

# Getting started with Red Hat OpenShift Service Mesh

A guide to production deployment and Day 2 operations

By Stelios Kousouris, Senior Applications Architect,  
and Ortwin Schneider, Principal Product Marketing Manager

# Contents

## Introduction

### Part 1: Heading to production

Chapter 1: Service mesh requirements and prerequisites

Chapter 2: Development environment set up

Chapter 3: Production environment set up

Chapter 4: Onboarding an external travel portal

Chapter 5: Complying with a new security regulation

Chapter 6: Expanding partnerships and implementing broker services

### Part 2: Day 2 operations

Chapter 7: Troubleshooting the mesh

Chapter 8: Tuning the service mesh

Chapter 9: Upgrading your service mesh to a new version

# Introduction

Technology and applications are at the core of today's digital businesses, and many organizations are building modern service management architectures to support rapid, iterative application development and deployment. **Service meshes** play a key role in these architectures, connecting sets of discrete services into complete applications. They streamline application development, simplify troubleshooting, enhance observability, increase resiliency, and facilitate performance tuning.

Even so, building your service mesh is only the first step. To derive value from it, you must configure, tune, and maintain your service mesh according to your organization's requirements. This e-book provides guidance on governance, design practices, and configuring Red Hat® OpenShift® Service Mesh for production use and on performing Day 2 operations.

## Red Hat OpenShift Service Mesh basics

**Red Hat OpenShift Service Mesh** – included with Red Hat OpenShift – provides a uniform way to connect, manage, and observe microservices-based applications. It incorporates a set of open source projects for integrating, managing, tracing, monitoring, and analyzing traffic between microservices:

- ▶ **Istio**, an open source project for integrating and managing traffic between services.
- ▶ **Jaeger**, an open, distributed tracing system that tracks requests as they move between services.
- ▶ **Kiali**, an open source project for viewing configurations, monitoring traffic, and analyzing traces.
- ▶ Multiple networking interfaces.
- ▶ The **Red Hat 3scale Istio plugin** for integration with Red Hat 3scale API Management.

## Key features and capabilities

Red Hat OpenShift Service Mesh provides a number of key capabilities uniformly across a network of services:

- ▶ **Traffic management.** Control the flow of traffic and application programming interface (API) calls between services, make calls more reliable, and make the network more robust in the face of adverse conditions.
- ▶ **Service identity and security.** Provide services in the mesh with a verifiable identity and provide the ability to protect service traffic as it flows over networks of varying degrees of trustworthiness.
- ▶ **Policy enforcement.** Apply organizational policy to the interaction between services, ensure access policies are enforced and resources are fairly distributed among consumers.
- ▶ **Telemetry.** Gain understanding of the dependencies between services and the nature and flow of traffic between them, providing the ability to quickly identify issues.

With Red Hat OpenShift Service Mesh:

- ▶ Developers can focus on adding business value, instead of connecting services.
- ▶ Problems are easier to identify and diagnose via distributed request tracing and a visible infrastructure layer.
- ▶ Applications are more resilient to downtime, since a service mesh can reroute requests away from failed services.
- ▶ Communications in the runtime environment can be optimized using performance metrics and observability.

## Online resources for this e-book

In addition to this e-book, we have also created an [online repository](#) of associated resources. The repository includes additional details, setup scripts, and code examples that you can use with this e-book. Links are included throughout the text, and you can click on the link icon next to each section to go directly to that section in the online repository.

## Further reading

Learn more about Red Hat OpenShift Service Mesh and how it can help your teams connect, manage, and observe microservices-based applications:

- ▶ [Red Hat OpenShift Service Mesh product documentation](#)
- ▶ [Red Hat OpenShift Service Mesh product information](#)
- ▶ [Service mesh or API management e-book](#)

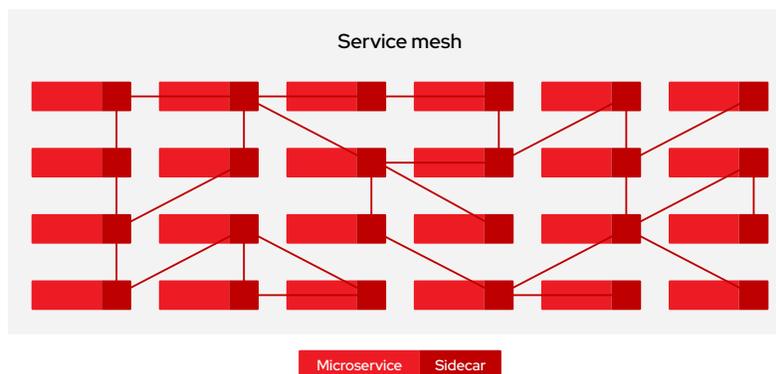


Figure 1. Conceptual service mesh architecture

### Try Red Hat OpenShift Service Mesh

Get started with Red Hat OpenShift Service Mesh in a local environment. Install Red Hat OpenShift Local on your laptop to follow along with the exercises in this e-book. With Red Hat OpenShift Local, you can also experiment with a service mesh – and other Red Hat OpenShift features – on your own.

[Download Red Hat OpenShift Local.](#) →

# Part 1: Heading to production

Organizations deploy service meshes for many use cases, but most have common underpinnings. To better illustrate how Red Hat OpenShift Service Mesh deployment and operations looks in real life, we'll follow a fictional organization—Travel by Keyboard—as they set up, operate, and maintain a service mesh.

## About Travel by Keyboard

Travel by Keyboard is an agency specializing in all-in-one travel packages to worldwide destinations. The agency works with airlines, hotels, car rental services, and insurance providers to create complete travel experiences for its customers.

Travel by Keyboard's IT department operates a travel portal and booking platform system architecture to support the business. Currently, the travel portal and booking platform do not use a service mesh, but the organization would like to deploy a service mesh to improve flexibility and resiliency while streamlining development and operations.

Travel by Keyboard's target service mesh architecture is shown in Figure 2. It features a microservices backend architecture with services for travels, flights, hotels, cars, and insurances (contained in the `travel-agency` namespace). The `travel-portal` namespace will include tenants for customers and partners. Travel by Keyboard also wants the option to integrate with external travel portals via APIs to allow them to expand and extend their services as needed.

Several teams and roles within the Travel by Keyboard organization will interact with the service mesh architecture:

- ▶ The Travel Portal product team will own the travel portal infrastructure and applications.
- ▶ The Travel Services product team will own the travel services infrastructure and applications.
- ▶ A single product owner will own the complete travel agency platform.
- ▶ A platform admin will maintain the Red Hat OpenShift platform and operators.
- ▶ A service mesh operator will manage the service mesh control plane.
- ▶ Several mesh developers will create traffic configurations for the service mesh.

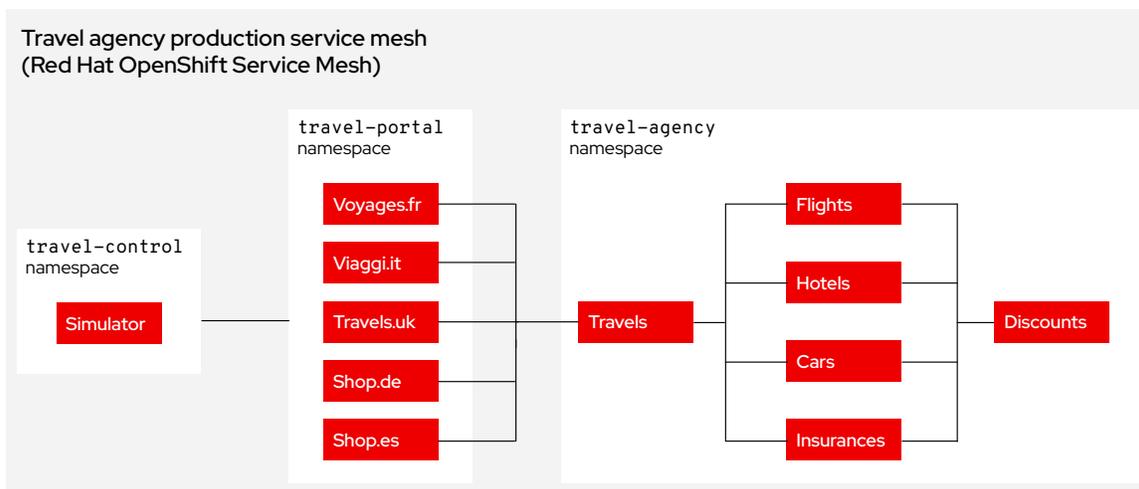


Figure 2. Travel by Keyboard's target service mesh architecture

## Chapter 1

# Service mesh requirements and prerequisites

The first step in deploying a service mesh into production is identifying your business and architecture requirements. You can collect these requirements in a variety of ways, but it's important to engage and understand the expectations of all stakeholders.

## Identify service mesh requirements

In our example, the key stakeholders are the development, product, security, and platform teams. At the project kick-off meeting, the stakeholders identified the following requirements.

### Development team requirements:

- ▶ Ability to trace every request during development and sample 20% of traffic in production
- ▶ Features and configurations to improve application resiliency, stability, and reliability

### Product team requirements:

- ▶ Ability to collect reportable performance and use metrics
- ▶ Ability to store metrics for up to one week

### Security team requirements:

- ▶ Mutual transport layer security (mTLS) for all intramesh and intermesh communications

### Platform team requirements:

- ▶ Centrally managed security
- ▶ Declarative approach – like GitOps – for configuring and managing the service mesh

## Define service mesh roles and responsibilities

The next step is to define the personas, roles, and responsibilities that will be involved with building and managing your service mesh. Several organizational, operational, and governance factors can impact how these items are set up. Factors include:

- ▶ **The type of clusters and meshes used.** Clusters can be multidomain application clusters or focused clusters, while meshes can be multitenant or single tenant.
- ▶ **The type of automation used for cloud service configuration.** Choices include pipelines, GitOps, automation platforms, scripts, and management tools.
- ▶ **The platform or service mesh operating model.** This defines how the service mesh is accessed. A producer-consumer platform lets administrators and operators deploy all configurations for developers to consume, while a self-service platform allows developers to provision preapproved configurations themselves.
- ▶ **Whether you have adopted DevSecOps approaches for application and cloud configuration delivery.** [DevSecOps approaches](#) bring together development, security, and operations into a collaborative shared-responsibility paradigm. The goal is to break down barriers between roles, disciplines, and teams across an organization to encourage collaboration and work toward common goals. A DevSecOps approach encompasses people, processes, technology, and governance.

## Operational setup

Our example will use the following operational setup:

- ▶ Cluster type: Focused clusters
- ▶ Mesh type: Cluster-wide single tenant
- ▶ Automation: Scripts, moving to GitOps over time
- ▶ Operating model: Self-service (restricted)
- ▶ DevSecOps: Adopted

Now you can define your enterprise personas and assign service mesh roles and responsibilities to them. [Setup scripts](#) for each role in this example are linked in the table and also provided in the online resource repository.

**Table 1. Enterprise personas, service mesh roles, and responsibilities**

Enterprise persona	Service mesh role	Responsibilities
Platform administrator	Cluster admin (default admin role)	<p>The cluster admin owns and operates Red Hat OpenShift clusters and sets organizational policies. They have all cluster privileges.</p> <ul style="list-style-type: none"> <li>▶ Add, remove, and update cluster operators.</li> <li>▶ Install container images to the image registry.</li> <li>▶ Update OpenShift versions.</li> <li>▶ Set up cluster infrastructure and configurations, including routers, networking, nodes, and resources.</li> <li>▶ Set up security.</li> <li>▶ Provide service mesh resources to retrieve logs for troubleshooting.</li> </ul>
Mesh operator	<b>Mesh operator</b>	<p>The mesh operator operates the service mesh control plane namespaces. They can administer one or more service meshes, depending on your operational setup.</p> <ul style="list-style-type: none"> <li>▶ Add, remove, and update the service mesh control plane, member role, member, and Istio resources within the owned namespaces.</li> <li>▶ Configure the service mesh observability stack.</li> <li>▶ Set up mesh security and certificates.</li> <li>▶ Configure ingress and egress gateways for both north-south and east-west traffic.</li> </ul>
Domain owner (technical lead)	<b>Mesh developer</b>	<p>The domain owner understands the internal and external dependencies for applications that interact with the service mesh and onboards developers.</p> <ul style="list-style-type: none"> <li>▶ Define environments for domain-based applications.</li> <li>▶ Define Istio configurations for the service mesh data plane.</li> <li>▶ Work with mesh operators to configure ingress and egress traffic and resources.</li> </ul>

Enterprise persona	Service mesh role	Responsibilities
Developer	<b>Mesh application viewer</b> (development environments)	<p>Developers use the configured platform, mesh, and observability stack to review and troubleshoot application functionality and performance.</p> <ul style="list-style-type: none"> <li>▶ Monitor the health, performance, and functional capabilities of applications.</li> <li>▶ Access Kiali visualizations, Jaeger telemetry, Prometheus metrics, and pod logs in the development environment.</li> </ul>
Application operator	<b>Mesh developer</b> (non-development environments)	<p>Application operators monitor and maintain applications deployed in the cluster and within the domain hosted mesh.</p> <ul style="list-style-type: none"> <li>▶ Access logs and envoy proxy configurations, telemetry, and traces to validate issues within their domain.</li> <li>▶ Extract information and work with mesh operators and developers to identify issues.</li> <li>▶ Suggest Istio configurations to mesh operators and developers.</li> </ul>
Product owner	<b>Mesh application viewer</b> (non-development environments)	<p>Product owners monitor the applications that comprise their product, both within and outside of the mesh, using the observability stack.</p> <ul style="list-style-type: none"> <li>▶ Monitor product health, use, cost, and other metrics within their domain.</li> <li>▶ Access observability stack information for up to one week after collection.</li> </ul>

## Map enterprise users to service mesh roles

Once you've defined your personas, roles, and responsibilities, you can map them to your team members. Following best practices, we'll map users to personas, roles, and namespaces within our development and production environments. The following tables show the development and production environment roles for Travel by Keyboard. The tables also include the usernames and passwords to be used in the examples throughout this e-book. For convenience in our examples, each person's username is their name (all lowercase), and their password is the same as their username.

### Logins for examples

This e-book provides example code, scripts, and interfaces that must be used by specific personas. Login information (username and password) for each person in our example scenario are shown in Tables 2 and 3.

**Table 2. User mapping for development environment**

Name	Username/password	Enterprise persona	Service mesh role	Namespaces
Phillip	phillip/phillip	Platform admin	Cluster admin	dev-istio-system
Emma	emma/emma	Mesh operator	Mesh operator	dev-istio-system
Cristina	cristina/cristina	Travel portal domain owner (technical lead)	Mesh developer	dev-travel-portal, dev-travel-control
Farid	farid/farid	Travel services domain owner (technical lead)	Mesh developer	dev-travel-agency
John	john/john	Travel portal developer	Mesh application viewer	dev-travel-portal, dev-travel-control

Name	Username/password	Enterprise persona	Service mesh role	Namespaces
Mia	mia/mia	Travel services developer	Mesh application viewer	dev-travel-agency
Mus	mus/mus	Product owner	Mesh application viewer	dev-travel-portal, dev-travel-control, dev-travel-agency

**Table 3. User mapping for production environment**

Name	Username/password	Enterprise persona	Service mesh role	Namespaces
Phillip	phillip/phillip	Platform admin	Cluster admin	prod-istio-system
Emma	emma/emma	Mesh operator	Mesh operator	prod-istio-system
Cristina	cristina/cristina	Travel portal domain owner (technical lead)	Mesh developer	prod-travel-portal, prod-travel-control
Farid	farid/farid	Travel services domain owner (technical lead)	Mesh developer	prod-travel-agency
Craig	craig/craig	Application operator (platform team)	Mesh developer	prod-travel-portal, prod-travel-control, prod-travel-agency
Mus	mus/mus	Product owner	Mesh application viewer	prod-travel-portal, prod-travel-control, prod-travel-agency

## Determine your service mesh deployments

Finally, you need to identify the deployment components to be used within your service mesh. Table 4 shows the components and number of instances of each that will be used in Travel by Keyboard's development environments. We will define the final production setup as we move through the deployment process in the following chapters.

**Table 4. Development environment service mesh architecture components**

Component	Number of instances	Component	Number of instances
grafana	1	jaeger	1
istiod	1	kiali	1
istio-egressgateway	1	prometheus	1
istio-ingressgateway	1	wasm-cacher-client-side-tenant	1

## Chapter 2

# Development environment set up

In this chapter, we'll set up the development environment for the travel portal and travel agency teams. Specific roles need to perform each task. In this case, the platform admin, mesh operator, and domain owners will set up the environment. The product owner, domain owners, developers, and mesh operators will use the observability stack to verify the setup.

Example code and scripts are provided for each step throughout this and the following chapters. [Additional information and the scripts](#) themselves are included in the online resource repository. →

## Development environment set up

This section provides information on setting up the development environment. In our example, Phillip, Emma, Farid, and Cristina perform these tasks.

### Define your cluster login URL

The platform admin (as the cluster admin role), Phillip, is responsible for setting up the overall cluster and platform.

1. Define the cluster URL to allow users to log in:

```
export CLUSTER_API=<YOUR-CLUSTER-API-URL>
```

### Prepare service mesh operators, namespaces, roles, and users

The platform admin (as the cluster admin role), Phillip, is also responsible for setting up the operators, namespaces, and roles within the development environment.

1. Change to the development set up directory and log in as Phillip:

```
cd ossm-heading-to-production-and-day-2/scenario-2-dev-setup
./login-as.sh phillip
```

2. Add the Red Hat OpenShift Service Mesh operators to the cluster via the [Red Hat OpenShift marketplace](#):

```
../common-scripts/add-operators-subscriptions-sm.sh
```

3. Create the required development travel agency namespaces:

```
../common-scripts/create-travel-agency-namespaces.sh dev
```

4. Create the [service mesh roles](#) as defined in Chapter 1.

5. Create the [development environment service mesh users](#) and assign roles as defined in Chapter 1.

## Create the service mesh control namespace

The mesh operator, Emma, is responsible for creating the control plane namespace and `ServiceMeshControlPlane` (or `smcp`) resource within the development environment.

1. Create the `dev-basic` service mesh control plane in the development environment:

```
./login-as.sh emma
./scripts/create-dev-smcp.sh dev-istio-system dev-basic
```

## Configure service mesh membership and deploy applications

The domain owners (as mesh developer roles), Farid and Cristina, are responsible for enrolling namespaces as members in the service mesh and for deploying applications within the development environment.

### Travel services domain (`dev-travel-agency` namespaces)

Farid owns the travel services domain and deploys applications to it as a mesh developer.

1. Check project labels prior to service mesh membership configuration:

```
./login-as.sh farid
../common-scripts/check-project-labels.sh dev-travel-agency
```

2. Add membership for the `dev-travel-agency` namespace to the `dev-basic` service mesh control plane by adding a `ServiceMeshMember` resource:

```
../common-scripts/create-membership.sh dev-istio-system dev-basic dev-travel-agency
../common-scripts/check-project-labels.sh dev-travel-agency
```

3. Deploy applications to the `dev-travel-agency` namespaces:

```
./scripts/deploy-travel-services-domain.sh dev dev-istio-system
```

### Travel portal domain (`dev-travel-portal` and `dev-travel-control` namespaces)

Cristina owns the travel portal domain and deploys applications to it as a mesh developer.

1. Check project labels prior to service mesh membership configuration:

```
./login-as.sh cristina
../common-scripts/check-project-labels.sh dev-travel-control
../common-scripts/check-project-labels.sh dev-travel-portal
```

2. Add membership for the `dev-travel-control` and `dev-travel-portal` namespaces to the `dev-basic` service mesh control plane by adding a `ServiceMeshMember` resource for each:

```
../common-scripts/create-membership.sh dev-istio-system dev-basic dev-travel-control
../common-scripts/check-project-labels.sh dev-travel-control
../common-scripts/create-membership.sh dev-istio-system dev-basic dev-travel-portal
../common-scripts/check-project-labels.sh dev-travel-portal
```

3. Deploy applications to the `dev-travel-control` and `dev-travel-portal` namespaces:

```
./scripts/deploy-travel-portal-domain.sh dev dev-istio-system
```

## Create a gateway for connections to the travel agency portal

The mesh operator, Emma, is responsible for creating the Istio Gateway resource to expose connections to the travel agency portal.

1. Create the Istio Gateway resource:

```
./login-as.sh emma
./scripts/create-ingress-gateway.sh dev-istio-system
```

## Development observability stack use and setup verification

This section provides information on using the observability stack to verify the setup of the development environment. It also provides information on how each role can use the observability stack. [Exercises for using the observability stack](#) as each role, as well as screenshots of expected outcomes, are detailed in the online resource repository. →

### Access the travel control dashboard

The cluster admin, mesh operator, and mesh developer roles – Phillip, Emma, Cristina, and Farid in our example – are responsible for capturing the URL of the travel control dashboard:

```
./login-as.sh <choose user>
echo "http://$(oc get route istio-ingressgateway -o jsonpath='{.spec.host}' -n dev-istio-system)"
```

You can then enter the URL in a web browser to access the travel control dashboard.

### Access the observability stack

The cluster admin, mesh operator, and mesh developer roles – Phillip, Emma, Cristina, and Farid in our example – can capture the URLs of the observability stack, including the Kiali, Jaeger, Prometheus, and Grafana components:

```
./login-as.sh <choose user>
echo "http://$(oc get route kiali -o jsonpath='{.spec.host}' -n dev-istio-system)"
echo "https://$(oc get route jaeger -o jsonpath='{.spec.host}' -n dev-istio-system)"
echo "https://$(oc get route prometheus -o jsonpath='{.spec.host}' -n dev-istio-system)"
echo "https://$(oc get route grafana -o jsonpath='{.spec.host}' -n dev-istio-system)"
```

You can also access the Grafana and Jaeger URLs from the Kiali dashboard by clicking on the question mark to the left of your name at the top of the window, and then selecting *About* from the pop-up menu. See the [interface](#). →

### Use the observability stack as the product owner

The product owner (as a mesh application viewer), Mus, has access to all three data plane namespaces (`dev-travel-portal`, `dev-travel-control`, and `dev-travel-agency`) and the control plane namespace via the Kiali dashboard. See the [Kiali dashboard](#) for Mus. →

As the product owner mesh application viewer, you can:

- ▶ **View traces for the overall solution.** Select *Distributed tracing* from the Kiali menu at the left side of the dashboard and log in with your credentials to access the tracing console.
- ▶ **View metrics for the overall solution.** Select an application workload from the *Workloads* section to see inbound and outbound metrics. You can also log in directly to the Prometheus dashboard using the URL from the previous section. You can apply the following commands to the graph view:
  - ▶ `istio_requests_total{destination_workload="discounts-v1", \app="discounts"}`
  - ▶ `istio_request_duration_milliseconds_count{app="discounts"}`
  - ▶ `istio_response_bytes_bucket`
- ▶ **View Grafana dashboards for the overall solution.** Log in to the Grafana URL using your credentials. You can see the status and performance of the travel agency solution by selecting *Dashboards > Manage > Istio > Istio Mesh Dashboard*. See the [Istio Mesh dashboard](#) and [performance view](#). →

As the product owner, you are allowed to view, but not modify, Istio configurations. You are not allowed to view Istio logs. You can verify that your role-based access controls are configured correctly by trying to view a configuration or log using the *Istio Config* and *Workloads* options from the left-side menu.

## Use the observability stack as the domain owner

The domain owners (as mesh developers), Cristina and Farid, have access to their respective domain namespaces. In our example, Cristina can access two data plane namespaces (`dev-travel-control` and `dev-travel-portal`) in the Travel Portal domain and the control plane namespace. Farid can access one data plane namespace (`dev-travel-agency`) in the Travel Services domain and the control plane namespace. See the Kiali dashboards for [Cristina](#) and [Farid](#). →

As a domain owner mesh developer, you can:

- ▶ **View traces for the overall solution.** Select *Distributed tracing* from the Kiali menu at the left side of the dashboard and log in with your credentials to access the tracing console.
- ▶ **View metrics for the overall solution.** Log in to the Prometheus dashboard using the URL from the previous section. You can apply the following commands to the graph view:
  - ▶ `istio_requests_total{destination_workload="discounts-v1", \app="discounts"}`
  - ▶ `istio_request_duration_milliseconds_count{app="discounts"}`
  - ▶ `istio_response_bytes_bucket`
- ▶ **View logs for the workloads in your domain.** Select a workload from the *Workloads* section in the left-side menu. The *Logs* tab shows both proxy and pod logs. See the [logs view](#). →
- ▶ **View and modify Istio configurations in your domain.** Select *Istio Config* from the left-side menu to see the configurations available to your role. See the [Istio config view](#). →
- ▶ **View Grafana dashboards.** Select *Dashboards > Manage > Istio > Istio Service Dashboard* or *Dashboards > Manage > Istio > Istio Workloads Dashboard* to check the status of services or workloads, respectively, in your domain. See the [Istio service](#) and [Istio workloads dashboards](#). →

## Use the observability stack as a developer

The developers (as mesh application viewers), Mia and John, have access to tracing, metrics, and dashboards within permitted namespaces. In our example, Mia is a developer for the Travel Services domain and can access information about the `dev-travel-agency` namespace. As a developer for the Travel Portal domain, John can access information about the `dev-travel-control` and `dev-travel-portal` namespaces. See the Kiali dashboards for [Mia](#) and [John](#). →

As a developer mesh application viewer, you can:

- ▶ **View traces for workloads in your domain.** Select *Distributed tracing* from the Kiali menu at the left side of the dashboard and log in with your credentials to access the tracing console.
- ▶ **View metrics for workloads in your domain.** Select an application workload from the *Workloads* section to see inbound and outbound metrics.
- ▶ **View Grafana dashboards for the overall solution.** Log in to the Grafana URL using your credentials. You can see the status and performance of the travel agency solution by selecting *Dashboards > Manage > Istio > Istio Mesh Dashboard*.

As a developer, you are allowed to view, but not modify, Istio configurations. You are not allowed to view Istio logs. You can verify that your role-based access controls are configured correctly by trying to view configuration details or logs using the *Istio Config* and *Workloads* options from the left-side menu.

## Use the observability stack as the mesh operator

The mesh operator, Emma, has full access to all three data plane namespaces (`dev-travel-control`, `dev-travel-portal`, and `dev-travel-agency`) and the control plane namespace. See the [Kiali dashboard](#) for Emma. →

As the mesh operator, you can:

- ▶ **View all namespaces in Kiali.** Select *Graphs > App Graph > Display* from the Kiali menu to show request distributions, namespace boxes, traffic animation, and security settings.
- ▶ **View logs for all workloads.** Select a workload from the *Workloads* option in the left-side menu. The *Logs* tab shows both proxy and pod logs.
- ▶ **View and modify Istio configurations.** Select *Istio Config* from the left-side menu to access and modify any configuration.
- ▶ **View dashboards.** Access Prometheus, Jaeger, and Grafana dashboards via the URLs defined earlier in this chapter. Select *Dashboards > Manage > Istio > Istio Control Plane Dashboard* to visualize the state of the service mesh control plane. See the [Istio control plane dashboard](#). →

## Chapter 3

# Production environment set up

In this chapter, we'll set up the production environment for the travel portal and travel agency teams according to the requirements shown in Chapter 1. As with the development environment, specific roles need to perform each task. In this case, the platform admin and mesh operator will set up and configure the production environment. Other enterprise personas and service mesh roles will have access to specified parts of the observability stack as before.

## Production environment set up

This section describes how to set up the production environment. In our example, Phillip and Emma perform these tasks.

### Define your cluster login URL

As with the development environment, the platform admin (as the cluster admin role), Phillip, is responsible for setting up the overall production cluster and platform.

1. Define the cluster URL to allow users to log in:

```
export CLUSTER_API=<YOUR-CLUSTER-API-URL>
```

### Prepare service mesh operators, namespaces, roles, and users

The platform admin (as the cluster admin role), Phillip, is responsible for setting up the operators, namespaces, and roles within the development environment. If you do not have separate Red Hat OpenShift clusters for development and production, you do not need to add the Red Hat OpenShift Service Mesh operators (step 2). Refer to the [Map enterprise users to service mesh roles](#) section in Chapter 1 for details.

1. Change to the production set up directory and log in as Phillip:

```
cd ossm-heading-to-production-and-day-2/scenario-2-dev-setup
./login-as.sh phillip
```

**IMPORTANT:** Skip steps 2 and 3 if you do not have separate development and production clusters.

2. Add the Red Hat OpenShift Service Mesh operators to the cluster via the [Red Hat OpenShift marketplace](#):

```
../common-scripts/add-operators-subscriptions-sm.sh
```

3. Create the [production service mesh roles](#) as defined in Chapter 1.

4. Create the required production travel agency namespaces:

```
../common-scripts/create-travel-agency-namespaces.sh prod
```

5. Create the [production environment service mesh users](#) and assign roles as defined in Chapter 1.

## Configure Jaeger, Elasticsearch, and the service mesh control plane for production

The mesh operator, Emma, is responsible for configuring the service mesh control plane and tracing within the production environment. There are two options for performing these tasks:

- ▶ Option 1: [Minimal Elasticsearch parameters via the control plane](#)
- ▶ Option 2: [Fully customized Elasticsearch parameters via Jaeger configuration](#)

Our example will use option 2 to fully customize the service mesh control plane and tracing deployment. Refer to the following Jaeger documentation for additional information about configuring an external Jaeger resource:

- ▶ [Understanding custom resource definitions](#)
- ▶ [Command line interface \(CLI\) flags for jaeger-collector with elasticsearch storage](#)

**IMPORTANT:** OpenSSL must be installed and accessible in the path to use the example scripts.

1. Configure Jaeger:

```
./login-as.sh emma
./scripts/create-prod-smcp-1-tracing.sh prod-istio-system production
```

This script creates a Jaeger resource in the `prod-istio-system` namespace with the following settings:

- ▶ [Production-focused setup](#)
- ▶ [Elasticsearch for persistent storage](#)
- ▶ [Indexes deleted on a rolling basis after 7 days](#)
- ▶ [One Elasticsearch node](#)
- ▶ [1Gi Elasticsearch index size](#)
- ▶ [Node resources are requested and limited](#)
- ▶ [No Elasticsearch node redundancy](#)

2. Confirm that the Jaeger resource was created:

```
oc get jaeger/jaeger-small-production -n prod-istio-system
```

3. Confirm that all components – a Jaeger collector, a Jaeger query, and an Elasticsearch deployment – were created:

```
oc get deployment -n prod-istio-system
```

NAME	READY	UP-TO-DATE	AVAILABLE	AGE
elasticsearch-cdm-prodistiosystemjaegersmallproduction-1	1/1	1	1	7m27s
jaeger-small-production-collector	1/1	1	1	7m25s
jaeger-small-production-query	1/1	1	1	7m25s

The `create-prod-smcp-1-tracing.sh` script also configures the service mesh according to the [production requirements](#) defined in Chapter 1. See the [complete YAML file](#). →

- ▶ [20% of all traces](#) collected for sampling
- ▶ [No external outgoing communications](#) to unregistered hosts allowed
- ▶ Traces stored in a Jaeger resource within the service mesh
- ▶ Elasticsearch provides persistence for traces
- ▶ Integration and use of an [external Jaeger resource](#)

## Configure the data plane and deploy applications into production

The domain owners (as mesh developer roles), Farid and Cristina, are responsible for configuring the data plane and deploying applications into production. In our example, the service mesh will contain:

- ▶ An `istio-proxy` sidecar container, for intercepting all communications into and out of the main application container and for applying service mesh configurations.
- ▶ A `jaeger-agent` sidecar container, to allow multitenancy in our Red Hat OpenShift cluster. See the [Jaeger deployment best practices](#) documentation for more details.

The `one-step-add-prod-deployments.sh` script allows Farid and Cristina to deploy applications to the travel agency and travel services domains:

```
./scripts/one-step-add-prod-deployments.sh <OCP CLUSTER DOMAIN eg. apps.example.com>
```

If you prefer to deploy applications into production in a step-by-step manner, refer to the [Setting up the production environment](#) section in the online resource repository.

## Configure Istio gateway access and security certificates

The mesh operator, Emma, is responsible for configuring an Istio Gateway resource and security certificates to allow access to the travel agency dashboard. Our example uses a Route resource to expose the travel agency dashboard over transport layer security (TLS). The certificate is hosted via the `control-gateway` resource in the ingress gateway, and the `route` resource is set to pass-through mode. This allows a separate certificate to be defined for each exposed service, but also requires additional maintenance, as each certificate must also be rotated separately by the mesh operator role.

Create the Route and Gateway resources and the certificates:

```
./login-as.sh emma
./scripts/create-https-ingress-gateway.sh prod-istio-system /
  <OCP CLUSTER DOMAIN eg. apps.example.com>
```

See the resulting [travel control dashboard](#). →

## Configure Prometheus for production

The mesh operator, Emma, is responsible for configuring Prometheus monitoring for the production environment. While Red Hat OpenShift offers an observability stack, the current version of Red Hat OpenShift Service Mesh does not integrate or federate with it. There are several options for getting around this:

- ▶ **Option 1:** Enhance the existing Prometheus deployment by adding a persistent volume
- ▶ **Option 2:** [Create an external Prometheus deployment using the Prometheus operator](#)
- ▶ **Option 3:** Collect data plane metrics only via integration with the Red Hat OpenShift monitoring stack
- ▶ **Option 4:** Configure and use an external monitoring tool to collect metrics

Our example will use option 1 to extend metric retention to seven days and add a persistent volume for long-term metric storage:

```
./login-as.sh emma
./scripts/update-prod-smcp-2-option1-prometheus.sh prod-istio-system
```

[Additional information](#) for the other three options is located in the online resource repository. →

## Final service mesh production setup

This section provides the final production setup, including tracing, metrics, scaled control plane components, and runtime resource allocations. The objectives, principles, and requirements defined in Chapter 1 should guide the final setup and serve as a starting point for establishing general rules and guidelines for the deployment and use of service mesh IT resources and assets. In our example, the travel agency architects have revisited and finalized the following service mesh objectives and architecture principles:

### Service mesh objectives

- ▶ Secure service-to-service communications.
- ▶ Monitor use and health of service-to-service communications.
- ▶ Allow teams to work separately to deliver parts of a solution.

### Service mesh architecture principles

- ▶ Use an external configuration mechanism for traffic encryption, authentication, and authorization.
- ▶ Transparently integrate additional services to expand functionality.
- ▶ Use an external traffic management and orchestration mechanism.
- ▶ Configure all components with high availability in mind.
- ▶ Use observability capabilities to verify system operation, rather than auditing.

### Final production service mesh setup

Based on these purposes and principles, here is our final production service mesh setup:

- ▶ **Tracing:** **5% of all traces** will be sampled and stored for seven days in an Elasticsearch cluster for debug purposes.
- ▶ **Metrics:** Collected metrics will be stored for seven days, and archived afterwards to allow for historical comparisons.
- ▶ **Grafana:** Grafana will use **persistent storage** for analytics.
- ▶ **Ingress and egress:** **Two ingress and two egress pod instances** will be deployed for high availability.
- ▶ **Istiod:** **Two Istiod instances** will be deployed for high availability.

The mesh operator, Emma, is responsible for updating the production deployment to meet the finalized requirements:

```
./login-as.sh emma
./scripts/update-prod-smcp-3-final.sh prod-istio-system production
```

## Production observability stack use and setup verification

As in the development environment, you should verify the setup of the production environment using the observability stack. Refer to the [Development observability stack use and setup verification](#) section in Chapter 2 for details about verifying the setup and how each role can use the stack. Chapter 4 will provide details about using the observability stack to tune service mesh performance and size service mesh components accordingly.

## Chapter 4

# Onboarding an external travel portal

In this chapter, we'll connect an external travel portal to Travel by Keyboard's environment to address a new business opportunity.

## Define requirements for the new opportunity

Travel by Keyboard's business development team has entered into a business deal with an external travel organization, Global Travel Organization (GTO), that will source offers from Travel by Keyboard's travel services domain. The following technical requirements will apply:

- ▶ The GTO travel portal will connect to Travel by Keyboard's travel services domain via APIs.
- ▶ All communications with the external client will be performed over mTLS.
- ▶ Authentication using a valid JSON web token (JWT) will be used to support additional authorization policies in the future.

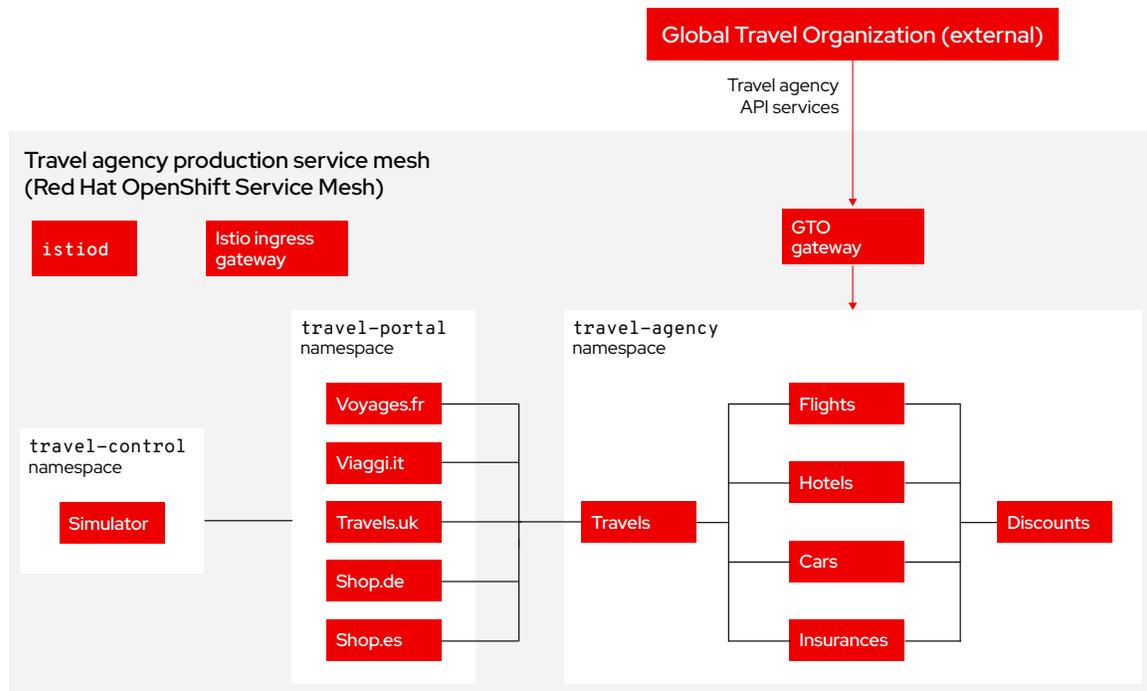


Figure 3. Travel by Keyboard's target service mesh architecture with connections to GTO's travel portal via APIs

## Create an additional ingress for API services

The mesh operator, Emma, is responsible for creating the additional Istio ingress gateway and certificates to allow external access to the travel agency services via mTLS.

1. Define the cluster URL to allow users to log in:

```
export CLUSTER_API=<YOUR-CLUSTER-API-URL>
```

2. Create the Istio Route resource, CA root keys and certificates, client/server certificates, and Gateway resource:

```
cd ossm-heading-to-production-and-day-2/scenario-4-onboard-new-portal-with-authentication
./login-as.sh emma
```

```
# Add to SMCP production resource in prod-istio-system
```

```
additionalIngress:
  gto-external-ingressgateway:
    enabled: true
    runtime:
      deployment:
        autoScaling:
          enabled: false
    service:
      metadata:
        labels:
          app: gto-external-ingressgateway
      selector:
        app: gto-external-ingressgateway
```

```
./scripts/create-external-mtls-https-ingress-gateway.sh prod-istio-system \
  <OCP CLUSTER DOMAIN e.g. apps.example.com>
```

This creates a new `gto-external-ingressgateway` pod, a TLS pass-through Route resource for access, and the `travel-api-gateway` Istio configuration. The `travel-api-gateway` Istio configuration defines the mTLS requirement and secrets that contain the needed certificates and keys.

See the script output: [Routes](#), [pods](#), [Istio configuration](#), and [YAML](#) →

## Deploy the new Istio configuration

The travel services domain owner (as a mesh developer role), Farid, is responsible for deploying the new `travel-api` `VirtualService` Istio configuration into the `prod-travel-agency` namespace to allow requests from the new gateway to reach the `cars`, `insurances`, `flights`, `hotels`, and `travels` services in the service mesh:

```
./login-as.sh farid
./scripts/create-client-certs-keys.sh curl-client
./scripts/deploy-external-travel-api-mtls-vs.sh prod prod-istio-system
curl -v -X GET --cacert ca-root.crt --key curl-client.key --cert curl-client.crt \
  https://gto-external-prod-istio-system.apps.<CLUSTERNAME>.<DOMAINNAME>/flights/Tallinn
```

This creates a mTLS handshake between the GTO client and the travel agency APIs via the newly created `gto-external-ingressgateway`.

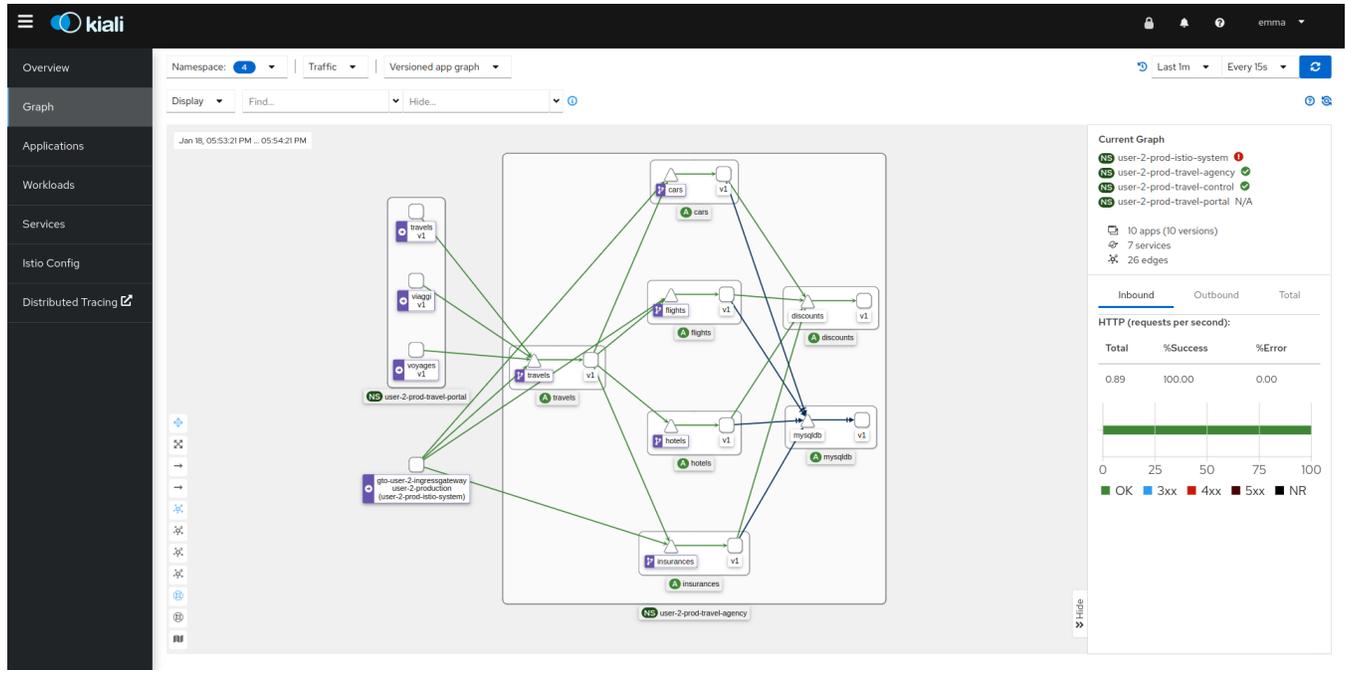


Figure 4. Kiali visualization of the new Istio configuration

## Enable JWT authentication

The intended final authentication workflow for GTO requests will use both mTLS handshakes and JWT:

1. The user authenticates with **Red Hat Single Sign-On** and receives a JWT.
2. The user initiates a request to `https://<route>/<service>` and passes the JWT with the request. Services can be `cars`, `insurances`, `flights`, `hotels`, or `travels`.
3. The `istio-proxy` container of the `gto-external-ingressgateway` pod checks the validity of the JWT token as defined by the `RequestAuthentication` object and permissions defined in the `AuthorizationPolicy` object.
4. If the JWT is valid, the user can access the path. Otherwise, an error message is returned to the user.

This approach is simple – only two custom resources (CRs) need to be deployed – and offers detailed authorization based on JWT fields. However, it does not use an OpenID Connect workflow, and the user is required to obtain and pass on the JWT themselves. You also need to define `RequestAuthentication` and `AuthorizationPolicy` objects for each protected application within the service mesh.

## Set up Red Hat Single Sign-On

The platform admin (as the cluster admin role), Phillip, is responsible for setting up the Red Hat Single Sign-On server and mounting the TLS certificate to `istiod`.

1. Configure the prerequisites for the single sign-on server:

```
./login-as.sh phillip
./prerequisites-setup.sh <CLUSTERNAME> <BASEDOMAIN> (eg. for apps.ocp4.example.com \
prerequisites-setup.sh ocp4 example.com)
```

2. Mount the Red Hat Single Sign-on TLS certificate to istiod:

```
./scripts/mount-rhssso-cert-to-istiod.sh prod-istio-system production <CLUSTERNAME> <BASEDOMAIN>
```

## Configure authentication and authorization for the gateway

The mesh operator, Emma, is responsible for creating the RequestAuthentication and AuthorizationPolicy objects to allow access to `gto-external-ingressgateway`. The RequestAuthentication object will require authentication and authorization requests to use only JWTs issued by Red Hat Single Sign-On, while the AuthorizationPolicy object will require JWTs to be issued by Red Hat Single Sign-On to be valid.

```
./login-as.sh emma
oc -n prod-istio-system apply -f approach_1/yaml/istio/jwt/01_requestauthentication.yaml
oc -n prod-istio-system apply -f \
  approach_1/yaml/istio/jwt/02_authpolicy_allow_from_servicemesh-lab_realm.yaml
```

This requires all requests coming in via the `gto-external` route and `gto-external-ingressgateway` to contain both a TLS certificate and a valid JWT to be allowed to proceed.

## Test travel agency API access using a JWT

Once you have set up authentication and authorization using JWTs, you should test to be sure everything is working as expected.

1. Ensure that you can no longer access the travel agency services without a valid JWT:

```
export GATEWAY_URL=$(oc -n prod-istio-system get route gto-external -o jsonpath='{.spec.host}')
curl -v -X GET --cacert ca-root.crt --key curl-client.key --cert curl-client.crt \
  https://$GATEWAY_URL/cars/Tallinn |jq
curl -v -X GET --cacert ca-root.crt --key curl-client.key --cert curl-client.crt \
  https://$GATEWAY_URL/travels/Tallinn |jq
curl -v -X GET --cacert ca-root.crt --key curl-client.key --cert curl-client.crt \
  https://$GATEWAY_URL/flights/Tallinn |jq
curl -v -X GET --cacert ca-root.crt --key curl-client.key --cert curl-client.crt \
  https://$GATEWAY_URL/insurances/Tallinn |jq
curl -v -X GET --cacert ca-root.crt --key curl-client.key --cert curl-client.crt \
  https://$GATEWAY_URL/hotels/Tallinn |jq
```

```
HTTP/1.1 403 Forbidden
```

- Retrieve a JWT for user *gtouser*. Note that `bcd06d5bdd1dbaaf81853d10a66aeb989a38dd51` is the `CLIENT_SECRET` defined during the Red Hat Single Sign-On [client secret creation](#) in our example.

```
TOKEN=$(curl -Lk --data "username=gtouser&password=gtouser&grant_type=password&client_id=istio\
&client_secret=bcd06d5bdd1dbaaf81853d10a66aeb989a38dd51" \
https://keycloak-rhssso.apps.ocp4.rhlab.de/auth/realms/servicemesh-lab/protocol/openid-connect/token \
| jq .access_token)
```

```
echo $TOKEN
```

- Use the JWT to access the travel services APIs as *gtouser*:

```
./scripts/call-via-mtls-and-jwt-travel-agency-api.sh prod-istio-system gto-external $TOKEN
```

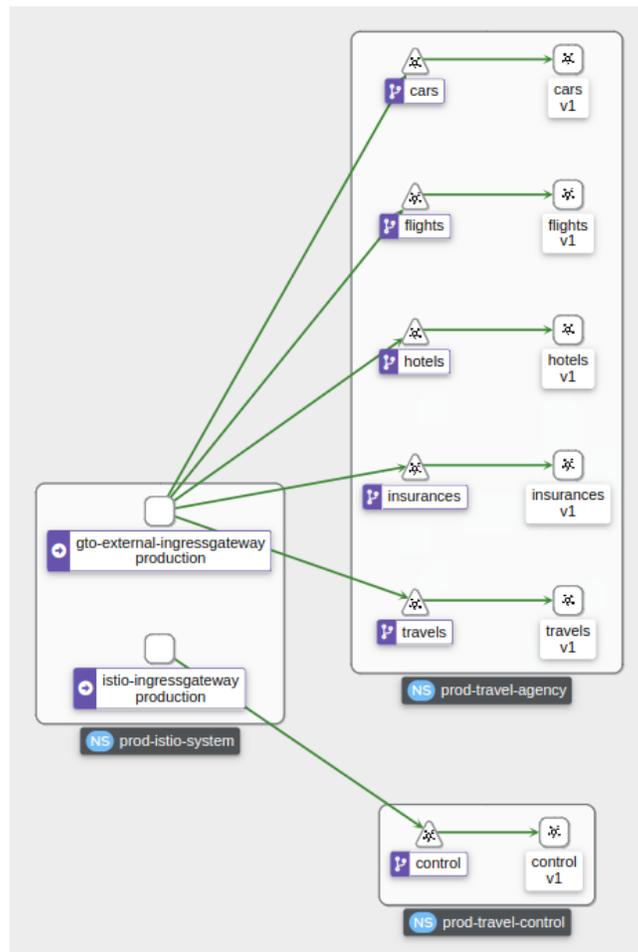


Figure 5. GTO users can access Travel by Keyboard's services via the travel services APIs.

## Chapter 5

# Complying with a new security regulation

In this chapter, we'll refine and enhance service mesh and environment security to address a new security regulation.

## Define requirements for the new regulation

A new security regulation has been released, and Travel by Keyboard's security and product teams have found that they'll need to make some changes to comply:

- ▶ A corporate (intermediate) certificate authority (CA) will be used to create mTLS certificates for intramesh communications.
- ▶ Additional authorization will be required to access specific services within the mesh.

## Implement mTLS across all services

In order to comply with the new security regulation, we will enhance the service mesh to generate and rotate intramesh mTLS certificates based on a corporate CA and an intermediate CA. In the following sections, we will create the certificates, keys, and certificate chain; add them to the production service mesh control plane resource; and validate the change.

### Create the corporate certificate authority and certificates

The platform admin (as the cluster admin role), Phillip, is responsible for creating the corporate CA and certificates required in the production environment. Note that your security team would usually provide the CA and certificates in an actual production scenario.

Change to the Chapter 5 directory and create the CA root, intermediate keys and certificates, and a chain certificate:

```
cd ossm-heading-to-production-and-day-2/scenario-5-new-regulations-mtls-everywhere
```

Follow the detailed instructions for [creating the root, keys, and certificates](#) in the resource repository. →

### Update the production service mesh control plane tenant

The mesh operator, Emma, is responsible for modifying the production service mesh control plane tenant to use the new corporate CA, intermediate CA, and chain certificates.

Start by verifying the issuer of the certificates currently used by the control plane and data plane. In each case, this should be something like `Issuer: 0 = cluster.local`.

1. Verify the issuer of the default service mesh certificate:

```
./login-as.sh emma
oc get -o yaml secret istio-ca-secret -n prod-istio-system | grep ca-cert | awk '{print $2}' | \
  base64 -d | openssl x509 -noout -text
```

2. Verify the issuer of the certificates used for communication with `istiod`:

```
oc exec "$(oc get pod -l app=istio-ingressgateway -n prod-istio-system -o \
  jsonpath={.items..metadata.name})" -c istio-proxy -n prod-istio-system -- openssl s_client \
  -showcerts -connect "$(oc get svc istiod-production -o jsonpath={.spec.clusterIP}):15012"
```

3. Verify the issuer of the certificates used in pod-to-pod communications:

```
oc exec "$(oc get pod -l app=hotels -n prod-travel-agency -o jsonpath={.items..metadata.name})" \
  -c istio-proxy -n prod-travel-agency -- openssl s_client -showcerts -connect \
  "$(oc get svc discounts -o jsonpath={.spec.clusterIP}):8000"
```

Next, create a secret named `cacerts` that contains the `ca-cert.pem`, `ca-key.pem`, `root-cert.pem`, and `cert-chain.pem` input files. Additional information can be found in the [Red Hat OpenShift documentation](#). →

4. Create the `cacerts` secret and add the input files:

```
oc create secret generic cacerts -n prod-istio-system \
  --from-file=ca-cert.pem=certs-resources/intermediate/certs/intermediate.cert.pem \
  --from-file=ca-key.pem=certs-resources/intermediate/private/intermediate.key.pem \
  --from-file=root-cert.pem=certs-resources/certs/ca.cert.pem \
  --from-file=cert-chain.pem=certs-resources/intermediate/certs/ca-chain.cert.pem
```

In our example, we will have:

- ▶ `intermediate.cert.pem` (`ca-cert.pem`), the certificate for the intermediate CA.
- ▶ `intermediate.key.pem` (`ca-key.pem`), the key for the intermediate CA certificate.
- ▶ `ca.cert.pem` (`root-cert.pem`), the root CA certificate.
- ▶ `ca-chain.cert.pem` (`cert-chain.pem`), the chain that includes both certificates.

The `istio-system-ca` secret created by the service mesh control plane by default can result in `istiod` choosing the incorrect enterprise certificates, so we'll remove it to avoid interference.

5. Remove the `istio-system-ca` secret:

```
oc get secret istio-ca-secret -n prod-istio-system -o yaml > istio-ca-secret-default.yaml
oc delete secret istio-ca-secret -n prod-istio-system
```

6. Add the `cacerts` secret to the production service mesh control plane resource in the `prod-istio-system` namespace:

```
# Add in production SMCP in prod-istio-system
security:
  certificateAuthority:
    type: Istiod
  istiod:
    type: PrivateKey
  privateKey:
    rootCADir: /etc/cacerts
```

Once you have added the new certificates, you need to restart the control plane and data plane resources to implement the changes.

- Restart the control plane `istiod` and gateway pods:

```
oc -n prod-istio-system delete pods \
  -l 'app in (istiod,istio-ingressgateway, istio-egressgateway,gto-external-ingressgateway)'
oc -n prod-istio-system get -w pods
```

- Restart the data plane pods:

```
oc -n prod-travel-control delete pods --all
oc -n prod-travel-agency delete pods --all
oc -n prod-travel-portal delete pods --all
```

- Verify that the cacerts issuer is Issuer: C = GB, ST = England, L = London, O = Travel Agency Ltd, OU = Travel Agency Ltd Certificate Authority, CN = Travel Agency Ltd Root CA:

```
oc get -o yaml secret cacerts -n prod-istio-system | grep ca-cert | awk '{print $2}' \
  | base64 -d | openssl x509 -noout -text
```

- Verify that the data plane communications are secured with the new corporate certificates:

```
./verify-dataplane-certs.sh
```

- Verify that the control plane communications are secured with the new corporate certificates:

```
./verify-controlplane-certs.sh
```

## Disable STRICT mTLS for specific services

Some workloads running in your service mesh may offer their own mTLS certificates, or may not support mTLS. You can configure your service mesh to handle these situations.

A [detailed exercise for disabling mTLS for the cars service](#) is included in the online repository. →

## Implement new authorization policies

We will also implement new authorization policies to align with the new security regulation. In our example, the mesh operator (Emma) and domain owners (Cristina and Farid) will restrict access to specific services based on best practices and business requirements.

[Detailed information](#) about authorization is available in the Istio documentation. →

### Check the default authorization policies

First, confirm that the default service mesh authorization policies allow all communications.

- Test the control `.prod-travel-control` communication:

```
https://travel-prod-istio-system.apps.<CLUSTERNAME>.<BASEDOMAIN>/
```

- Run the script to test the remaining communications:

```
./scripts/check-authz-all.sh ALLOW prod-istio-system <CLUSTERNAME> <BASEDOMAIN> \
  <CERTS_LOCATION> (CERTS_LOCATION ../scenario-4-onboard-new-portal-with-authentication)
```

The output should show that all communications are allowed by default:

```
Authorization prod-istio-system --> prod-travel-agency
-----
[ALLOW] gto-external-ingressgateway --> travels.prod-travel-agency
[ALLOW] gto-external-ingressgateway --> cars.prod-travel-agency
[ALLOW] gto-external-ingressgateway --> flights.prod-travel-agency
[ALLOW] gto-external-ingressgateway --> insurances.prod-travel-agency
[ALLOW] gto-external-ingressgateway --> hotels.prod-travel-agency

Authorization prod-travel-control --> prod-travel-agency
-----
[ALLOW] control.prod-travel-control --> travels.prod-travel-agency
[ALLOW] control.prod-travel-control --> cars.prod-travel-agency
[ALLOW] control.prod-travel-control --> flights.prod-travel-agency
[ALLOW] control.prod-travel-control --> insurances.prod-travel-agency
[ALLOW] control.prod-travel-control --> hotels.prod-travel-agency

Authorization prod-travel-portal --> prod-travel-agency
-----
[ALLOW] viaggi.prod-travel-portal --> travels.prod-travel-agency
[ALLOW] viaggi.prod-travel-portal --> cars.prod-travel-agency
[ALLOW] viaggi.prod-travel-portal --> flights.prod-travel-agency
[ALLOW] viaggi.prod-travel-portal --> insurances.prod-travel-agency
[ALLOW] viaggi.prod-travel-portal --> hotels.prod-travel-agency

Authorization prod-travel-agency --> prod-travel-agency
-----
[ALLOW] travels.prod-travel-portal --> discounts.prod-travel-agency
[ALLOW] travels.prod-travel-portal --> cars.prod-travel-agency
[ALLOW] travels.prod-travel-portal --> flights.prod-travel-agency
[ALLOW] travels.prod-travel-portal --> insurances.prod-travel-agency
[ALLOW] travels.prod-travel-portal --> hotels.prod-travel-agency
```

## Apply a deny-all policy

Best practices for service mesh security use a [default-deny pattern](#). In Chapter 4, we configured the `authpolicy-gto-external` resource to allow requests via the `gto-external-ingressgateway`. Now, we will deny all requests by default and allow only requests specified by an authorization policy. The domain owners (as mesh developer roles), Farid and Cristina, can implement these changes.

1. Apply the default-deny pattern to the `prod-travel-control` and `prod-travel-agency` namespaces:

```
cd ossm-heading-to-production-and-day-2/scenario-5-new-regulations-mtls-everywhere
./login-as.sh cristina
oc apply -f authz-resources/01-default-deny-travel-portal-access.yaml
./login-as.sh farid
oc apply -f authz-resources/01-default-deny-travel-agency-access.yaml
```

2. Verify that access is denied via web browser. The following link should result in a message saying `RBAC: access denied`.

```
https://travel-prod-istio-system.apps.<CLUSTERNAME>.<BASEDOMAIN>/
```

3. Run the script to check the remaining communications:

```
./scripts/check-authz-all.sh DENY prod-istio-system <CLUSTERNAME> <BASEDOMAIN> \
  <CERTS_LOCATION> (CERTS_LOCATION ../scenario-4-onboard-new-portal-with-authentication)
```

The output should show that all communications are denied by default:

```
Authorization prod-istio-system --> prod-travel-agency
-----
[DENY] gto-external-ingressgateway --> travels.prod-travel-agency
[DENY] gto-external-ingressgateway --> cars.prod-travel-agency
[DENY] gto-external-ingressgateway --> flights.prod-travel-agency
[DENY] gto-external-ingressgateway --> insurances.prod-travel-agency
[DENY] gto-external-ingressgateway --> hotels.prod-travel-agency

Authorization prod-travel-control --> prod-travel-agency
-----
[DENY] control.prod-travel-control --> travels.prod-travel-agency
[DENY] control.prod-travel-control --> cars.prod-travel-agency
[DENY] control.prod-travel-control --> flights.prod-travel-agency
[DENY] control.prod-travel-control --> insurances.prod-travel-agency
[DENY] control.prod-travel-control --> hotels.prod-travel-agency

Authorization prod-travel-portal --> prod-travel-agency
-----
[DENY] viaggi.prod-travel-portal --> travels.prod-travel-agency
[DENY] viaggi.prod-travel-portal --> cars.prod-travel-agency
[DENY] viaggi.prod-travel-portal --> flights.prod-travel-agency
[DENY] viaggi.prod-travel-portal --> insurances.prod-travel-agency
[DENY] viaggi.prod-travel-portal --> hotels.prod-travel-agency

Authorization prod-travel-agency --> prod-travel-agency
-----
[DENY] travels.prod-travel-portal --> discounts.prod-travel-agency
[DENY] travels.prod-travel-portal --> cars.prod-travel-agency
[DENY] travels.prod-travel-portal --> flights.prod-travel-agency
[DENY] travels.prod-travel-portal --> insurances.prod-travel-agency
[DENY] travels.prod-travel-portal --> hotels.prod-travel-agency
```

You can also [verify the communication policies visually](#) via Kiali. →

## Apply business-justified authorization policies

Now, we will allow access to specific services and requests based on business needs.

### Prod-travel-control namespace

The mesh operator, Emma, is responsible for allowing access to the service mesh via the default `istio-ingressgateway`:

```
./login-as.sh emma
oc apply -f authz-resources/02-travel-portal-allow-external-traffic.yaml
```

### Travel portal domain

The travel portal domain owner (as a mesh developer role), Cristina, is responsible for allowing access from the `istio-ingressgateway` to the travel portal. We will set the `istio-ingressgateway` to allow calls on any path, and the principal `cluster.local/ns/prod-istio-system/sa/istio-ingressgateway-service-account` will be allowed to make calls to the `prod-travel-control` namespace:

```
./login-as.sh cristina
oc apply -f authz-resources/02-travel-portal-allow-istio-ingressgateway-traffic.yaml
```

Access to the travel portal interface will be reinstated within a few seconds. As before, you can test the connection via `https://travel-prod-istio-system.apps.<CLUSTERNAME>.<BASEDOMAIN>/`.

### Travel services domain

The travel services domain owner (as a mesh developer role), Farid, is responsible for allowing access from the `gto-external-ingressgateway` to the `travels.prod-travel-agency`, `hotels.prod-travel-agency`, `cars.prod-travel-agency`, `insurances.prod-travel-agency`, and `flights.prod-travel-agency` services so that GTO can perform search requests. The previous section allowed calls to `gto-external-ingressgateway` on any path for any external principal (`requestPrincipals[*]`). In this section, we will allow only the principal `cluster.local/ns/prod-istio-system/sa/gto-external-ingressgateway-service-account` to make calls to the `prod-travel-agency` namespace.

1. Allow access for GTO:

```
./login-as.sh farid
oc apply -f authz-resources/03-gto-external-travels-to-travel-agency-allow.yaml
```

2. Check whether travel searches from GTO are allowed. You should receive an HTTP 500 error with "invalid connection" as the reason:

```
TOKEN=$(curl -Lk --data "username=gtouser&password=gtouser&grant_type=password&\
client_id=istio&client_secret=bcd06d5bdd1dbaaf81853d10a66aeb989a38dd51" \
https://keycloak-rhssso.apps.ocp4.rhlab.de/auth/realms/servicemesh-lab/protocol/openid-connect/token \
| jq .access_token)

./scenario-4-onboard-new-portal-with-authentication/scripts/call-via-mtls-and-jwt-travel-agency-api.sh \
prod-istio-system gto-external $TOKEN
```

3. Access from GTO is now authorized and there are no longer 403 errors, but connections are failing. Inspect the communications in Kiali. You should find that requests from `gto-external-ingressgateway` to `*.prod-travel-agency` are allowed, but that connections from the `flights` service to the `mysqldb` service are not allowed. You can check this in the `istio-proxy` logs for the `flights-v1` workload.
4. Allow communications within the `prod-travel-agency` namespace so that the principal `cluster.local/ns/prod-travel-agency/sa/default` is allowed to make calls to the `prod-travel-agency` namespace:

```
./login-as.sh farid
oc apply -f authz-resources/04-intra-travel-agency-allow.yaml
```

5. Verify that travel searches from GTO are now allowed:

```
TOKEN=$(curl -Lk --data "username=gtouser&password=gtouser&grant_type=password&\
client_id=istio&client_secret=bcd06d5bdd1dbaaf81853d10a66aeb989a38dd51" \
https://keycloak-rhso.apps.<CLUSTERNAME>.<BASEDOMAIN>/auth/realms/servicemesh-lab/protocol/openid-connect/token \
| jq .access_token)

./scenario-4-onboard-new-portal-with-authentication/scripts/call-via-mtls-and-jwt-travel-agency-api.sh \
prod-istio-system gto-external $TOKEN
```

6. Test the intra-namespace communications:

```
./scripts/check-authz-all.sh 'ALLOW intra' prod-istio-system <CLUSTERNAME> <BASEDOMAIN> \
<CERTS_LOCATION> (CERTS_LOCATION ../scenario-4-onboard-new-portal-with-authentication)
```

The output should show the updated communications authorizations:

```
Authorization prod-istio-system --> prod-travel-agency
-----
[ALLOW] gto-external-ingressgateway --> travels.prod-travel-agency
[ALLOW] gto-external-ingressgateway --> cars.prod-travel-agency
[ALLOW] gto-external-ingressgateway --> flights.prod-travel-agency
[ALLOW] gto-external-ingressgateway --> insurances.prod-travel-agency
[ALLOW] gto-external-ingressgateway --> hotels.prod-travel-agency

Authorization prod-travel-control --> prod-travel-agency
-----
[DENY] control.prod-travel-control --> travels.prod-travel-agency
[DENY] control.prod-travel-control --> cars.prod-travel-agency
[DENY] control.prod-travel-control --> flights.prod-travel-agency
[DENY] control.prod-travel-control --> insurances.prod-travel-agency
[DENY] control.prod-travel-control --> hotels.prod-travel-agency

Authorization prod-travel-portal --> prod-travel-agency
-----
[DENY] viaggi.prod-travel-portal --> travels.prod-travel-agency
[DENY] viaggi.prod-travel-portal --> cars.prod-travel-agency
[DENY] viaggi.prod-travel-portal --> flights.prod-travel-agency
[DENY] viaggi.prod-travel-portal --> insurances.prod-travel-agency
[DENY] viaggi.prod-travel-portal --> hotels.prod-travel-agency
```

```
Authorization prod-travel-agency --> prod-travel-agency
```

```
-----
[ALLOW] travels.prod-travel-portal --> discounts.prod-travel-agency
[ALLOW] travels.prod-travel-portal --> cars.prod-travel-agency
[ALLOW] travels.prod-travel-portal --> flights.prod-travel-agency
[ALLOW] travels.prod-travel-portal --> insurances.prod-travel-agency
[ALLOW] travels.prod-travel-portal --> hotels.prod-travel-agency
```

7. The domain owner (as a mesh developer role), Farid, can then allow the principal `cluster.local/ns/prod-travel-portal/sa/default` to make calls to the `prod-travel-agency` namespace:

```
./login-as.sh farid
oc apply -f authz-resources/05-travel-portal-to-travel-agency-allow.yaml
```

## Verify authorization policies

The authorization policies applied in the previous sections restored business-justified connections within our example environment.

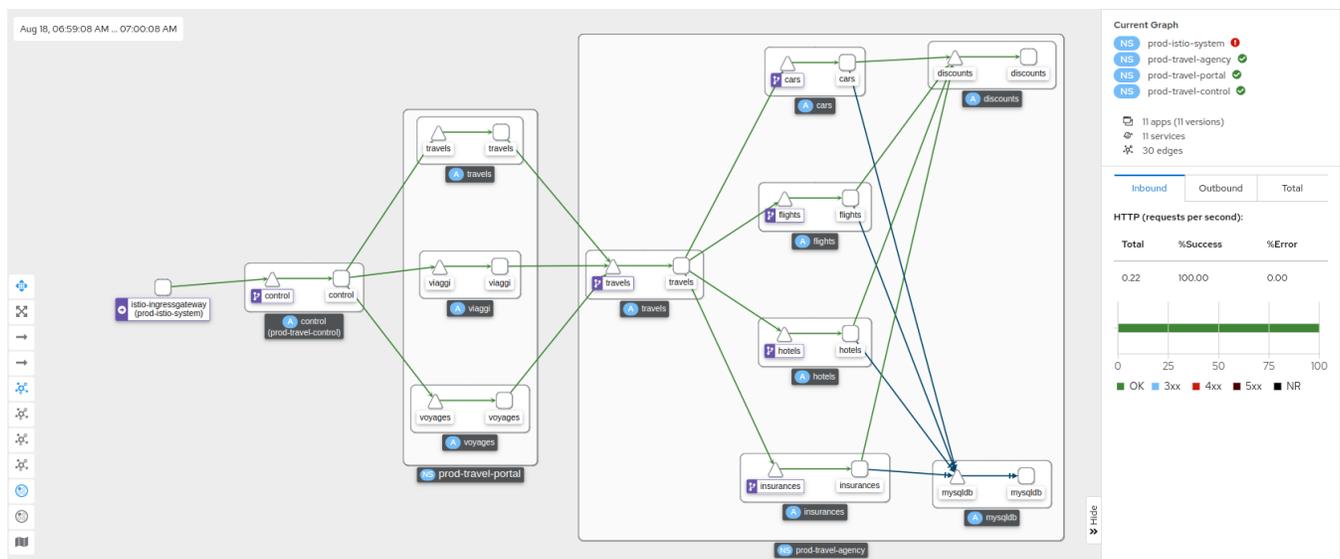


Figure 6. Kiali visualization of authorization policies, including business-justified connections

This configuration is different from the allow-all policy we started with in this chapter. The mesh operator, Emma, can verify that the `travel-portal-control` namespace still does not have access to the `travel-portal-agency` namespace:

```
./login-as.sh emma
./scripts/check-authz-all.sh 'ALLOW intra' prod-istio-system <CLUSTERNAME> <BASEDOMAIN> \
  <CERTS_LOCATION> (CERTS_LOCATION ../scenario-4-onboard-new-portal-with-authentication)
```

## Apply detailed authorization policies for the partner portal

In our example, Travel by Keyboard's business development team has finalized the business agreement with their partner, GTO. As a result, GTO will be able to source offers *only* from the `flights` and `insurances` services via the travel agency APIs. In this section, we'll configure the authorization policy to align with this business agreement. The mesh operator, Emma, is responsible for configuring the detailed authorization policies.

1. Deny communications from `gto-external-ingressgateway` to all paths *except* `/flights` and `/insurances` using the `notPaths` operation:

```
./login-as.sh emma
oc apply -f authz-resources/06-gto-external-travels-only-flights-insurances-paths-allow.yaml
```

2. Check that the correct connections are now allowed:

```
./scripts/verify-fine-grained-authz.sh prod-istio-system <CLUSTERNAME> <BASEDOMAIN> \
  <CERTS_LOCATION> (CERTS_LOCATION ../scenario-4-onboard-new-portal-with-authentication)
```

```
[DENY] GTO --> /travels
[DENY] GTO --> /cars
[ALLOW] GTO --> /flights
[ALLOW] GTO --> /insurances
[DENY] GTO --> /hotels
```

Figure 7 shows the final authentication and authorization configuration for our example.

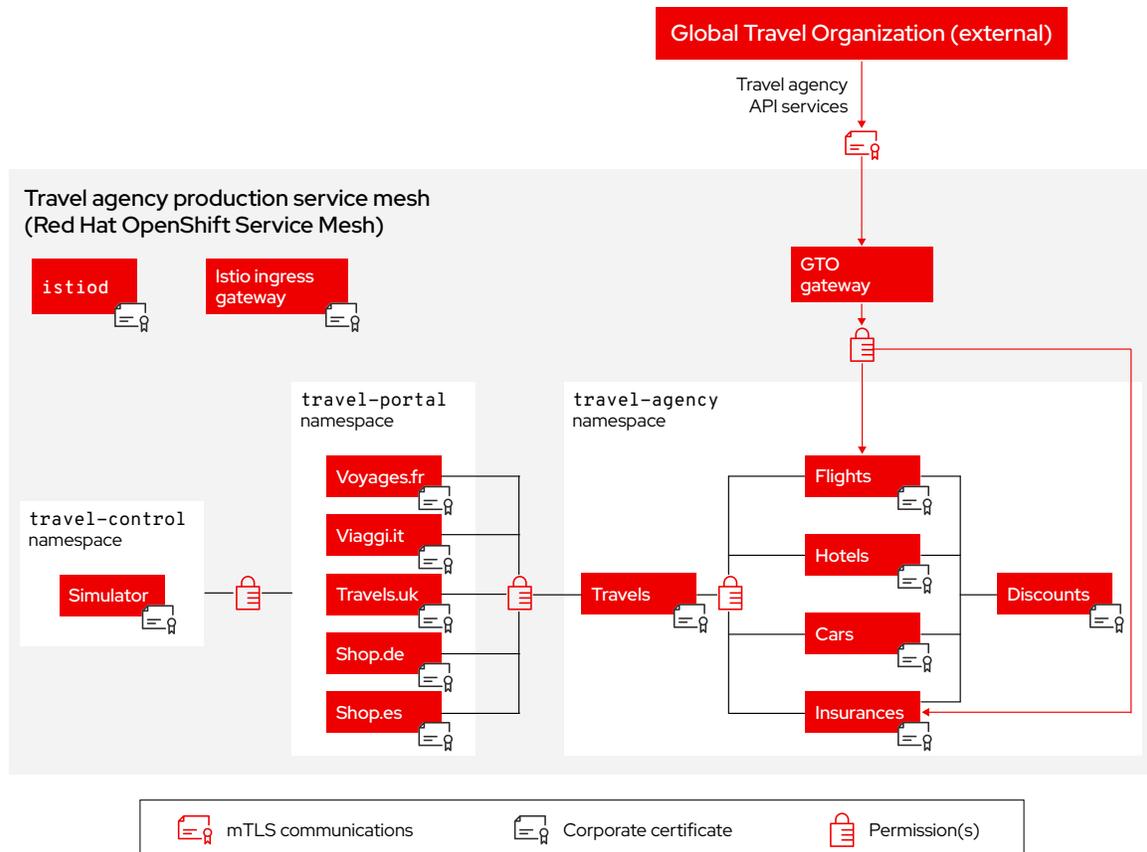


Figure 7. Final authentication and authorization configuration for the travel agency service mesh

## Chapter 6

# Expanding partnerships and implementing broker services

In this chapter, we'll onboard an insurance broker partner and forward certain insurance requests to that partner to offer premium insurance packages to customers traveling to popular destinations.

## Define requirements for the new opportunity

Travel by Keyboard's business development team has partnered with an external insurance broker, Super Insurance, to offer additional premium insurance packages to customers. As a result, certain insurance requests will be forward to the partner:

- ▶ All requests for insurance related to travels to the London, Rome, Paris, Berlin, Munich, and Dublin destinations will be forwarded to Super Insurance.
- ▶ All communications with the external partner services will be performed over mTLS.

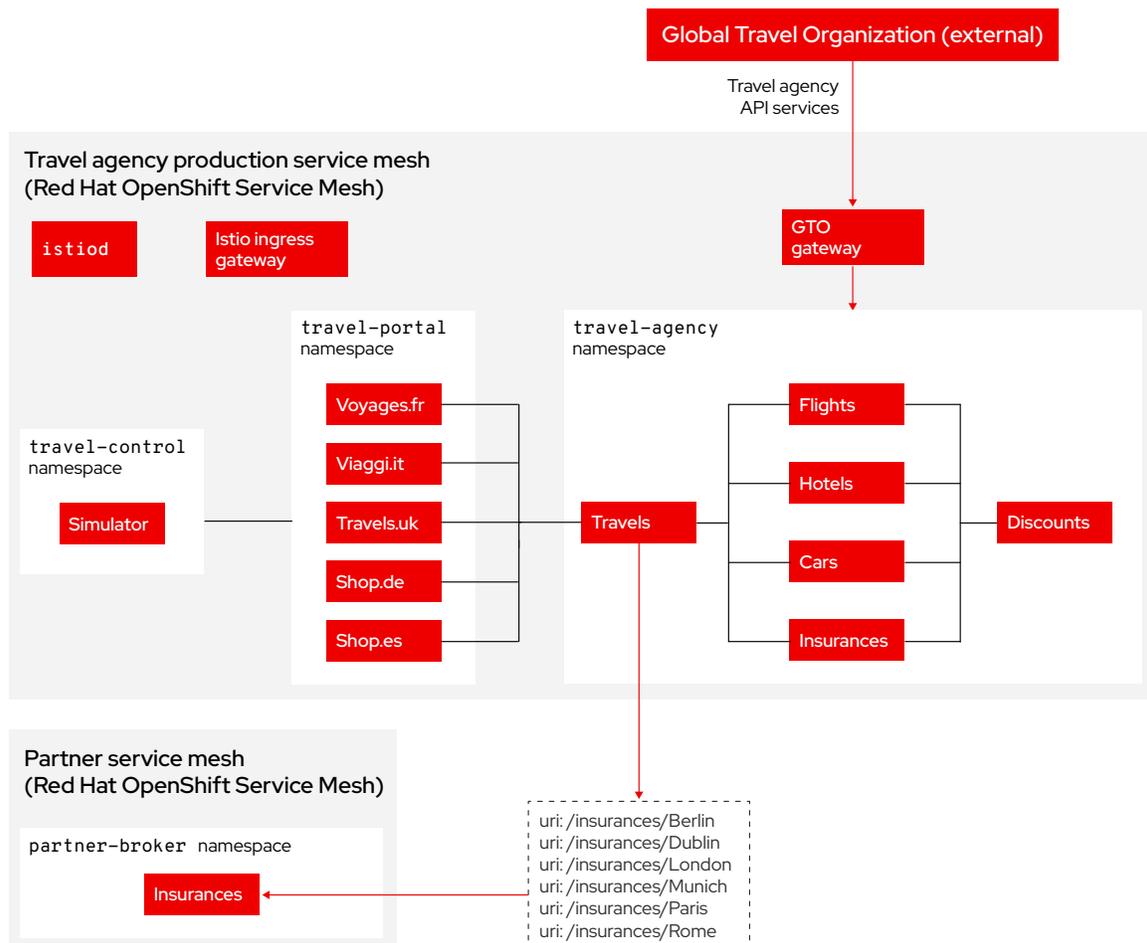


Figure 8. Travel agency service mesh architecture with connections to an external insurance broker partner system

## Deploy the partner insurance services

First, we need to deploy and configure the service mesh components within Super Insurance's system. We will perform these steps as the cluster admin role (Phillip), though the external organization should also have a set of enterprise personas with similar service mesh roles and responsibilities to those we have set up for Travel by Keyboard.

The cluster admin role will be responsible for deploying the required service mesh components, including:

- ▶ Deploying the service mesh control plane `partner` resource in the `partner-istio-system` namespace.
- ▶ Creating the `premium-broker` namespace and adding service mesh membership for the namespace to the partner's service mesh control plane.
- ▶ Deploying the insurance service in the `premium-broker` namespace.
- ▶ Enabling traffic routing to the insurance service in the `premium-broker` namespace.
- ▶ Verifying that traffic is routed to the partner insurance service.

In our example, we'll perform all of these tasks using scripts:

```
cd ossm-heading-to-production-and-day-2/scenario-6-partner-agency-multi-mesh
./login-as.sh phillip

./create-premium-insurance-broker.sh <PARTNER INSURANCE NAMESPACE> \
  <PARTNER ISTIO CP NAMESPACE> <CLUSTER.DOMAIN eg. apps.ocp4.rhlab.de> <PARTNER SMCP NAME>
./create-premium-insurance-broker.sh premium-broker partner-istio-system \
  <CLUSTER.DOMAIN eg. apps.ocp4.example.com> partner
```

## Route requests to the insurance broker partner

Next, we'll configure Travel by Keyboard's systems to separate and forward insurance requests for the premium destinations defined earlier to Super Insurance. There are three options for achieving this, none of which require changes to our applications:

- ▶ **Option 1: Non-mTLS external insurance service call**  
This option applies [VirtualService](#) and [DestinationRule](#) resources to route traffic to the remote service location and creates a `ServiceEntry` resource for the remote service destination location.
- ▶ **Option 2: Multitenancy using mTLS without federation**  
This option uses the egress gateway of the production service mesh to redirect calls to the remote `premium-broker/insurances` service, sharing and applying remote certificates at the gateway. An example of this setup is available in [this blog post](#).
- ▶ **Option 3: Federation between the production mesh and the partner mesh**  
This option federates the production and partner service mesh instances and imports the partner insurance service into the production service mesh. Important [considerations for planning service mesh federation](#) are included in the online resource repository. →

Our example will use option 3 to federate the two service meshes. To do so, we need to:

- ▶ Update the production service mesh control plane resource to declare two additional gateways – `partner-mesh-egress` and `partner-mesh-ingress` – for the federated connection.
- ▶ Update the partner service mesh control plane resource to declare two additional gateways – `production-mesh-egress` and `production-mesh-ingress` – for the federated connection.
- ▶ Extract the configuration map `istio-ca-root-cert` from each of the meshes and share it on the control plane namespace of the opposite mesh to support the TLS handshake.
- ▶ Create the `partner ServiceMeshPeer` resource in `prod-istio-system` to initiate the peering from the production mesh to the partner mesh.
- ▶ Create `production ServiceMeshPeer` resource in `partner-istio-system` to initiate the peering from the partner mesh to the production mesh.
- ▶ Export the `insurances` service (via `ExportedServiceSet`) from the `premium-broker` namespace and import it (via `ExportedServiceSet`) into the `prod-travel-agency` namespace.

In our example, the cluster admin, Phillip, will perform all of these tasks using a script:

```
./scripts/option-3-execute-federation-setup.sh <1_SMCP_NAMESPACE> <1_SMCP_NAME> \
  <2_SMCP_NAMESPACE> <2_SMCP_NAME> <PREMIUM_NAMESPACE>
./scripts/option-3-execute-federation-setup.sh prod-istio-system production partner-istio-system \
  partner premium-broker
```

The outcome of this configuration will be that insurance quotes requests arriving at the `prod-travel-agency/travels` service will be forwarded to the federated `insurances.premium-broker.svc.partner-imports.local` service.

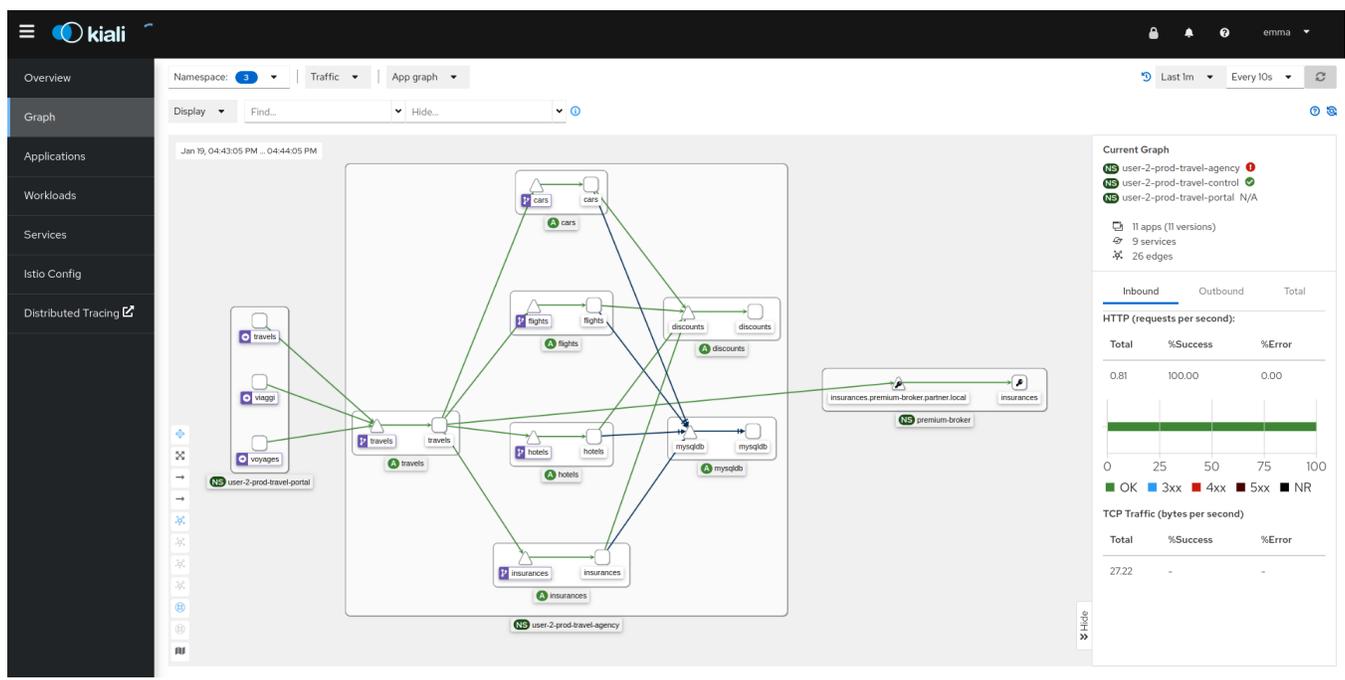


Figure 9. Production mesh federating requests to the imported service `insurances.premium-broker.svc.partner-imports.local`

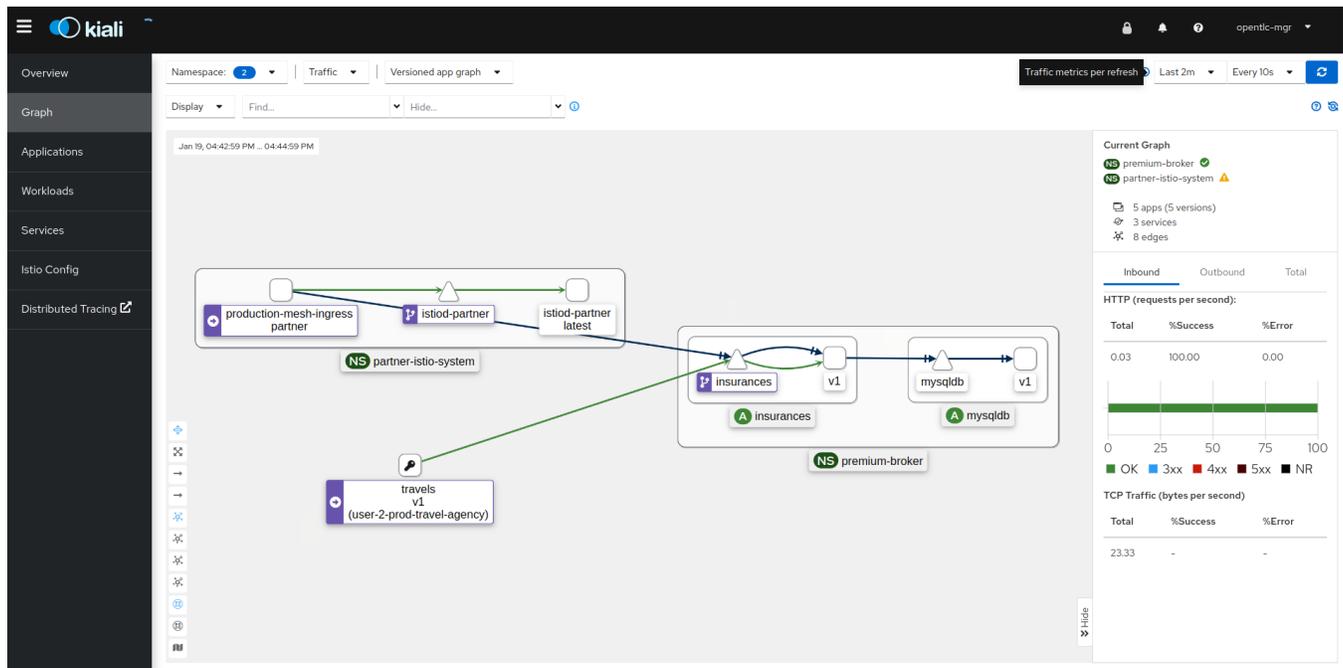


Figure 10. Partner mesh serving requests via exported service `insurances.premium-broker` to `travels-v1.prod-travel-agency` in the production mesh

The logs of the `insurances.premium-broker` pod should show that it serves only the specified premium destinations, while all others are served by the `insurances.prod-travel-agency` pod.

## Chapter resources

- ▶ Red Hat OpenShift documentation: [Exporting a service from a federated mesh](#)
- ▶ Blog post: [Security options with mTLS for egress edge traffic](#)
- ▶ Blog post: [Federation automated setup](#)

# Part 2: Day 2 operations

Once you have deployed your service mesh into production, you need to optimize, tune, and maintain it. Day 2 operations are critical for the ongoing efficiency, stability, and utility of any technology, and service meshes are no different. In Part 1 of this e-book, we laid the foundations for Day 2 operations by configuring the platform for security, adaptability, and observability. In Part 2, we'll explore methods and techniques for troubleshooting, tuning, and upgrading your service mesh.

## Chapter 7

# Troubleshooting the mesh

In this chapter, we'll explore several techniques for troubleshooting your service mesh.

## Key questions to ask

Each Istio network configuration debugging journey is different, but it helps to start by asking ten key questions:

1. Is the Istio network configuration syntactically valid?
2. Does the network configuration have an error or warning status set?
3. Is the resource name right? Is the resource in the right namespace?
4. Are the resource selectors correct?
5. Did Envoy accept (ACK) the configuration?
6. Did the configuration appear as expected in Envoy?
7. Did `istiod` (Pilot) log errors?
8. Are the Red Hat OpenShift Service Mesh operator and control plane healthy?
9. Is your application part of the expected mesh instance?
10. Are the relevant certificates valid?

The following sections provide guidance for gathering the information needed to answer each of these questions and, hopefully, resolve service mesh issues in the process.

## Understand the components of your mesh

The first step in nearly every debugging journey is knowing what you're dealing with. It's important to understand which **mesh components** you have, their purpose, their configurations, and from where those configurations are applied. You also need to understand the expected flow of traffic in your mesh and which tools your mesh operator and application operations team personas can use in debugging efforts. In our example, we'll focus on the following mesh components:

- ▶ **Ingress components** that allow traffic into the mesh. In our example, `istio-ingressgateway` and `gto-ingressgateway` carry north inbound client traffic and `partner-mesh-ingress` carries federated west inbound traffic.
- ▶ **Egress components** that allow traffic out of the mesh. In our example, `istio-egressgateway` carries south outbound traffic and `partner-mesh-egress` carries federated east outbound traffic.
- ▶ **istio-proxy sidecar containers** – one for each pod that is part of the mesh – that intercept TCP, HTTP, HTTP/2, and GRPC<sup>1</sup> traffic into and out of the main application containers.
- ▶ **istiod components** that apply the mesh configurations for traffic, security, observability, and more.

Two of the main actions of a service mesh are to apply Red Hat OpenShift network policies that govern and restrict access to the mesh and to rewrite the IP tables rules so that each pod can only be accessed via the service mesh `istio-proxy` sidecar. The `NetworkPolicy` resource can enforce mesh isolation and multitenancy to achieve this.

More information about network policies is available in the [Red Hat OpenShift documentation](#):

- ▶ [Multitenancy versus cluster-wide mesh installations](#)
- ▶ [Understanding mesh network policies](#)
- ▶ [Setting the correct network policy](#)

### Control plane network policies

In our example, the following `NetworkPolicy` resources have been applied in the `prod-istio-system` namespace:

**Table 5. NetworkPolicy resource application for the prod-istio-system namespace**

Name	Pod selector
<code>gto-external-ingressgateway</code>	<code>app=gto-external-ingressgateway,istio=ingressgateway</code>
<code>istio-expose-route-production</code>	<code>maistra.io/expose-route=true</code>
<code>istio-grafana-ingress</code>	<code>app=grafana</code>
<code>istio-ingressgateway</code>	<code>app=istio-ingressgateway,istio=ingressgateway</code>
<code>istio-istiod-production</code>	<code>app=istiod,istio.io/rev=production</code>

<sup>1</sup> Transmission Control Protocol (TCP), Hypertext Transfer Protocol (HTTP), HTTP/2, and Google Remote Procedure Call (GRPC)

Name	Pod selector
<code>istio-jaeger-ingress</code>	<code>app.kubernetes.io/component in (all-in-one,query), app.kubernetes.io/instance=jaeger-small-production, app.kubernetes.io/managed-by=jaeger-operator, app.kubernetes.io/part-of=jaeger</code>
<code>istio-kiali-ingress</code>	<code>app=kiali</code>
<code>istio-mesh-production</code>	<code>&lt;none&gt;</code>
<code>istio-prometheus-ingress</code>	<code>app=prometheus</code>
<code>partner-mesh-ingress</code>	<code>app=partner-mesh-ingress, federation.maistra.io/ingress-for=partner-mesh-ingress, istio=ingressgateway</code>

### Data plane network policies

The following NetworkPolicy resources have been applied in the data plane namespaces:

**Table 6. NetworkPolicy resource application for the data plane namespaces**

Name	Pod selector
<code>istio-expose-route-production</code>	<code>maistra.io/expose-route=true</code>
<code>istio-mesh-production</code>	<code>&lt;none&gt;</code>

### istio-proxy ports and purposes

Finally, these are the key ports exposed by `istio-proxy` and their purpose:

- ▶ 15001: Carry traffic out of the sidecar towards the main workload container, or as traffic response/request from the workload container
- ▶ 15006: Carry traffic into the sidecar
- ▶ 15000: Access diagnostics
- ▶ 15020: Access merged Prometheus telemetry from the Istio agent, Envoy, and application

More information about the [ports used by Istio](#) is available on the Istio website. →

## Check your mesh configurations

While there are many resources that provide configurations for the service mesh, only the `istiod` component applies the configurations to the data plane and `istio-proxy` sidecars. The mesh operator, Emma, can use the `istiocli` CLI tool to see and troubleshoot the states of the currently applied configurations.

## Check whether the configurations been applied

First, check that the service mesh configurations have actually been applied.

1. Use the `analyze` command to check for errors in the overall service mesh configurations. Note that this can take a long time for very large data planes.

```
istioctl analyze
```

2. Check that XDS protocol of the discovery services for the service mesh Envoy proxy are in sync between the control plane and data plane:

```
istioctl proxy-status istio-ingressgateway-6b948db88c-2sqth -i prod-istio-system \
-n prod-istio-system
```

or

```
istioctl proxy-status -i prod-istio-system -n prod-istio-system
```

## Confirm which configurations have been applied and where

If you suspect that configurations may be missing, misplaced, or simply wrong, you can use the observability stack and CLI tools to investigate.

1. Use Kiali to [visualize your Istio configurations](#) in each namespace and see any validation errors.
2. Check the `ServiceMeshControlPlane` resource and workload deployment for any configurations that have been added to enhance or override the expected configurations. More information about Istio proxy annotations is available on the [Istio website](#) and in the [Red Hat OpenShift documentation](#). →
3. Verify that the workload is part of the service mesh and is a member of the correct mesh:
  - ▶ [Validate sidecar injection](#) for your workload.
  - ▶ Check that the `maistra.io/member-of` label of the namespace that contains the workload points to the correct control plane namespace.
4. Use Prometheus to check for traffic type issues. Check that the `istio_agent_pilot_duplicate_envoy_clusters` and `istio_agent_pilot_destrule_subsets` metrics are greater than 0, and check for duplicates. This can often solve issues reported as "no healthy upstream."
5. Generate a readout of the pod and service exposed ports and applied Istio configurations using Kiali or `istioctl`.
  - ▶ In Kiali, select the workload you want to view from the *Workloads* section.
  - ▶ Use the `describe` command in the `istioctl` CLI:

```
istioctl experimental describe pod cars-v1-594b79cfbf-wlcg9.prod-travel-agency -i \
prod-istio-system -n prod-travel-agency
```

## Detailed configuration analysis

A [detailed configuration analysis and debugging exercise](#) is available in the online repository. You can also find more information in the [Debugging Envoy and Istiod](#) section of the Istio documentation.

## Check runtime health and security

When you experience runtime issues in your service mesh, it's important to check the health of both the control plane and data plane. The mesh operator, Emma, and the mesh developers, Farid and Cristina, can use Kiali, Jaeger, and Prometheus to investigate the health of the components within their respective domains.

### Check control plane health

The mesh operator, Emma, can check the control plane health.

1. Verify the installation of your [service mesh operators](#) and check your [control plane installation](#).
2. Check the `istiod` logs for all instances:

```
oc logs -f istiod-production-<POD-HASH>
```

### Check data plane health

The mesh developers, Farid and Cristina, can check the data plane health within their respective domains.

1. Use Kiali to see [overview](#) and [detailed views](#) of the health and state of applications, Istio configurations, services, and workloads.
2. Use Jaeger to [identify potential health issues](#). →
3. Use Prometheus to query Envoy and application [metrics](#). Additional information about [using metrics in Prometheus](#) is available in the online repository. →
4. Enhance `istio-proxy` logging levels for additional insight. You can view the expanded content in Kiali or via the `oc` CLI.
  - ▶ [Enable Envoy access logs](#) for your entire service mesh, including `istio-proxy` containers and ingress/egress gateways.

- ▶ Check current `istio-proxy` logging levels:

```
./istioctl proxy-config log <POD NAME>
```

- ▶ Apply new levels as needed:

```
./istioctl proxy-config log <POD NAME> --level http2:debug,grpc:debug
```

Be sure to consider possible [response flags](#) when [parsing the resulting logs](#). It may also be necessary to configure [ingress access logging](#) to get a complete view of the incoming traffic flows.

### Check security health

As service mesh security requirements become more complex, it's important to ensure that the certificates used for authentication and traffic encryption are correct and valid.

1. In Chapter 5, we defined two scripts—[verify-controlplane-certs.sh](#) and [verify-dataplane-certs.sh](#)—that verify these certificates. The mesh operator, Emma, can use these scripts to ensure that the expected certificates are in place.

2. You can also install and use the [ksniff community tool](#) to verify how TLS is applied at the traffic level:

```
WORKLOAD="istio-egressgateway-8598cbf7cb-nl68z"  
NAMESPACE="istio-system-egressgw-mtls-client"  
oc sniff $WORKLOAD -p -n $NAMESPACE -o output.pcap
```

3. Check your `istiod` logs for possible issues with provided certificates. For example, the log readout below shows a common error that occurs when the intermediate CA key is password-protected:

```
2022-09-16T11:50:06.830472Z          error  failed to create discovery service: failed to create CA:  
failed to create an istiod CA: failed to create CA KeyCertBundle (failed to parse private key PEM:  
failed to parse the RSA private key)
```

```
Error: failed to create discovery service: failed to create CA: failed to create an istiod CA:  
failed to create CA KeyCertBundle (failed to parse private key PEM: failed to parse the RSA private  
key)
```

## Chapter resources

- ▶ Istio documentation: [How does Envoy-based tracing work?](#)
- ▶ Istio documentation: [Debugging Envoy and Istiod](#)
- ▶ Red Hat Developer: [Service Mesh Troubleshooting Tools](#)
- ▶ Red Hat Customer Portal: [Packet capture inside Pod using community ksniff with OpenShift 4](#)
- ▶ Red Hat Customer Portal: [Consolidated Troubleshooting Article OpenShift Service Mesh 2.x](#)

## Chapter 8:

# Tuning the service mesh

In this chapter, we'll explore techniques for tuning your service mesh.

## Define desired outcomes for your tuning efforts

Travel by Keyboard's application and platform teams have provided non-functional requirements and desired outcomes for tuning the service mesh.

### Application team requirements

The application team – including the product owner, technical lead, and mesh developer roles – need to tune the mesh from the application aspect to handle the expected customer load. In this case, the expected load is 250,000 requests per day with a peak of 250 requests per second (rps).

### Platform team requirements

The platform team – including the cluster operator, mesh operator, and platform (application ops) roles – need to tune the mesh from the control plane aspect to optimize observability, ingress/egress application runtime, and configurations. In our example, the platform team wants to define the best practices, sizing guides, and benchmarks for assessing and improving the service mesh to deliver the best possible experience to project teams. To do this, they want to answer the following questions:

- ▶ How do we define the maximum capacity of a Red Hat OpenShift Service Mesh instance?
- ▶ How do we define the maximum number of applications that can join the mesh in the future?
- ▶ Which criteria and metrics should we use to govern and assess capacity?
- ▶ What are the current control plane and data plane limits of a Red Hat OpenShift Service Mesh instance?

## Understand service mesh tuning prerequisites

The objectives, architecture principles, and production setup of your service mesh determine the type of tuning required. For our example, we have defined these aspects in the [Final service mesh production setup](#) section of Chapter 3.

In cloud-based environments, there are many components – firewalls, load balancers, container platforms, and more – that should be tuned. In our example, however, we will focus on tuning two primary areas:

- ▶ **The data plane:** Consists of all `istio-proxy` (Envoy) sidecars and ingress and egress gateway components, and is responsible for handling workloads' incoming traffic.
- ▶ **The control plane:** Includes the `istiod` service and the observability stack. The control plane is responsible for keeping proxies updated with the latest configurations and certificates.

The following sections provide guidance for testing the performance of your service mesh, measuring sizing needs, and overall tuning options.

## Size the service mesh data plane

The application team needs to tune the data plane components—ingress and egress workloads and `istio-proxy` sidecars—for memory, CPU, and threads to meet the latency and throughput targets defined earlier. Correctly sizing the service mesh data plane requires testing a set of scenarios based on actual expected load. Different configurations should be load tested, assessed, and adapted until the desired performance outcome is reached.

A [detailed data plane tuning exercise](#) is available in the online resource repository. In the example, we configure the mesh for testing; test the default data plane values, tune `istio-proxy` threads, CPU resources, MEM resources, and the `mysql` database; and finally test the tuned configurations. →

### Monitoring the data plane

To make informed decisions about sizing, you need to monitor specific aspects of the data plane. Istio defines a [set of metrics](#) that allow you to monitor HTTP, HTTP/2 GRPC, and TCP traffic. Key metrics include:

- ▶ `istio_requests_total`, a counter that measures the total number of requests.
- ▶ `istio_request_duration_milliseconds`, a distribution that measures the latency of requests.

You can also monitor these metrics using [Grafana](#) and [Kiali](#), and set Prometheus alerts against them to help tune the data plane. Additional details about these [metrics](#) are available in the online resource repository.

Finally, you can retrieve information about the CPU and memory of the `istio-proxy` and main workload containers using the [containers-mem-cpu.sh script](#) provided in the [detailed data plane tuning exercise](#). As with the Istio metrics, you can also access these through Prometheus.

### Data plane tuning recommendations

While each tuning exercise is different, we recommend considering the following as a start.

#### Availability considerations

- ▶ Check pod priority and preemption, as most important pods have scheduling priority.
- ▶ Configure liveness, readiness, and startup probes.
- ▶ Ensure that you set realistic compute resources for containers, using the known limits for each container, and configure the [horizontal pod autoscaler \(HPA\)](#) accordingly.
- ▶ Check your deployment strategy selection.
- ▶ Configure and tune managed application and database connection pools.

## Proxy (Envoy) considerations

- ▶ Increase application concurrency when it is too thin to improve throughput.
- ▶ Upgrade traffic to HTTP2 to multiplex several requests over the same connection and avoid the overhead associated with creating new connections.
- ▶ Tune the client pool connections via Istio `DestinationRules` configurations to improve the performance of the network.
- ▶ Ensure that only the required configurations are sent to each proxy to avoid unneeded overhead in both the data plane and the control plane.

High-throughput workloads place additional demands upon your service mesh and can benefit from additional tuning. The online resource repository includes additional information about [tuning for high-throughput demands](#). →

## Size the service mesh control plane

The goal of tuning the control plane is to ensure that it can support the data plane effectively, handle some amount of expanded data plane capacity, and provide the required resources for the observability stack. The online resource repository includes additional information about [tuning the control plane](#), as well as [detailed tuning exercises](#). →

## Monitoring the control plane

As with the data plane, metrics are key to tuning your control plane. Istio also provides [metrics](#) that help with control plane tuning. Key metrics include:

- ▶ `pilot_xds`, the number of endpoints connected to istiod using xDS, or simply the number of clients that need to be kept up-to-date by the control plane.
- ▶ `pilot_xds_pushes`, the count of xDS messages sent, including build and send errors.
- ▶ `pilot_proxy_convergence_time`, the time it takes to push new configurations to Envoy proxies (in milliseconds).

## istiod tuning recommendations

Because one of the core functions of the control plane is to support the data plane, you need to adapt your `istiod` tuning when your data plane size increases significantly, for example by hundreds of pods. In this scenario, we recommend that you:

1. Review the [deployment model](#) of your service mesh, as a different model may better suit your new requirements.
2. Apply the `Sidecar` resource to allow separated visibility of resource configurations for unrelated namespaces within the same mesh.
3. Adjust the HPA settings for `istiod` components to allow for a predefined increase in the number of xDS clients.

## Capacity planning for the observability stack

Your observability stack setup will also impact control plane sizing and tuning. Capacity planning for the observability stack involves sizing your runtime components—including Kiali, Jaeger, Elasticsearch, Prometheus, and Grafana—and the associated long-term storage for metric, graph, and trace persistence. These capacity requirements are directly dependent

upon the size of your data plane, the number of incoming requests, and the capture and retention settings for metrics and traces. The online resource repository includes a [detailed observability stack tuning exercise](#) in which we check the final production setup for our example, and adjust the settings to align with the established non-functional requirements. →

## Tune across service mesh layers

We have provided guidance for uncovering capacity needs and basic tuning. However, to fine-tune your service mesh, you also need to consider aspects that cross both the control plane and data plane layers. A detailed understanding of these aspects – for example, TLS settings for communication into and out of clusters, service-to-service communication requirements, and bootstrapping configurations – is needed to create a stable set of benchmarks for fine-tuning.

## Chapter resources

### `istio-proxy` metrics

- ▶ Istio documentation: [Proxy-level metrics](#)
- ▶ Istio documentation: [Istio standard metrics](#)
- ▶ Istio documentation: [Envoy statistics](#)
- ▶ Envoy documentation: [Statistics overview](#)

### `istiod` metrics

- ▶ Istio documentation: [Exported metrics](#)

### Observability stack

- ▶ Kiali documentation: [Prometheus tuning](#)
- ▶ Istio documentation: [Using Prometheus for production-scale monitoring](#)

### Mesh Performance Testing Resources

- ▶ IstioCon presentation: [Scaling to 1M RPS with multicluster Istio](#)
- ▶ Tool: [Locust load-testing tool](#)

### Sizing and Tuning Practices

- ▶ Istio documentation: [Performance and scalability](#)
- ▶ Istio blog: [Best practices: Benchmarking Service Mesh performance](#)
- ▶ Red Hat OpenShift documentation: [Optimizing networking](#)

## Chapter 9:

# Upgrading your service mesh to a new version

In this chapter, we'll explore how to upgrade your service mesh to a new version.

## Identify which versions you currently use

When creating a service mesh upgrade plan of action, it's important to understand which service mesh component versions you currently use. Details about versions are provided in the [Understanding Service Mesh versions](#) and [How versioning affects Service Mesh upgrades](#) sections of the Red Hat OpenShift documentation.

Red Hat OpenShift Service Mesh contains four components that are affected by versioning:

- ▶ The Red Hat OpenShift Service Mesh operator.
- ▶ The Red Hat OpenShift Service Mesh control plane.
- ▶ The data plane sidecar image.
- ▶ The observability stack component operators.

### Red Hat OpenShift Service Mesh operator version

The Red Hat OpenShift Service Mesh operator version determines:

- ▶ The pod image used to operate the mesh.
- ▶ Which [versions of certain components](#) – like Istio, Envoy Proxy, Jaeger, and Kiali – are included by default.
- ▶ Which Red Hat OpenShift Service Mesh control plane versions are supported.
- ▶ Which [custom resource definitions \(CRDs\)](#) are supported.

Identify the current Red Hat OpenShift Service Mesh operator in our example:

```
oc -n prod-istio-system get csv | grep servicemesh
```

### Red Hat OpenShift Service Mesh control plane version

The Red Hat OpenShift Service Mesh control plane version determines the availability and configuration of control plane features. It also determines which version of Red Hat OpenShift Service Mesh will be used and, as a result, what the Red Hat OpenShift Service Mesh operator will deploy. An update to the control plane version affects the versions of all control plane components – including ingress and egress gateways and `istiod` – and the version of the sidecar image injected into data plane pods.

The latest Red Hat OpenShift Service Mesh operator currently supports multiple control plane versions, so it will not automatically update supported versions, and you may not need to update your control plane version in conjunction with an upgrade to your Red Hat OpenShift Service Mesh operator. More information about control plane versions is available in the [Upgrading the control plane](#) section of the Red Hat OpenShift documentation.

## Observability stack component operator versions

Red Hat OpenShift Service Mesh relies on the Jaeger, Elasticsearch, Prometheus, and Kiali components to provide observability for the mesh. The Red Hat OpenShift Service Mesh operator and control plane versions determine which versions of the observability components are required. Additionally, the resources that each of the observability component operators manage must be compatible with the installed version of that operator. As a result, these resources will need to be updated in conjunction with operator upgrades.

## Upgrade your service mesh

Based on the version dependencies, a service mesh upgrade incorporates both mandatory and optional component upgrades.

**Table 7. Required and optional component upgrades for a service mesh upgrade**

Component	Upgrade need
Red Hat OpenShift Service Mesh operator	Mandatory upgrade
Observability component operators	Mandatory upgrade
Red Hat OpenShift Service Mesh control plane	Optional upgrade
Sidecar image (via restart)	Optional upgrade

## Upgrade the Red Hat OpenShift Service Mesh operator

Red Hat OpenShift automatically creates a new installation plan whenever a new catalog source containing new operator versions is installed. If you selected manual update approvals during the initial Red Hat OpenShift installation, you will need to manually accept the installation plan to start the update. Information and additional resources regarding [upgrades in disconnected environments](#) is available in the online resource repository.

Once the installation plan has been either automatically or manually accepted, the upgrade procedure is started. First, the operators [are upgraded](#). Then all control plane components are restarted to receive the new Red Hat OpenShift Service Mesh operator version into the current control plane version. The new Red Hat OpenShift Service Mesh operator version becomes active only after the control plane is upgraded.

## Upgrade the control plane

The next step is to upgrade the control plane. To do so, you need to update the version field (`.spec.version`) of the control plane resource:

```
$ oc patch smcp <control_plane_name> --type json \
  --patch '[{"op": "replace", "path": "/spec/version", "value": "v2.2"}]'
```

Then, all control plane components – ingress and egress gateways and `istiod` – are restarted automatically and upgraded to the newest version images. Verify that the new control plane version is deployed:

```
$ oc get smcp -n istio-system
```

## Upgrade the data plane

Next, it's time to upgrade the application sidecar containers and Envoy proxies and configurations. To do so, restart the application pods:

```
$ oc rollout restart deployment $DEPLOYMENT-NAME
```

## Chapter resources

- ▶ Red Hat OpenShift documentation: [Upgrading Service Mesh](#)
- ▶ Red Hat OpenShift documentation: [Updating sidecar proxies](#)

### **About Stelios Kousouris, Senior Applications Architect**

With 20 years of experience in delivering software solutions, Stelios Kousouris has always been interested in bringing business solutions into production and optimizing software runtimes. He currently focuses on understanding how Red Hat OpenShift Service Mesh networking and the Serverless deployment paradigm can be best utilized, and guides teams in application modernizations for cloud environments.

### **About Ortwin Schneider, Principal Product Marketing Manager**

Ortwin Schneider has more than 20 years of professional experience in developing individual software solutions for various industries. In the past, he led an agile development team, supporting customers to provide business value and sustain agility. He is highly interested in software architectures that use hybrid cloud and cloud-native models. Ortwin is currently working as Principal Technical Marketing Manager at Red Hat on the Red Hat OpenShift team, focusing on creating cloud-native solution patterns.