content is often manually entered again and again across entries. When this content requires an update to change a value across all entries, someone must either write a script to programmatically change the values or change each manually. It's much more elegant to abstract those values into a referenced entry.

 **Flat Product Page** A flat product page acting as a blueprint for a Product Page. [Edit](#)


Fields (7)

JSON preview

Sidebar

⋮	 <b>Title</b> Short text	Entry title	<a href="#">Settings</a>	⋮
⋮	 <b>Slug</b> Short text		<a href="#">Settings</a>	⋮
⋮	 <b>Metadata Keywords</b> Short text, list		<a href="#">Settings</a>	⋮
⋮	 <b>Body</b> Rich text		<a href="#">Settings</a>	⋮
⋮	 <b>Button Name</b> Short text		<a href="#">Settings</a>	⋮
⋮	 <b>Button Color</b> Short text		<a href="#">Settings</a>	⋮
⋮	 <b>Button Italics</b> Boolean		<a href="#">Settings</a>	⋮





This is an example of an abstracted content type for a product web page.

 **Abstracted Product Page** An abstracted content type acting as a blueprint for a Product Page. [Edit](#)

Fields (4)

JSON preview

Sidebar

⋮	 <b>Title</b> Short text	Entry title	<a href="#">Settings</a>	⋮
⋮	 <b>SEO Object</b> Reference		<a href="#">Settings</a>	⋮
⋮	 <b>Product</b> Reference		<a href="#">Settings</a>	⋮
⋮	 <b>Button</b> Reference		<a href="#">Settings</a>	⋮



## Understandability vs. omnichannel reusability

People often design a content model clearly geared toward one medium, which can cause omnichannel problems down the road.

For example, we often see a content model that mimics a website's layout with landing page, category page, product page, etc. as content types. It's easy to understand for all stakeholders, users can be quickly onboarded, and this maps nicely to the front-end code for a website. We see this many times, and it suits many customers well.

The downside to this pattern is poor content reusability. The mixing of hard content (a product) with context (a page describing the product) means the wrong content is pulled from the API in a different medium. If you want to reuse the hard content in a mobile app or digital signage when displaying product info using our REST APIs, you would likely pull an excessive amount of content from a product page entry that doesn't make sense for a non-web medium. Using the GraphQL API could be a solution, but this may create unwanted cognitive overhead when maintaining your app over time.

## Extensibility automation vs. authoring usability

Abstracting fields into their own content type offers some benefits to developers while challenging content authors and editors.

Another example of this relates to automation through webhooks. As a content model becomes more structured through assemblies, it is possible to finely tune when a webhook fires to another service. This is a similar strategy of developers including extra divs on a web page for styling reasons. However, like the question of omnichannel reusability, this introduces added complexity to a content creator or editor.

## Automation vs. space limits and maintainability

From time to time, customers map content types to React components in their design system. On a technical level, this seems like a good idea. Having one-to-one mapping between Contentful and the front end makes sense. However, we generally do not endorse this pattern for two reasons. First, it often results in a high content-type count, which pushes against the limits of purchase limits. Second, components and content types will change as the app goes through its life cycle. Keeping them in sync may become difficult to maintain, especially if the app is only partially redesigned. Instead, we suggest using higher order components or selectors in redux to do this mapping.

## Malleable content models

With Contentful, the content model is not set. It may need to change over time due to agile iteration, honest mistakes, or changing business needs. Our powerful, easy-to-learn migration tooling allows you to refine or overhaul it as needed.

Content models require modification for any number of reasons. Maybe you need to move some fields into their own content type so that content can be reused and data entry time diminished, such as in [this video](#).

So how is this done?

Combining our [content migration DSL and tooling](#) with [space environments](#) allows your dev team to prepare migration scripts for the next version of your content model independent of your editors and programmatically apply these changes — [CMS as code](#).


As you adjust your content model, we suggest using an expand/contract process. In this, one migration script will create the necessary new fields, references, and new content types. A second script will adjust the corresponding entries so that the new and old schemas temporarily exist together, with old fields disabled for editing and the new fields possibly omitted from the API response. Once all applicable parties have approved the new version of the content model, a final script runs, cleaning out the old content.

## Conclusion

Contentful allows you to define your content model and equips you to ensure your content model promotes intended outcomes, even as your digital products evolve. However, a content model is just another software architecture choice, and there are trade-offs between functionalities. Keeping this in mind will help you develop your optimal content model.

# Configuring automatic backups

Edwin Maldonado



How to set up automatic backups for Contentful spaces using AWS Lambda Function and CloudWatch for scheduling.

Contentful maintains a [command line interface \(CLI\)](#) which offers a `space export` command that lets you download a backup of your Space. If you are interested in reading about the different options that you can pass to the space export command, go to the readme on [Github](#).

```
contentful space export -- space-id {SPACE-ID} -- environment-id {SPACE-ENV} -- management-  
Token {CMA-TOKEN} -- exportDir {YOUR-PATH}
```

This will create a backup that is a JSON file with the whole space structure (content model, entries, assets, roles & permissions and webhooks).

The CLI `space export` command works great but we have bigger plans. We want to automate this process by scheduling a routine to do it for us.

Thankfully, there is yet another tool called the [Contentful export tool](#), which is a Javascript library that helps you to backup your space (as the `space export` CLI) in a programmatic way.

## Process

1. Generate a backup of your space. By writing a NodeJS script that can collect the export options, download the backup and save it as a JSON file locally.
2. Upload the backup to AWS S3. Once you have the JSON dump locally, you can upload it to a S3 Bucket where you can collect the backups.
3. Wrap your backup script in a Lambda function so that you have the advantage of triggering it in a variety of ways. One good option is AWS CloudWatch, which can be configured to run as a Cron Job.
4. For the purpose of this tutorial, I will schedule my script on a daily basis. You are free to adapt it to your needs.

## Implementation

Create a new AWS Lambda function from the [AWS Console](#). You will quickly notice that there are three ways to create a lambda deployment package:

1. Use the built in editor
2. Upload a .zip file
3. Upload a file from Amazon S3

Use the `Upload a .zip file` option, which means that you will need to upload a zip file that contains our function `.js` file and the `node\_modules` required for our function.

Create a `package.json` file with the below dependencies:

### PACKAGE.JSON

```
{  
  "dependencies": {  
    "aws-sdk": "2.450.0",  
    "file-system": "",  
    "contentful-export": ""  
  }  
}
```

Create another file called `index.js` This file will contain your main function — now you're only downloading our backup file to our local `tmp` folder:

## INDEX.JS

```
'use strict'

const AWS = require('aws-sdk');
const fs = require('fs');
const contentfulExport = require('contentful-export');
exports.handler = async (event) => {
  const local_backup_path = '/tmp/'
  const file_name = 'contentful_backup.json'
  const options = {
    spaceId: process.env.SPACE_ID,
    environmentId: process.env.SPACE_ENV,
    managementToken: process.env.MANAGEMENT_TOKEN,
    contentFile: file_name,
    exportDir: local_backup_path,
    useVerboseRenderer: false,
    saveFile: true
  }
  try {
    const result = await contentfulExport(options);
    return sendResponse(200, 'File downloaded');
  } catch (err) {
    console.log('Oh no! Some errors occurred!', err);
  }

  return sendResponse(200, 'End of the Function');
};
```

Note that we are using environment variables on the script, you can set these `ENV` variables on the AWS Lambda function page.

**Environment variables**

You can define environment variables as key-value pairs that are accessible from your function code. These are useful to store configuration settings without the need to change function code. [Learn more](#)

MANAGEMENT_TOKEN		Remove
S3_BUCKET_NAME		Remove
SPACE_ENV		Remove
SPACE_ID		Remove
Key	Value	Remove

► Encryption configuration

You still need to upload the dump file to an S3 bucket. In the AWS Console, create a new bucket (or pick an existing one). Once you have your bucket in place, you need to set up the communication between our lambda function and the S3 bucket. In order to do that we have two options:

1. Use your API credentials in the lambda function to talk to S3.
2. Or set up a new IAM Role.

I lean towards the IAM Role option to keep things clean in our function. So, again from the AWS Console, go to the `Security, Identity, & Compliance` section and choose IAM.

Create the new IAM role and make sure to attach the AmazonS3 policy to it.

**Permissions** | Trust relationships | Tags | Access Advisor | Revoke sessions

▼ Permissions policies (1 policy applied)

Attach policies [Add inline policy](#)

Policy name ▼	Policy type ▼
► AmazonS3FullAccess	AWS managed policy

► Permissions boundary (not set)

Link the new IAM Role to your lambda execution option:


### Execution role

Choose a role that defines the permissions of your function. To create a custom role, go to the [IAM console](#).

Use an existing role ▼

### Existing role

Choose an existing role that you've created to be used with this Lambda function. The role must have permission to upload logs to Amazon CloudWatch Logs.

super\_lambda\_role ▼ 

[View the super\\_lambda\\_role role on the IAM console.](#)

This allowed our lambda function to talk directly with our S3 bucket. Let's write the uploader part then, and this is how the final function looks:

```
'use strict'

const AWS = require('aws-sdk');
const s3 = new AWS.S3();
const fs = require('fs');
const contentfulExport = require('contentful-export');

exports.handler = async (event) => {
  const local_backup_path = '/tmp/'
  const file_name = 'contentful_backup.json'

  const options = {
    spaceId: process.env.SPACE_ID, // Replace it
    environmentId: process.env.SPACE_ENV, // Replace it
    managementToken: process.env.MANAGEMENT_TOKEN, // Replace it

    contentFile: file_name,
    exportDir: local_backup_path,
    useVerboseRenderer: false,
    saveFile: true
  }
}
```

```

    try {
      const result = await contentfulExport(options);
      console.log('Your space data has been downloaded!');
      // Save file to AWS S3
      console.log('Preparing file for AWS S3!');
      const outputFile = local_backup_path + file_name;
      let fileBuffer = new Buffer(fs.readFileSync(outputFile));
      fs.unlinkSync(outputFile);

      await uploadFile(fileBuffer);

      return sendResponse(200, outputFile);
    } catch (err) {
      console.log('Oh no! Some errors occurred!', err);
    }

    return sendResponse(200, 'End of the Function');
  };

const sendResponse = (status, body) => {
  var response = {
    statusCode: status,
    body: body
  };

  return response;
};

const uploadFile = async (buffer) => {
  let params = {
    Bucket: process.env.S3_BUCKET_NAME,
    Key: 'backups/' + Date.now().toString() + '.json',
    Body: buffer
  };

  return await s3.putObject(params).promise();
}

const uploadFile = async (buffer) => {

```



```

let params = {
  Bucket: process.env.S3_BUCKET_NAME,
  Key: 'backups/' + Date.now().toString() + '.json',
  Body: buffer
};

return await s3.putObject(params).promise();
}

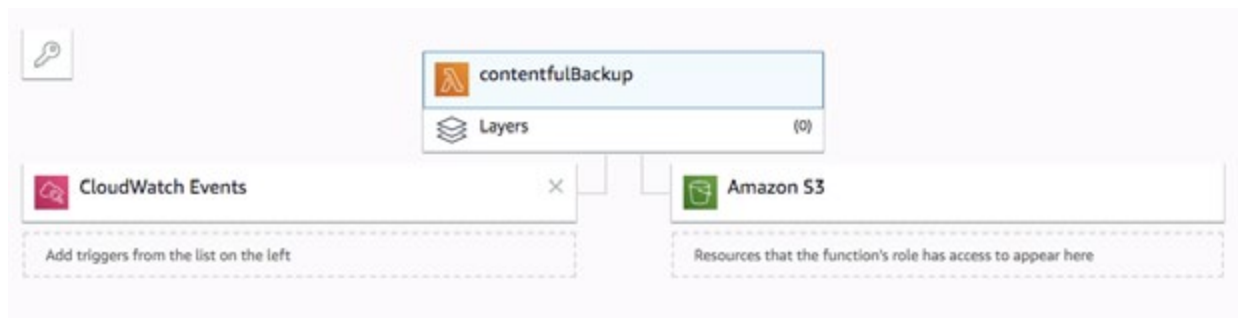
```

## Scheduling the function

Pick a trigger for our lambda function, specifically a CloudWatch Event. Create the trigger and schedule it as you wish, I used the daily option:

```
`Schedule expression: rate(1 day)`
```

Make sure to compress your deployment package (`index.js` and `node\_modules`), upload the zip file to the lambda's function code section and test it!



This is how the final solution looks on the AWS lambda page.

# Data ownership with Contentful

Arjun Ganesh



How to own your data and make sure it is always available no matter what.

Contentful provides 99.96 to 99.99 availability depending on your enterprise SLA. This is sufficient in most cases, but sometimes you need even more reliability. Enterprises that store critical data in Contentful may need to satisfy legal or compliance requirements associated with data ownership and business continuity. Regulated and public enterprises might need to have backup protection in the unlikely event that Contentful ceases to operate in the future.

Contentful and AWS S3 offer a reliable, secure and out-of-the-box integration that solves these demands. It backs up data on your own infrastructure and allows you to maintain an audit log. You could also build a custom HA/DR/BPC procedure with Contentful.

## Solution

Integrating AWS and Contentful begins by setting up an S3 bucket on AWS. Below are screenshots that show how an S3 bucket is created via the AWS console:

A screenshot of the AWS Management Console 'Create bucket' wizard. The interface is dark-themed with a blue header. The header shows the title 'Create bucket' and a close button. Below the header is a progress bar with four steps: 1. Name and region, 2. Configure options, 3. Set permissions, and 4. Review. The first step, 'Name and region', is active. It contains three input fields: 'Bucket name' with the value 'mybucket', 'Region' with the value 'Asia Pacific (Sydney)', and 'Copy settings from an existing bucket' with the value 'You have no buckets'. At the bottom of the form are three buttons: 'Create', 'Cancel', and 'Next'.

You will need to enable versioning, as seen below. This stores multiple versions of the same file as one:

**Create bucket**

1 Name and region    2 **Configure options**    3 Set permissions    4 Review

**Properties**

**Versioning**  
☒ Keep all versions of an object in the same bucket. [Learn more](#)

**Server access logging**  
☐ Log requests for access to your bucket. [Learn more](#)

**Tags**  
You can use tags to track project costs. [Learn more](#)

Key	Value
<a href="#">+ Add another</a>	

**Object-level logging**  
☐ Record object-level API activity using AWS CloudTrail for an additional cost. [See CloudTrail pricing](#) or [learn more](#)

**Default encryption**  
☐ Automatically encrypt objects when they are stored in S3. [Learn more](#)

► **Advanced settings**

**Management**

[Previous](#) [Next](#)


Use an existing user with the right permissions, or generate a new Access ID and Access Key that can be used from Contentful to write in this folder. The latter is the better option.


That's all that is required from AWS. The remaining configuration can be done from the Contentful web app. You need to be an admin or an owner of the space that you're using to do so.


Create a new webhook within Contentful using the AWS S3 template.


Webhook Templates


TEMPLATES (18)


 Bitbucket  
Trigger a pipeline


 AWS Lambda  
Invoke a function  
Enterprise plan only


 Google Cloud  
Run a function


 Webtask  
Run a function


 Slack  
Notify a channel


 Twilio  
Send a SMS


 Mailgun  
Send an email


 AWS SQS  
Send a message  
Enterprise plan only

 PubNub  
Publish a message

 AWS S3  
Store entries  
Enterprise plan only

 Algolia  
Index entries

 Elasticsearch  
Index entries

 Jira  
Create a task

### AWS S3

- Stores entries in an S3 bucket every time they are modified
- Scoped to events in the master environment
- Object key is the entry ID and with .json suffix
- Object contains JSON data with all topic, user ID and the entity itself
- Should be used with a versioned bucket

AWS region (required)

The AWS region of your queue. For example: eu-west-1.

Bucket (required)

The name of a bucket you want to store your entries.

AWS Access Key Id (required)

Use a keypair with minimal access. The only required policy action is `s3:PutObject`

Secret Access Key (required)

Secret Access Key of the keypair used above. This value can't be revealed once stored.

All properties can be updated later.

Create webhook

Cancel

Enter the region and the name of the bucket, along with the access credentials.

When an entry is created or updated, a JSON file is stored in this S3 bucket with all the info of the entry. The JSON file is versioned — you always see the latest version. If you want to publish content and aren't concerned with draft content, then update the events when the webhook is triggered under the webhook settings.

This covers most of the use cases:

1. Backup of Contentful data in your AWS S3.
  - a. You own your data, it is on your cloud.
  - b. You're ensured compliance and data availability in the unlikely event that Contentful does not exist in the future
3. Create an audit log, because the data contains user information for each entry. Notice each JSON is associated with the user ID who made the change.

We still need to do one more thing to satisfy HA, DR and BCP requirements. If you are accessing Contentful directly from a client:

1. Create a lambda or similar function that can be called from the client
2. In the lambda make the Contentful call
3. If this fails, make the call to the S3 bucket to get the corresponding JSON
4. Massage the JSON from the S3 and remove unwanted elements
5. Return it to the client

If you have a service that merges Contentful content and some form of React/Mustache templates, update this service to call the S3 bucket (similar to above) if Contentful is down.

Now — even if Contentful is down — your digital experience will still function with the latest published data from the S3 bucket.

## Conclusion

This simple, out-of-the-box integration between Contentful and AWS S3 is reliable and solves complex enterprise requirements. It demonstrates just how flexible and reliable SaaS- and API-based integrations are.

## Useful links

[Webhooks](#)

[Contentful Delivery APIs](#)

[AWS S3 bucket \(create. read\)](#)

# Blue-Green deployment

Sia Rad



User experience plays a crucial role in the success of every business. Displaying the latest content to customers without interrupting their experience is a win-win for both businesses and customers.

Some businesses struggle to make their latest changes available to customers with the agility that the market demands due to the redundant and inherent dependencies in their developer operations (DevOps) process. Time to market increases when platforms cannot eliminate the dependencies between making and deploying iterative changes.

Business leaders and marketing teams need an agile approach to let customers know about products and services just in time. Developers need a platform that enables them to deploy changes while future changes are being developed without impacting customer experience. And consistent and reliable access to the latest content and features keeps customers happy.

## Solution summary

You can leverage Contentful's out-of-the-box capabilities to build the solution. Refer to [CMS as Code](#) to learn how to use spaces and space environments if your DevOps is not already doing so.

Two spaces can be dedicated to your agile DevOps (Blue and Green) and a config file (self-hosted) to switch from the Blue to Green or vice versa. Doing so enables you to serve your front end with the content from one space (Blue acting as the active space, serving as production) while your dev team works on new changes in the other space (Green acting as the standby space, serving preproduction) without impacting production. Once the changes are ready to go live, a config change would promote the preproduction to production and vice versa. As a result, Blue becomes preproduction and Green becomes production until the next round of continuous deployment.

## Prerequisite concepts

Take some time to review and learn about the following concepts before you continue:

1. [CMS as Code](#), also known as Infrastructure/Content as Code
2. [Managing multiple environments](#)
3. [Contentful CLI](#) and [CMA SDKs](#)

## Use case

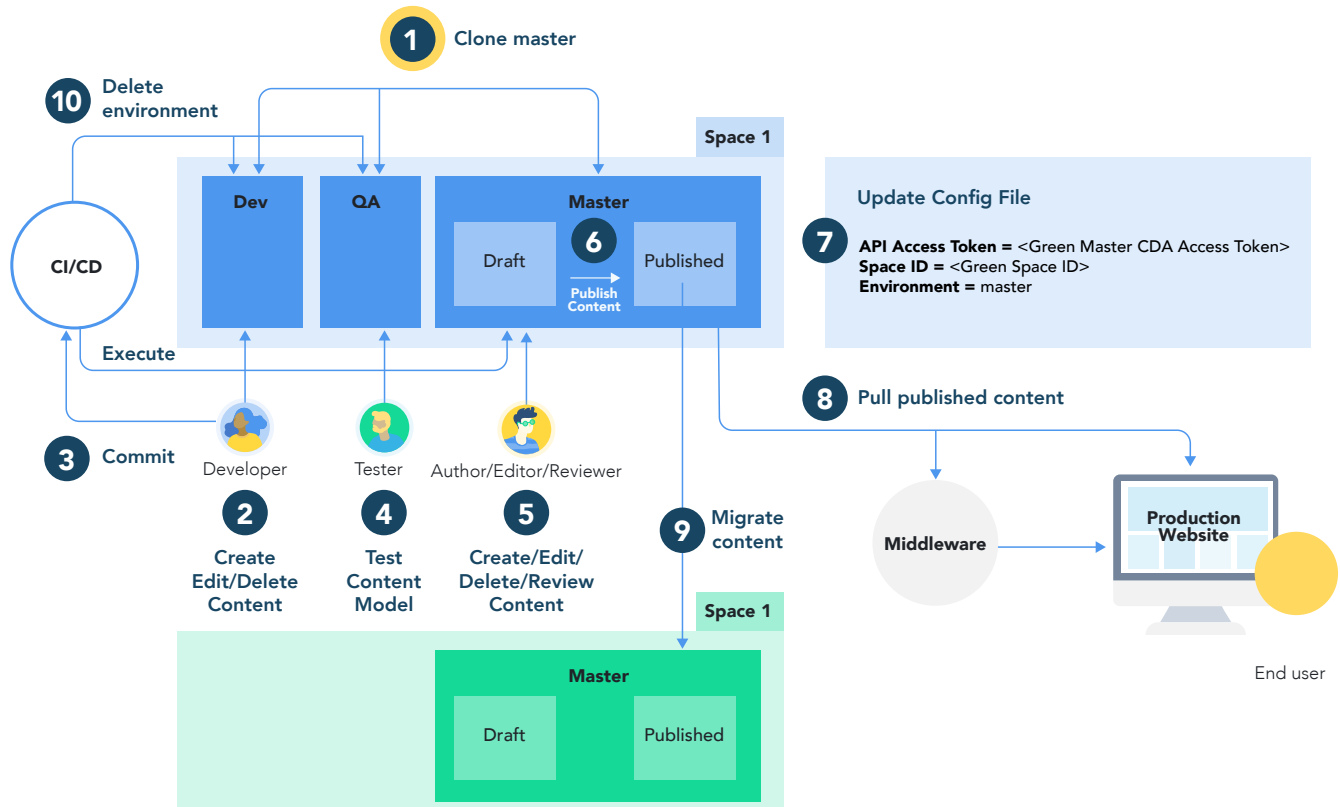
A healthcare provider company in North America with more than 75 million members wanted to build a platform that had a high level of availability for customers. They also wanted their IT team to build and deploy content models, and revamp, review and publish content without interrupting the customer experience. Previously, they made iterative releases and pushed changes from development to test and production. But each deployment of a new release caused downtime. Some of the deployments took hours and needed to happen overnight to minimize the impact on the customer experience.

The requested platform required a DevOps process that uses the following components:

1. Two spaces (known as Blue and Green)
2. Space environments (e.g. Dev, QA, Master)  
*Note: The number of environments and the naming convention could vary for different businesses.*
3. Contentful CLI/CMA scripts to apply changes from Dev, QA, Master
4. Code repository to commit the scripts and check out and run as needed
5. Script to publish the draft content, if applies per release
6. A self-hosted config file containing:
  - a. Prod Space API Access Token
  - b. Prod Space ID
  - c. Prod Environment Name (this is currently optional, since it isn't possible to rename the "master" environment at this point. Let's keep it here for forward compatibility.)

## Blue-Green deployment process

Now that you have a way to build the Blue-Green architecture, let's look at what process should be used to properly build and deploy changes in this model.



- Let's assume that the Blue space (production) is serving the production front end and the Dev and QA envs are created as clones of the Master on the Green Space (Pre-Prod).
- Script the changes, run it on the Dev environment, and do unit testing.
- Commit the script into your CI/CD code repository.
- Deploy the changes on the QA env by running the same script and review and QA the changes.
- Deploy the changes on the Master by running the same script and preview and validate the changes.
- Publish any content that needs to be part of this release. Note: At this point, you are ready to serve production front end by the latest changes from the Green space.
- Update the config file with the Green space properties (i.e. Space ID, CDA Access Token, and master environment name).
- Start serving the front end by pulling the latest changes from the Green space.
- Assuming you have already built a script to refresh everything on the Master environment of the Blue space with a copy of the Master environment from the Green space, run the script and validate the results.
- Dispose of the environments on the Green space if your normal practice is to recreate those environments for every sprint.
- Repeat the same process for the next release by switching Blue and Green space roles.



## Expected outcome

Implementing this solution will enable businesses and their IT teams to streamline their DevOps process and deploy changes to production without having production downtime and interrupting their customer experience. It will improve time to market and increase customer satisfaction and sales.

## Conclusion

A Blue-Green deployment offers a solution for IT teams unable to deploy changes to production without degrading customer experience. A Blue-Green deployment provides the architecture to automate the process, eliminate dependencies, and deploy changes without scheduled production outages or unscheduled interruptions to service.

## Useful links

[CMS as Code](#)

[Managing multiple environments](#)

[Contentful CLI](#)

[Contentful Delivery APIs](#)

[Contentful Management APIs](#)

# Smart web applications

Arjun Ganesh



Your web application should have a top-down approach from the view of your customers, putting your customers at the top. It can be as simple as a sleek, user-friendly layout of interactive images sprinkled with personalized data. The latter is what adds intelligence to your site. Customers want to see more of the things they like in an enjoyable, immersive experience.

Various cloud-based services make this possible. Using Contentful, Optimizely and Google Analytics, you can achieve simple, mildly personalized web applications that can then be moderated based on performance statistics. It's like Lego pieces for building a site.

## Where to start?

Every web application needs content. Let's start with that. You need somewhere to securely store the content that allows you to retrieve and edit it quickly without developers. Get Contentful. Your editors and authors can manage content within Contentful, and your developers can retrieve content via APIs and integrate it with their front end code of choice.

Content creators are always trying to figure out what content to display when, on which page and to whom. An experimentation platform enables them to do this without the guesswork. Optimizely seamlessly integrates with Contentful and equips you with the capability of targeting audiences and the APIs to access them. Google Analytics provides the data that content creators need to make strategic decisions. Data about what performs well, what makes the highest sale and generates the most traffic.

## Build build build

Create a content model inside Contentful that can power the data for your user experience and front-end code. Get content creators to start populating content. Numerous articles cover this topic in detail, and I have listed some at the end of this article.

To use the Optimizely feature in Contentful, the fields that content creators would like to experiment on need to be references. For example, if a call-to-action button is required, make it a child reference type to your home page parent type.

Next, set up your experimentation platform. Create variations and audiences. Define targeting criteria such as cookies or other context parameters such as device or location).

Contentful Integration Web

## Audiences


Saved   Attributes

---

Active ▼
Browse by Name
Create New Audience...

Name ▼	Created	Modified	
Developer	Apr 11, 2018	Apr 11, 2018	...
Digital Leader	Apr 12, 2018	Apr 12, 2018	...
Editor	Apr 12, 2018	Apr 12, 2018	...

To connect Optimizely to Contentful, go to apps in the Contentful web app and enable the Optimizely app. You should see a screen similar to below:

 App: Optimizely

Save

Let us know how we can improve the Optimizely app. [Give feedback](#)

### Features

Optimizely integration enables:

- Loading experiments from Optimizely
- Adding content to variations
- Seeing all experiments connected with Contentful (drafts, running, ended)

Check out [documentation](#) for more information.

### Connect Optimizely

Connect your Optimizely account to get started.

Personal Access Token (required)

.....

Disconnect account

Get a personal access token in the [Optimizely web app](#).

### Optimizely Project

Works only with Optimizely Full Stack projects

Project (required)

[Demo] The Example App

### Content Types

Select the content types for which you want to enable A/B testing.

Add content type

Layout > Highlighted Course

Course

Edit Delete

Enter your Optimizely Personal Access Token and connect your account. This will bring up a list of your projects. Select the one you created. Add the relevant Contentful content types that you want to run this experiment against. If you select the parent type, you should see the referenced child that you want to experiment with.

Edit content type

Content Type (required)

Layout > Highlighted Course

Reference Fields

☒ Course

Save

Cancel

When you create an entry, you will see a drop-down menu that gives you the option to select a variation container.

Course (required)

+ Create new entry and link

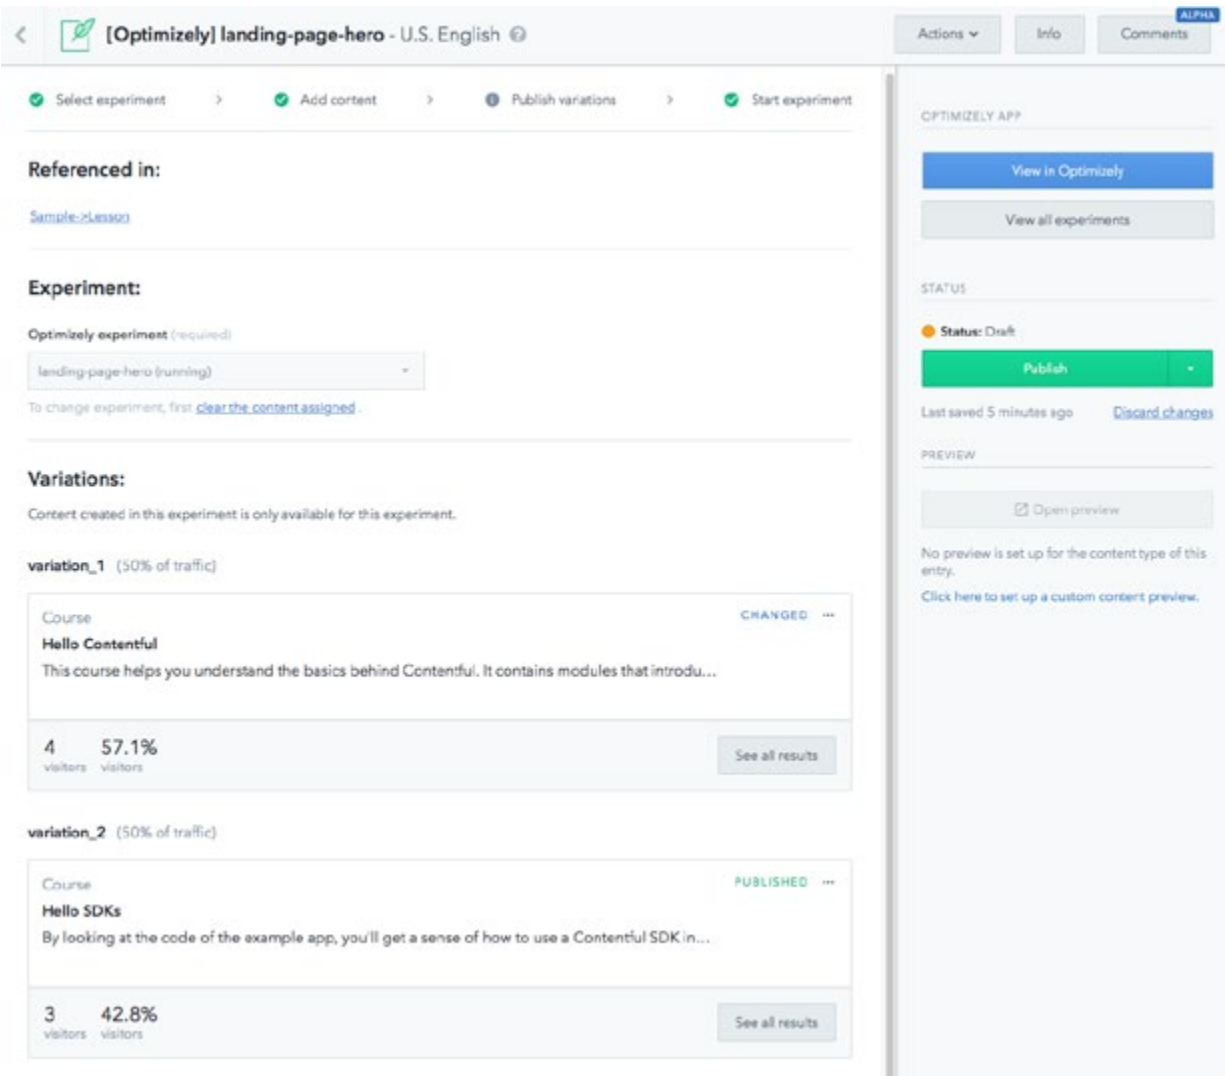
Link existing entry

ALL CONTENT TYPES

Course

Variation Container

A variation container brings up the Optimizely variations and allows you to select the content you would like to display for each variation.



Publish the entry and let the front-end code work with the Optimizely API.

Surface your Google Analytics data inside the Contentful web app as a custom sidebar UI extension. You need to build a JS page that surfaces stats from Google Analytics and install it as a sidebar UI extension in the web app. I have included additional resources under useful links.

## Conclusion

Contentful provides a single platform for practitioners to create content, run Optimizely experiments and make decisions based on data derived from Google Analytics. This reduces complexity and increases go-to-market speed. It empowers practitioners to create content based on numerous experiments — resulting in an awesome customer experience.

## Useful links

[Content Modeling Basics](#)

[Contentful Delivery APIs](#)

[Contentful Management APIs](#)

[Sample Google Analytics Integration](#)