

- N

Northstar

Admin
- Q

Find

⌘K
- <

Clusters
- staging-eu
- 🏠

Overview
- 🔌

Add-ons
- 📦

Stacks
- 📄

Operations
- 🔌

Pods
- 📋

Nodes
- >_

Kubectl access
- ⚙️

Settings

STAGING-EU

Operations

Everything the platform has run against this cluster, newest first, with how far each got and how it ended.

All statuses	☐ Needs attention only	4 operations
GitOps reconcile Write · Outcome not known 1 of 2 steps succeeded · 1 failed Manifest ownership conflict; review the merged spec. Started 14 Oct, 10:30 UTC · Recorded example	Failed	Open in Classic Ankra ↗
Deploy monitoring Write · Outcome not known 2 of 5 steps succeeded · 1 failed Chart download timed out after 30 seconds. Started 14 Oct, 10:18 UTC · Recorded example	Failed	Open in Classic Ankra ↗
Refresh inventory Read · Nothing to change 1 of 1 steps succeeded Started 14 Oct, 10:14 UTC · Recorded example	Succeeded	
Reconcile ingress Write · Changed the cluster 4 of 4 steps succeeded Started 14 Oct, 10:10 UTC · Recorded example	Succeeded	

Existing Next list and Classic handoff · Anonymised sample records

- N

Northstar

Admin
- Find

⌘K
- <

Clusters
- staging-eu
- Overview
- Add-ons
- Stacks
- Operations
- Pods
- Nodes
- Kubectl access
- Settings

STAGING-EU

Operations

Everything the platform has run against this cluster, newest first, with how far each got and how it ended.

All statuses	☐ Needs attention only	4 operations
GitOps reconcile Write · Outcome not known 1 of 2 steps succeeded · 1 failed Manifest ownership conflict; review the merged spec. Started 14 Oct, 10:30 UTC · Recorded example	Failed	View details
Deploy monitoring Write · Outcome not known 2 of 5 steps succeeded · 1 failed Chart download timed out after 30 seconds. Started 14 Oct, 10:18 UTC · Recorded example	Failed	View details
Refresh inventory Read · Nothing to change 1 of 1 steps succeeded Started 14 Oct, 10:14 UTC · Recorded example	Succeeded	View details
Reconcile ingress Write · Changed the cluster 4 of 4 steps succeeded Started 14 Oct, 10:10 UTC · Recorded example	Succeeded	View details

Proposed change: open native details from the existing list pattern.

STAGING-EU

GitOps reconcile failed

Recorded example · Write operation · Outcome not known

The merged spec was rejected

A manifest name is already associated with a stack. The error alone does not establish whether the cause is configuration or platform state.

Select a step to inspect its evidence →

- **Resolve source and includes**

Succeeded

Source read; merged spec produced
- **Validate manifest ownership**

Failed

Manifest ownership conflict

Review the source and merged stack ownership

Compare includes and resource names. If the source is valid, preserve the operation evidence for a platform investigation.

Review investigation

Back to operations

STAGING-EU

Resolve the configuration failure

A reviewed source change or a platform investigation comes before another attempt.

Inspect	Git revision, include paths and merged stack ownership
If duplicated	Prepare a Git change that gives the manifest one owner
If already valid	Attach operation ID and the redacted excerpt to an investigation
Before retry	Confirm the cause was resolved and review current preconditions
After execution	Read the new outcome and observed workload health separately

The storyboard does not claim a confirmed root cause

The live excerpt is evidence of a validation failure, not proof that renaming a manifest is the right fix.

Back to operations

Explore transient retry

STAGING-EU

Monitoring rollout failed

op-1842 · Git sync · monitoring · Started 10:18 UTC

Chart download timed out

The namespace and secret references were reconciled. Grafana and dependent checks did not run.

Select a step to inspect its evidence →

- Resolve Git revision

Succeeded

a31f2c8 · main
- Reconcile prerequisites

Succeeded

Namespace and secret references
- Download pinned chart

Failed

grafana 10.1.0 · timeout after 30 seconds
- Apply chart

Not started

Waiting for chart download
- Observe readiness

Not started

No new workload observation

Execution context

Initiator	Git sync
Revision	a31f2c8
Changed	Prerequisites
Health	Not verified
Failure type	Download timeout

Inspect failed step

Review retry

STAGING-EU

Review the retry

Retry op-1842 · Same cluster, pinned chart and reviewed revision

Target	Northstar / staging-eu / monitoring
Source	a31f2c8 · pinned Grafana chart 10.1.0
Retry scope	Platform retry policy; revalidate completed prerequisites
Expected change	Install chart and run dependent readiness checks
Existing changes	Namespace and secret references remain in place

Preconditions are checked again at execution

A changed Git base, lost access or another active operation requires a fresh review.

This is a new attempt linked to the original failure.

A deterministic configuration error should be fixed in Git before retrying.

Start sample retry

Back to failure

STAGING-EU

Follow the new attempt

op-1843 · Retry of op-1842 · Started 10:24 UTC

In progress

3 of 5 steps succeeded · 1 running · 1 pending

Select any step, including those that have not started →

● Revalidate context

Access, revision and prerequisites rechecked

Succeeded

● Resolve pinned chart

grafana 10.1.0 downloaded

Succeeded

● Apply chart

Release revision recorded

Succeeded

● Observe workloads

2 of 3 pods Ready at last observation

Running

● Close the attempt

Waiting for readiness observations

Pending

View later observation

Review cancellation

AM Alex Morgan
alex@example.test

NORTHSTAR / STAGING-EU

Live Next shell · Proposed flow · Sample data · 14 Oct 2026

Progress reports observed step counts. A cancellation request cannot promise reversal of completed mutations.

STAGING-EU

Review cancellation

op-1843 · Cancellation is a request, not a rollback

Some changes have already completed

The chart has been applied. Stopping observation will not uninstall it or undo the release.

Completed	Context checks, chart download and chart apply
In progress	Workload readiness observation
After request	Show cancellation requested until the backend confirms
Cluster outcome	Keep completed changes and any unknown outcome visible

Need to undo a change?

Prepare a separate reviewed Git change. Cancellation does not restore workload data.

Keep observing

View cancelled example

STAGING-EU

Execution cancelled

op-1843 · Terminal execution state · Completed changes preserved

Cancelled

Execution	Readiness observation stopped at 10:25 UTC
Cluster change	Chart apply completed before cancellation
Workload health	Unknown after the observation stopped
Pending work	Closing verification did not run

Inspect the current cluster state

Keep the audit receipt; a new read can establish present health without repeating a write.

Explore completed chart apply

Inspect nodes

Back to operations

STAGING-EU

The attempt completed

op-1843 · Alternate successful branch · Recorded 10:26 UTC

Succeeded	5 of 5 steps succeeded
Outcome	Changed the cluster
Source	a31f2c8 · Grafana 10.1.0
Execution	Completed 10:25:12 UTC
Observed readiness	3 / 3 pods Ready at 10:26 UTC
Original attempt	op-1842 · failed download retained in history

What this proves

The sample execution completed and pods were Ready at the recorded time. Ongoing health needs fresh observations.

Explore recorded chart apply

Inspect cluster nodes

Back to operations

STAGING-EU

Refresh inventory

op-1839 · staging-eu · read operation

Succeeded

1 of 1 steps succeeded · select any step

- Read cluster inventory

Succeeded
- Inventory was recorded. This read changed no cluster resources.

Recorded execution, not a live health check

Inventory reads do not change the cluster or prove current health.

Back to operations

Explore first step

STAGING-EU

Reconcile ingress

op-1838 · ingress stack · write operation

Succeeded

4 of 4 steps succeeded · select any step

- Resolve stack source

Succeeded

Pinned the ingress definition for this execution.
- Render and validate

Succeeded

The rendered manifests passed validation.
- Apply ingress release

Succeeded

The ingress release was reconciled.
- Observe controller readiness

Succeeded

The controller was Ready at 10:10:20; public routing was not checked.

Recorded execution, not a live health check

The controller was Ready at the recorded time. External DNS and TLS were not probed.

Back to operations

Explore first step

STAGING-EU

Resolve Git revision

op-1842 / Monitoring rollout · Step 1 of 5 · monitoring · a31f2c8

← Back to operation

Succeeded

10:18:00–10:18:02 · 2s

EXECUTION STEPS

1. Resolve Git revision

Succeeded

→

2. Reconcile prerequisites

Succeeded

→

3. Download pinned chart

Failed

→

4. Apply chart

Not started

→

5. Observe readiness

Not started

→

Steps preserve their original results.

A retry creates a separate attempt.

RECORDED RESULT

Pinned the source used by this attempt.

Depends on

No preceding step

Unblocks

Reconcile prerequisites



Recorded step output · UTC · bounded sample

```
01 10:18:00 resolve refs/heads/main
02 10:18:02 resolved a31f2c8
03 10:18:02 source checkout complete
```

Resources and observations

RESOURCE	RESULT	EVIDENCE
Git revision a31f2c8	Read	No cluster write

Next step →

STAGING-EU

Reconcile prerequisites

op-1842 / Monitoring rollout · Step 2 of 5 · monitoring · a31f2c8

← Back to operation

Succeeded

10:18:02–10:18:04 · 2s

EXECUTION STEPS

1. Resolve Git revision

Succeeded

→

2. Reconcile prerequisites

Succeeded

→

3. Download pinned chart

Failed

→

4. Apply chart

Not started

→

5. Observe readiness

Not started

→

Steps preserve their original results.

A retry creates a separate attempt.

RECORDED RESULT

Prerequisites remain in place after the later failure.

Depends on

Resolve Git revision

Unblocks

Download pinned chart



Recorded step output · UTC · bounded sample

```
01  10:18:03  namespace monitoring reconciled
02  10:18:04  secret references validated
03  10:18:04  secret values omitted from output
```

Resources and observations

RESOURCE	RESULT	EVIDENCE
Namespace / monitoring	Reconciled	Recorded write
Secret references	Validated	Values hidden

← Previous step

Next step →

STAGING-EU

Download pinned chart

op-1842 / Monitoring rollout · Step 3 of 5 · monitoring · a31f2c8

← Back to operation

Failed

10:18:05–10:18:35 · 30s

EXECUTION STEPS

1. Resolve Git revision

Succeeded



2. Reconcile prerequisites

Succeeded



3. Download pinned chart

Failed



4. Apply chart

Not started



5. Observe readiness

Not started



Steps preserve their original results.

A retry creates a separate attempt.

Review retry

RECORDED RESULT

Repository request timed out; no chart was applied.

Depends on

Reconcile prerequisites

Unblocks

Apply chart



Recorded step output · UTC · bounded sample

```
01 10:18:05 fetching grafana chart 10.1.0
02 10:18:35 ERROR request timed out after 30s
03 10:18:35 chart artifact unavailable
04 10:18:35 dependent steps will not start
```

Resources and observations

RESOURCE	RESULT	EVIDENCE
Grafana chart / 10.1.0	Unavailable	No cluster write

← Previous step

Next step →

STAGING-EU

Apply chart

op-1842 / Monitoring rollout · Step 4 of 5 · monitoring · a31f2c8

← Back to operation

Not started

No start or finish recorded

EXECUTION STEPS

1. Resolve Git revision

Succeeded

→

2. Reconcile prerequisites

Succeeded

→

3. Download pinned chart

Failed

→

4. Apply chart

Not started

→

5. Observe readiness

Not started

→

RECORDED RESULT

Blocked by the failed chart download.

Depends on

Download pinned chart

Unblocks

Observe readiness

This step has not produced output

It has not started. There is no log stream or step result to inspect yet.

Inspect preceding step

Resources and observations

RESOURCE	RESULT	EVIDENCE
Helm release / grafana	Planned only	No apply recorded

← Previous step

Next step →

Steps preserve their original results.

A retry creates a separate attempt.

NORTHSTAR / STAGING-EU

Live Next shell · Proposed flow · Sample data · 14 Oct 2026

STAGING-EU

Observe readiness

op-1842 / Monitoring rollout · Step 5 of 5 · monitoring · a31f2c8

← Back to operation

Not started

No start or finish recorded

EXECUTION STEPS

1. Resolve Git revision

Succeeded

→

2. Reconcile prerequisites

Succeeded

→

3. Download pinned chart

Failed

→

4. Apply chart

Not started

→

5. Observe readiness

Not started

→

Steps preserve their original results.

A retry creates a separate attempt.

RECORDED RESULT

The chart apply never ran; workload health was not observed.

Depends on

Apply chart

Unblocks

End of this execution

This step has not produced output

It has not started. There is no log stream or step result to inspect yet.

Inspect preceding step

Resources and observations

RESOURCE	RESULT	EVIDENCE
Workloads / monitoring	Not observed	Health unknown

← Previous step

STAGING-EU

Revalidate context

op-1843 / Monitoring retry · Step 1 of 5 · monitoring · retry of op-1842

← Back to operation

Succeeded

10:24:00–10:24:02 · 2s

EXECUTION STEPS

1. Revalidate context

Succeeded

→

2. Resolve pinned chart

Succeeded

→

3. Apply chart

Succeeded

→

4. Observe workloads

Running

→

5. Close the attempt

Pending

→

Steps preserve their original results.

A retry creates a separate attempt.

RECORDED RESULT

Access, source revision and prerequisites were rechecked.

Depends on

No preceding step

Unblocks

Resolve pinned chart



Recorded step output · UTC · bounded sample

```
01  10:24:00  linked original attempt op-1842
02  10:24:01  cluster scope and access revalidated
03  10:24:02  a31f2c8 and prerequisites verified
```

Resources and observations

RESOURCE	RESULT	EVIDENCE
Namespace / monitoring	Observed	Existing prerequisite

Next step →

STAGING-EU

Resolve pinned chart

op-1843 / Monitoring retry · Step 2 of 5 · monitoring · retry of op-1842

← Back to operation

Succeeded

10:24:02–10:24:05 · 3s

EXECUTION STEPS

1. Revalidate context

Succeeded

→

2. Resolve pinned chart

Succeeded

→

3. Apply chart

Succeeded

→

4. Observe workloads

Running

→

5. Close the attempt

Pending

→

Steps preserve their original results.

A retry creates a separate attempt.

RECORDED RESULT

The pinned chart was downloaded for this attempt.

Depends on

Revalidate context

Unblocks

Apply chart



Recorded step output · UTC · bounded sample

```
01  10:24:02  fetching grafana chart 10.1.0
02  10:24:05  chart artifact downloaded
03  10:24:05  ready for release apply
```

Resources and observations

RESOURCE	RESULT	EVIDENCE
Grafana chart / 10.1.0	Downloaded	No cluster write

← Previous step

Next step →

STAGING-EU

Apply chart

op-1843 / Monitoring retry · Step 3 of 5 · monitoring · retry of op-1842

← Back to operation

Succeeded

10:24:05–10:24:12 · 7s

EXECUTION STEPS

1. Revalidate context

Succeeded

→

2. Resolve pinned chart

Succeeded

→

3. Apply chart

Succeeded

→

4. Observe workloads

Running

→

5. Close the attempt

Pending

→

Steps preserve their original results.

A retry creates a separate attempt.

RECORDED RESULT

The release was applied. Readiness is a separate step.

Depends on

Resolve pinned chart

Unblocks

Observe workloads



Recorded step output · UTC · bounded sample

```
01  10:24:05  applying release monitoring/grafana
02  10:24:12  release revision 7 recorded
03  10:24:12  manifests submitted successfully
```

Resources and observations

RESOURCE	RESULT	EVIDENCE
Helm release / grafana	Revision 7	Recorded write
Deployment / grafana	Submitted	Readiness separate

← Previous step

Next step →

STAGING-EU

Observe workloads

op-1843 / Monitoring retry · Step 4 of 5 · monitoring · retry of op-1842

← Back to operation

Running

Started 10:24:12 · observed 10:24:36

EXECUTION STEPS

1. Revalidate context

Succeeded

→

2. Resolve pinned chart

Succeeded

→

3. Apply chart

Succeeded

→

4. Observe workloads

Running

→

5. Close the attempt

Pending

→

Steps preserve their original results.

A retry creates a separate attempt.

RECORDED RESULT

2 of 3 pods were Ready at the last recorded observation.

Depends on

Apply chart

Unblocks

Close the attempt



Recorded step output · UTC · bounded sample

```
01  10:24:12  readiness observation started
02  10:24:24  pods Ready: 1/3
03  10:24:36  pods Ready: 2/3
04  10:24:36  waiting for remaining pod
```

Resources and observations

RESOURCE	RESULT	EVIDENCE
Pods / monitoring	2 / 3 Ready	Observed 10:24:36

← Previous step

Next step →

STAGING-EU

Close the attempt

op-1843 / Monitoring retry · Step 5 of 5 · monitoring · retry of op-1842

← Back to operation

Pending

No start or finish recorded

EXECUTION STEPS

1. Revalidate context

Succeeded



2. Resolve pinned chart

Succeeded



3. Apply chart

Succeeded



4. Observe workloads

Running



5. Close the attempt

Pending



RECORDED RESULT

Waiting for the readiness observation to finish.

Depends on

Observe workloads

Unblocks

End of this execution

This step has not produced output

It has not started. There is no log stream or step result to inspect yet.

Inspect preceding step

Resources and observations

RESOURCE	RESULT	EVIDENCE
Execution receipt	Not produced	Dependency pending

← Previous step

Steps preserve their original results.

A retry creates a separate attempt.

STAGING-EU

Resolve source and includes

op-1841 / GitOps reconcile · Step 1 of 2 · cluster definition · merged ownership

← Back to operation

Succeeded

10:30:00–10:30:02 · 2s

EXECUTION STEPS

1. Resolve source and includes

Succeeded

→
2. Validate manifest ownership

Failed

→

RECORDED RESULT

The source was read and includes were merged.

Depends on

No preceding step

Unblocks

Validate manifest ownership

📄

 Recorded step output · UTC · bounded sample

01 10:30:00 reading cluster definition

02 10:30:01 resolving included stack definitions

03 10:30:02 merged spec ready for validation

Resources and observations

RESOURCE	RESULT	EVIDENCE
Cluster definition	Read	No write in this step

Steps preserve their original results.

A retry creates a separate attempt.

Next step →

STAGING-EU

Validate manifest ownership

op-1841 / GitOps reconcile · Step 2 of 2 · cluster definition · merged ownership

← Back to operation

Failed

10:30:02–10:30:03 · 1s

EXECUTION STEPS

1. Resolve source and includes

Succeeded

→

2. Validate manifest ownership

Failed

→

RECORDED RESULT

Ownership validation failed; the root cause needs investigation.

Depends on

Resolve source and includes

Unblocks

End of this execution



Recorded step output · UTC · bounded sample

```
01  10:30:03  manifest/example-config already belongs
02  10:30:03  to stack example-stack
03  10:30:03  manifest names must be unique per cluster
04  10:30:03  merged spec rejected
```

Resources and observations

RESOURCE	RESULT	EVIDENCE
Manifest / example-config	Conflict recorded	Inspect source
Stack / example-stack	Recorded owner	Verify merged spec

← Previous step

Steps preserve their original results.

A retry creates a separate attempt.

Review investigation

NORTHSTAR / STAGING-EU

Live Next shell · Proposed flow · Sample data · 14 Oct 2026

STAGING-EU

Read cluster inventory

op-1839 / Refresh inventory · Step 1 of 1 · staging-eu · read operation

← Back to operation

Succeeded

10:14:00–10:14:03 · 3s

EXECUTION STEPS

1. Read cluster inventory

Succeeded

→

RECORDED RESULT

Inventory was recorded. This read changed no cluster resources.

Depends on

No preceding step

Unblocks

End of this execution



Recorded step output · UTC · bounded sample

```
01  10:14:00  reading provider inventory
02  10:14:01  reading Kubernetes node inventory
03  10:14:03  observation recorded
04  10:14:03  no mutation requested
```

Resources and observations

RESOURCE	RESULT	EVIDENCE
Provider machines	Read	Observed 10:14:03
Kubernetes nodes	Read	Observed 10:14:03

Steps preserve their original results.

A retry creates a separate attempt.

STAGING-EU

Resolve stack source

op-1838 / Reconcile ingress · Step 1 of 4 · ingress stack · write operation

← Back to operation

Succeeded

10:10:00–10:10:01 · 1s

EXECUTION STEPS

1. Resolve stack source

Succeeded

→

2. Render and validate

Succeeded

→

3. Apply ingress release

Succeeded

→

4. Observe controller readiness

Succeeded

→

RECORDED RESULT

Pinned the ingress definition for this execution.

Depends on

No preceding step

Unblocks

Render and validate



Recorded step output · UTC · bounded sample

```
01 10:10:00 resolving ingress stack
02 10:10:01 source pinned for execution
```

Resources and observations

RESOURCE

RESULT

EVIDENCE

Ingress stack definition

Read

No cluster write

Steps preserve their original results.

A retry creates a separate attempt.

Next step →

STAGING-EU

Render and validate

op-1838 / Reconcile ingress · Step 2 of 4 · ingress stack · write operation

← Back to operation

Succeeded

10:10:01–10:10:03 · 2s

EXECUTION STEPS

1. Resolve stack source

Succeeded

→

2. Render and validate

Succeeded

→

3. Apply ingress release

Succeeded

→

4. Observe controller readiness

Succeeded

→

RECORDED RESULT

The rendered manifests passed validation.

Depends on

Resolve stack source

Unblocks

Apply ingress release



Recorded step output · UTC · bounded sample

```
01 10:10:01 rendering pinned chart and values
02 10:10:03 rendered manifests validated
03 10:10:03 secret values omitted
```

Resources and observations

RESOURCE

RESULT

EVIDENCE

Ingress manifests

Validated

No apply in this step

← Previous step

Next step →

Steps preserve their original results.

A retry creates a separate attempt.

NORTHSTAR / STAGING-EU

Live Next shell · Proposed flow · Sample data · 14 Oct 2026

STAGING-EU

Apply ingress release

op-1838 / Reconcile ingress · Step 3 of 4 · ingress stack · write operation

← Back to operation

Succeeded

10:10:03–10:10:09 · 6s

EXECUTION STEPS

1. Resolve stack source

Succeeded

→

2. Render and validate

Succeeded

→

3. Apply ingress release

Succeeded

→

4. Observe controller readiness

Succeeded

→

RECORDED RESULT

The ingress release was reconciled.

Depends on

Render and validate

Unblocks

Observe controller readiness



Recorded step output · UTC · bounded sample

```
01 10:10:03 reconciling ingress controller
02 10:10:09 release revision recorded
03 10:10:09 workload observation follows
```

Resources and observations

RESOURCE

RESULT

EVIDENCE

Ingress controller release

Reconciled

Recorded write

← Previous step

Next step →

Steps preserve their original results.

A retry creates a separate attempt.

STAGING-EU

Observe controller readiness

op-1838 / Reconcile ingress · Step 4 of 4 · ingress stack · write operation

← Back to operation

Succeeded

10:10:09–10:10:20 · 11s

EXECUTION STEPS

1. Resolve stack source

Succeeded

→

2. Render and validate

Succeeded

→

3. Apply ingress release

Succeeded

→

4. Observe controller readiness

Succeeded

→

RECORDED RESULT

The controller was Ready at 10:10:20; public routing was not checked.

Depends on

Apply ingress release

Unblocks

End of this execution



Recorded step output · UTC · bounded sample

```
01 10:10:09 observing ingress controller
02 10:10:20 deployment Ready
03 10:10:20 external DNS and TLS not probed
```

Resources and observations

RESOURCE	RESULT	EVIDENCE
Ingress controller	Ready	Observed 10:10:20
Public DNS / TLS	Not checked	No claim of reachability

← Previous step

Steps preserve their original results.

A retry creates a separate attempt.

NORTHSTAR / STAGING-EU

Live Next shell · Proposed flow · Sample data · 14 Oct 2026

STAGING-EU

Record desired count

op-2107 / Scale general workers · Step 1 of 4 · general · desired workers 4 → 6

← Back to operation

Succeeded

10:31:00–10:31:01 · 1s

EXECUTION STEPS

1. Record desired count

Succeeded

→

2. Create provider instances

Succeeded

→

3. Bootstrap and join

Running

→

4. Observe ready capacity

Pending

→

RECORDED RESULT

The accepted target is 6 workers; actual capacity is observed separately.

Depends on

No preceding step

Unblocks

Create provider instances



Recorded step output · UTC · bounded sample

```
01 10:31:00 validated general group capacity
02 10:31:01 desired count recorded: 6
```

Resources and observations

RESOURCE

RESULT

EVIDENCE

Node group / general

Desired 6

Definition updated

Steps preserve their original results.

A retry creates a separate attempt.

Next step →

STAGING-EU

Create provider instances

op-2107 / Scale general workers · Step 2 of 4 · general · desired workers 4 → 6

← Back to operation

Succeeded

10:31:01–10:31:26 · 25s

EXECUTION STEPS

1. Record desired count

Succeeded

→

2. Create provider instances

Succeeded

→

3. Bootstrap and join

Running

→

4. Observe ready capacity

Pending

→

RECORDED RESULT

Both instances exist at the provider; joining is still in progress.

Depends on

Record desired count

Unblocks

Bootstrap and join



Recorded step output · UTC · bounded sample

```
01  10:31:01  provider allocation requested
02  10:31:24  general-05 allocated
03  10:31:26  general-06 allocated
```

Resources and observations

RESOURCE	RESULT	EVIDENCE
Machine / general-05	Allocated	Provider observation
Machine / general-06	Allocated	Provider observation

← Previous step

Next step →

Steps preserve their original results.

A retry creates a separate attempt.

STAGING-EU

Bootstrap and join

op-2107 / Scale general workers · Step 3 of 4 · general · desired workers 4 → 6

← Back to operation

Running

Started 10:31:26 · observed 10:33:36

EXECUTION STEPS

1. Record desired count

Succeeded

→

2. Create provider instances

Succeeded

→

3. Bootstrap and join

Running

→

4. Observe ready capacity

Pending

→

RECORDED RESULT

general-05 joined; general-06 is waiting for SSH.

Depends on

Create provider instances

Unblocks

Observe ready capacity



Recorded step output · UTC · bounded sample

```
01 10:31:56 general-05 joined Kubernetes
02 10:33:36 general-06 SSH probe pending
03 10:33:36 provider Running is not node Ready
```

Resources and observations

RESOURCE	RESULT	EVIDENCE
Node / general-05	Joined	Ready observed
Machine / general-06	Waiting for SSH	Not joined

← Previous step

Next step →

Steps preserve their original results.

A retry creates a separate attempt.

Inspect provisioning

NORTHSTAR / STAGING-EU

Live Next shell · Proposed flow · Sample data · 14 Oct 2026

STAGING-EU

Observe ready capacity

op-2107 / Scale general workers · Step 4 of 4 · general · desired workers 4 → 6

← Back to operation

Pending

Waiting for bootstrap and join

EXECUTION STEPS

1. Record desired count

Succeeded

→

2. Create provider instances

Succeeded

→

3. Bootstrap and join

Running

→

4. Observe ready capacity

Pending

→

RECORDED RESULT

The latest fleet read is 4 / 6 Ready; this step has not completed.

Depends on

Bootstrap and join

Unblocks

End of this execution

This step has not produced output

It has not started. There is no log stream or step result to inspect yet.

Inspect preceding step

Resources and observations

RESOURCE	RESULT	EVIDENCE
Worker capacity	4 / 6 Ready	Separate fleet observation

← Previous step

Steps preserve their original results.

A retry creates a separate attempt.

STAGING-EU

Nodes

The machines behind staging-eu · UpCloud · kubeadm · Ankra provisions this infrastructure

- Machines
- Node groups
- Control plane
- Bastion & VMs

Machines

11

3 controllers · 6 workers · 2 other

Kubernetes ready

8 / 9

1 worker has not joined

Worker capacity

24 vCPU / 48 GiB

Provisioned capacity, not usage

Estimated run rate

€304 / month

Sample list estimate · all 11 priced

Needs you

general-04 is running, but has not joined Kubernetes

First boot is waiting for SSH. Last recorded probe: 38 seconds ago.

Inspect provisioning

MACHINE	ROLE / GROUP	PROVIDER	KUBERNETES	ZONE
general-04 4 vCPU · 8 GiB	Worker / general	Running	Not joined	fi-hel2
general-03 4 vCPU · 8 GiB	Worker / general	Running	Ready	fi-hel1
control-plane-01 4 vCPU · 8 GiB	Control plane	Running	Ready	fi-hel1
bastion-01 2 vCPU · 4 GiB	Bastion	Running	Not a node	fi-hel1

- View all machines
- Zone resilience

Provider observation 10:30:12 · Kubernetes observation 10:30:24 · Costs are estimates, not an invoice

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. One inventory joins provider machines to Kubernetes observations. A VM or bastion is never counted as an unready Kubernetes node.

- N

Northstar

Admin
- Q

Find

⌘K
- <

Clusters
- staging-eu
- 🏠

Overview
- 🔌

Add-ons
- 📶

Stacks
- 📄

Operations
- 🔌

Pods
- 📁

Nodes
- >_

Kubectl access
- ⚙️

Settings

STAGING-EU

Machines

One place for provider instances, Kubernetes nodes, bastions and standalone VMs

- Machines
- Node groups
- Control plane
- Bastion & VMs

Find a machine

Name, group, provider ID or address

Role

All roles

State

Needs attention first

MACHINE	ROLE / GROUP	PROVIDER	KUBERNETES	SIZE / ZONE
general-04 vm-104 · observed 10:30	Worker / general	Running	Not joined	4 vCPU · fi-hel2
general-03 vm-103 · observed 10:30	Worker / general	Running	Ready	4 vCPU · fi-hel1
control-plane-01 vm-001 · observed 10:30	Control plane	Running	Ready	4 vCPU · fi-hel1
control-plane-02 vm-002 · observed 10:30	Control plane	Running	Ready	4 vCPU · fi-hel2
bastion-01 vm-010 · observed 10:30	Bastion	Running	Not a node	2 vCPU · fi-hel1
build-cache vm-020 · observed 10:30	Standalone VM	Running	Not a node	4 vCPU · fi-hel1

Showing 6 of 11 machines · Active inventory · Removed machines are available in history

- View node groups
- Back to machines

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Search and role/group/state filters retain machine identities and keep provider state separate from node readiness. Removed infrastructure is an explicit historical view.

STAGING-EU

Node groups

Choose how capacity is provisioned, where it runs and which workloads can use it

- Machines
- Node groups**
- Control plane
- Bastion & VMs

Add node group

GROUP	OBSERVED / DESIRED	SCALING OWNER	PLACEMENT	EST. / MONTH
general 4 vCPU · 8 GiB per worker	4 / 4 · 3 Ready	Fixed count · Ankra	Zone spread	€96
analytics 4 vCPU · 8 GiB per worker	2 / 2 · 2 Ready	Fixed count · Ankra	Pinned fi-hel1	€48

2 pods are waiting for capacity

Check resource requests and placement before growing the general group.

Review capacity

Changes are reviewed together

Edit a group, review the cost and disruption, then follow each submitted operation.

- Open general group
- Review saved drafts

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Groups are first-class Ankra resources: instance type, actual/desired size, placement, scaling owner, labels and taints. Lifecycle actions live on the group, not every row.

STAGING-EU

general

Node group · Fixed count · Ankra owns the desired count

MachinesNode groupsControl planeBastion & VMs

Desired workers

4

Fixed count

Provider machines

4

All running

Kubernetes ready

3 / 4

general-04 has not joined

Estimated run rate

€96 / month

4 × €24 · sample list price

Instance	4 vCPU · 8 GiB · amd64
Placement	Spread across configured cluster zones
Labels	workload=general
Taints	None
Configuration	Ankra cluster definition · last changed by Alex

- Change count
- Autoscaling
- Placement rules
- Resize workers

- Inspect joining worker
- Back to groups

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. A focused group workspace makes count, scaling policy, placement and instance resize different decisions. Desired, allocated and Kubernetes-ready counts remain distinct.

STAGING-EU

Add capacity to general

2 checkout replicas are Pending · Decide whether more of this group can help

Machines	Node groups	Control plane	Bastion & VMs
WAITING WORKLOAD			
	checkout 2 replicas Pending	Insufficient memory	1 CPU · 3 GiB
	image-worker 1 replica Pending	Missing GPU resource	1 nvidia.com/gpu
			COMPATIBLE GROUP
			general
			No current group

Adding general workers addresses only the checkout shortage

The GPU workload needs a compatible GPU group and device plugin.

Plan GPU capacity

Desired worker count	Quota headroom	4 additional workers · observed 10:31
6 (currently 4)	Cost change	+ €48 / month · estimate
	Expected capacity	+ 8 vCPU · 16 GiB before reservations

Review scale-up

Back to group

STAGING-EU

Review scale-up

general · 4 → 6 workers · No existing machine is removed

Machines	Node groups	Control plane	Bastion & VMs
Change	Add 2 workers to general; retain the current instance type		
Placement	Use the configured zone spread		
Estimated cost	€96 → €144 / month · + €48 · sample catalogue estimate		
Configuration	Ankra cluster definition · worker count 4 → 6		
Workloads	New capacity becomes usable after kubelet joins and reports Ready		
Execution	Recheck provider quota, access and active infrastructure changes		

general-04 is still joining

Scaling does not repair this existing worker. Its provisioning operation stays visible separately.

Submit sample change

Keep current count

STAGING-EU

general is scaling

Operation op-2107 · Desired 6 · Provider machines 6 · Kubernetes ready 4

Machines

Node groups

Control plane

Bastion & VMs

●

Desired count recorded

Completed

Ankra cluster definition now requests 6 workers

●

Provider instances created

Completed

general-05 and general-06 allocated

●

Bootstrap and join

In progress

general-05 joined; general-06 waiting for SSH

●

Ready capacity observed

Waiting

4 of 6 workers Ready; general-04 and general-06 still joining

Follow the machine that is still starting

general-06: attempt 4 · 2m 10s of a 10m SSH budget · no answer yet.

Inspect machine

View observed result

Open operation history

STAGING-EU

Capacity is available

Later sample · op-2107 completed · Observations refreshed at 10:44 UTC

- Machines
- Node groups**
- Control plane
- Bastion & VMs

Desired workers

6

Fixed count

Provider machines

6

Running

Kubernetes ready

6 / 6

Observed 10:44 UTC

Pending checkout pods

0

2 new replicas Ready

Operation	op-2107 · added general-05 and general-06
Original worker	general-04 joined after the separate boot investigation
Cost estimate	€144 / month for this group · billing may differ
Unresolved workload	image-worker still requires a GPU group

The GPU workload still needs a placement decision

More general CPU workers cannot satisfy its GPU request.

Add a GPU group

- Back to groups
- Inspect workloads

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Receipt separates completed infrastructure work, actual node readiness and whether the original workloads recovered. This is a later fictional state, not an automatic resolution.

STAGING-EU

general-06

Provider machine is running · New worker from op-2107 · Not joined

- Machines
- Node groups
- Control plane
- Bastion & VMs

Provider

Running

Observed 10:33:36 UTC

Kubernetes

Not joined

No matched node record

Bootstrap

Waiting for SSH

Attempt 4 · recorded sample

Elapsed / budget

2m 10s / 10m

No live countdown

Machine	general-06 · group general
Last cause	SSH connection timed out; no answer from the private address
Access path	Ankra → bastion-01 → general-06
Owning operation	op-2107 · scale general from 4 to 6
Cloud-init	Output not yet readable over SSH

A provider instance exists. Usable Kubernetes capacity has not yet been observed for this machine.

- Back to bootstrap step
- Inspect bastion path

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Provisioning evidence stays bound to the selected machine and its operation. This is general-06 from the scale-up, not general-04 from the earlier failure.

STAGING-EU

general-04

Provider machine is running · Kubernetes node has not joined

MachinesNode groupsControl planeBastion & VMs

Provider Running Observed 10:30:12	Kubernetes Not joined No matched node record	Bootstrap Waiting for SSH Probe recorded 38 seconds ago	Elapsed / budget 2m 10s / 10m Attempt 4 · not a live countdown
---	---	--	---

Last cause	SSH connection timed out; no answer from the private address
Access path	Ankra → bastion-01 → general-04
Provider machine	vm-104 · fi-hel2 · group general
Owning operation	op-2103 · add worker · still running
Cloud-init	Output not yet readable over SSH

An active provisioning operation owns this machine. A second create or restart could interfere with first boot.

Investigate with Ankra

Inspect bastion path

STAGING-EU

Investigate general-04

Ask Ankra · Read-only investigation · This machine and its owning operation are in scope

MachinesNode groupsControl planeBastion & VMs

✦ Ankra

The instance exists, but it has not joined the cluster.

The recorded probe timed out through bastion-01. I cannot yet tell whether the route, SSH service or initial boot is the cause.

Next evidence to collect

1. Check the bastion health and route to fi-hel2.
2. Read the next SSH probe and cloud-init output.
3. Compare the machine state with operation op-2103.

No infrastructure change has been proposed or started.

Scope	general-04
Provider	Running · 10:30
Kubernetes	No matched node
Operation	op-2103
Mode	Ask · read-only

Inspect access path

Back to machine

STAGING-EU

Add a node group

Choose capacity for the workload · Provider catalogue and availability are account-specific

- Machines
- Node groups**
- Control plane
- Bastion & VMs

Example context: OVH self-managed cluster · GPU-enabled region · Eligible catalogue read at 10:32 UTC

Group name

inference

Instance type

GPU option · 1 GPU / node

Initial count

2 workers

Availability zone

Zone returned by the provider

Workload label

workload=inference

Scheduling taint

workload=inference:NoSchedule

Applications need matching placement

The workload must tolerate this taint and request a GPU. The device plugin must advertise the GPU before Kubernetes can schedule it.

Review GPU group

Back to groups

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Add-group supports an instance type, count, location/zone, labels and taints. GPU availability is a provider capability, not universal. This variant shows a catalogue option in a separate OVH example, not an UpCloud promise.

STAGING-EU

Review inference group

GPU-capable OVH example · 2 workers · Dedicated inference placement

Machines	Node groups	Control plane	Bastion & VMs
Group	inference · 2 GPU workers		
Hardware	Selected provider GPU type · 1 GPU per worker		
Placement	Selected OVH zone · workload=inference		
Taint	workload=inference:NoSchedule		
Before submission	Recheck stock, GPU quota and the current catalogue estimate		
Usable capacity	Nodes Ready and device plugin advertises the requested GPU resource		

A running GPU instance is not yet schedulable GPU capacity

After provisioning, inspect the advertised resource and the workload request. Metrics appear only when a compatible exporter is connected.

View later GPU worker

Edit group plan

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Review the same group entered in the add form. Stock, quota, GPU model and price must be supplied by the selected provider catalogue. Device-plugin readiness is an independent post-provision check.

STAGING-EU

Review workload placement

analytics · Pin data where its volumes can attach

Machines	Node groups	Control plane	Bastion & VMs
Zone	fi-hel1 · pinned to the volume location		
Group label	workload=analytics		
Taint	workload=analytics:NoSchedule		
Application requirement	nodeSelector workload=analytics; matching toleration		
Storage constraint	Volumes in fi-hel1 cannot be assumed attachable in another zone		
Existing workloads	Label and taint changes do not promise automatic migration		

Changing node spread does not make the application zone-resilient

Review replicas and storage topology before moving stateful workloads.

Inspect zone resilience

Review group changes

Back to group

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Zone pinning, group labels and taints exist in Ankra. Compatibility with workload affinity, tolerations and zonal storage is a proposed review aggregation. Editing labels does not move existing pods.

STAGING-EU

Let the group grow with demand

general · Choose who owns the worker count

Machines

Node groups

Control plane

Bastion & VMs

Scaling owner

Ankra autoscaler

▼

Current workers

4

Quota headroom

Up to 4 more

Estimate at minimum

€72 / month

Estimate at maximum

€192 / month

Minimum workers

3

Maximum workers

8

Autoscaling owns the count within these bounds

The fixed-count editor and quota-maximum action are unavailable while autoscaling is enabled. Provider quota can still prevent growth.

Moving back to fixed count is a separate reviewed ownership change. Show the observed count at that time.

Review policy change

Inspect quota limits

STAGING-EU

Review autoscaling policy

general · Fixed count of 4 → Ankra autoscaler with 3–8 workers

Machines	Node groups	Control plane	Bastion & VMs
Owner change	Fixed count → Ankra autoscaler		
Bounds	Minimum 3 · maximum 8 workers		
Current machines	4; the autoscaler will decide changes within the bounds		
Sample estimate	€72–€192 / month for this group; not a fixed bill		
Current quota	4 additional workers available at last provider read		
After acceptance	Observe policy and scaling decisions independently of node readiness		

The autoscaler replaces the manual count decision

To choose a fixed count again, review a separate ownership change using the current observed count.

Return to group workspace

Edit policy

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Ownership and limits are reviewed together. Enabling the policy does not immediately promise a specific observed count. Current quota is a check, not reserved future headroom.

STAGING-EU

How far can general grow?

Account quota · Shared with other clusters using this provider account

- Machines
- Node groups**
- Control plane
- Bastion & VMs

LIMIT	AVAILABLE	PER WORKER	ADDITIONAL WORKERS
vCPU quota	16 vCPU	4 vCPU	4
Memory quota	64 GiB	8 GiB	8
Instance quota	10 machines	1 machine	10

Binding limit	vCPU quota · maximum total 8 workers
Current / maximum	4 → 8 · 4 additional workers
Estimated cost increase	+ €96 / month at the maximum · sample list price
Availability	Quota is not stock. Check the selected type and zone again at execution

- Review maximum count
- Choose a smaller count

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Develop includes Max out. The redesign presents it as a reviewed quota ceiling, not a recommendation. Unknown quota or size disables the computed action; provider stock remains a separate limit.

STAGING-EU

Review scaling to the quota ceiling

general · 4 → 8 workers · An explicit capacity choice

Machines	Node groups	Control plane	Bastion & VMs
Proposed count	4 → 8 · add 4 workers		
Limit used	16 available vCPU / 4 vCPU per worker		
Estimated cost	€96 → €192 / month · sample list price		
Shared account impact	Consumes the currently reported vCPU headroom		
Stock	Provider instance and zone availability still need validation		
Execution gate	Refuse stale quota, lost authority or conflicting infrastructure work		

A quota ceiling is not a capacity recommendation

Choose the smaller count your workload requires. Using all headroom can prevent another cluster from growing.

Choose a smaller count

Back to quota evidence

STAGING-EU

Review node group changes

2 staged changes · Cost, ownership and disruption before submission

Machines	Node groups	Control plane	Bastion & VMs
GROUP	BEFORE → AFTER	EXPECTED EFFECT	ESTIMATED COST
general Fixed count	4 → 6 workers	Add 2 machines	+ €48 / month
analytics Placement	Add workload label / taint	Future scheduling changes	No machine change

Each change runs separately

If one change is refused, completed changes remain. Keep the failed draft and link each accepted change to its operation.

Changed definition

Ankra infrastructure configuration; Git-backed fields must show their source

Recheck

Current ownership, provider capabilities, permissions and active operations

Submit sample changes

Return to drafts

STAGING-EU

One change needs attention

Review result · One accepted change and one refused change

Machines	Node groups	Control plane	Bastion & VMs
CHANGE	RESULT	WHAT HAPPENED	NEXT STEP
general · 4 → 6	Accepted	op-2107 is provisioning workers	Follow operation
analytics placement	Not accepted	Another update changed the group	Review current configuration

The analytics draft is preserved

Refresh its current configuration and compare your edit. The accepted scale-up remains in progress.

Review failed draft

Follow accepted scale-up

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. A concrete partial-write state from Save all semantics. Do not retry a batch that already launched capacity changes; retry only a refused draft after fresh review.

STAGING-EU

Resize general workers

Change instance type · Rolling infrastructure change

Machines	Node groups	Control plane	Bastion & VMs
Before → after	4 vCPU / 8 GiB → 8 vCPU / 16 GiB per worker		
Scope	4 machines in general · preserve group placement rules		
Workload disruption	Drain one worker at a time; wait for it to return Ready		
Disruption budgets	Must permit eviction; a blocked drain stops the rollout		
Provider effect	Instance upgrade may be irreversible; no automatic downgrade		
Estimated cost	€96 → €192 / month · sample estimate		

Review capacity before the first drain

The remaining workers must be able to run the displaced workloads. A Ready node alone does not establish that there is enough capacity.

Inspect affected workloads

Back to group

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Group instance-type changes are distinct from count changes and may power-cycle or replace instances depending on provider. Irreversible effects, drain, headroom and capability must come from the operation contract.

STAGING-EU

Remove analytics capacity

Scale analytics from 2 → 1 · A disruption budget currently blocks removal

Machines	Node groups	Control plane	Bastion & VMs
Requested count	2 → 1 worker		
Current provider machines	2 · both retained		
Blocker	analytics-worker disruption budget allows 0 evictions		
Potential saving	€24 / month only after a machine is actually removed		
Workload owner action	Restore replica availability or review the workload policy		

No machine has been removed

The desired change is blocked. Preserve the operation and its drain evidence, then reassess the workload before trying again.

Remove worker

Inspect blocker

Back to groups

STAGING-EU

general-03

Worker · general · UpCloud / fi-hel1 · Provider vm-103

Machines

Node groups

Control plane

Bastion & VMs

Provider

Running

Observed 10:30:12

Kubernetes

Ready

Observed 10:30:24

Scheduling

Available

18 pods on this node

Memory usage

7.1 / 7.6 GiB

Metrics · 5 minute window

Memory pressure needs investigation

MemoryPressure is True. Read workload requests and recent usage before choosing more or larger capacity.

Inspect workloads

RESOURCE	CAPACITY	ALLOCATABLE	POD REQUESTS / USAGE
CPU	4 vCPU	3.8 CPU	2.9 CPU requested · 1.4 CPU measured
Memory	8 GiB	7.6 GiB	6.2 GiB requested · 7.1 GiB measured
Pods	110	110	18 scheduled

Review maintenance

Resize the group

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Machine identity links provider and Kubernetes data. Capacity, allocated pod requests and measured use are different quantities. Actions belong to the correct layer: group resize, node scheduling, provider restart.

STAGING-EU

Workloads on general-03

What uses this capacity, and what can safely move?

Machines	Node groups	Control plane	Bastion & VMs
WORKLOAD	MEMORY REQUEST / USE		DISRUPTION
checkout commerce · Deployment	2 GiB / 2.4 GiB		1 eviction allowed
analytics-worker analytics · Deployment	4 GiB / 4.6 GiB		0 evictions allowed
node-exporter monitoring · DaemonSet	64 MiB / 48 MiB		Excluded from drain

Usage: 5 minute window ending 10:30 UTC · Requests: pod spec · Disruption budgets: observed 10:30 UTC

analytics-worker cannot currently be evicted

Fix replica availability or review its policy with the application owner before maintenance.

Ask Ankra

- Review maintenance
- Review group scale-down

STAGING-EU

Investigate memory pressure

Ask Ankra · general-03 and the workloads scheduled on it

- Machines
- Node groups
- Control plane
- Bastion & VMs

The node is Ready, but memory pressure is reported

Node evidence	MemoryPressure=True at 10:30 UTC
Measured usage	7.1 GiB of 7.6 GiB allocatable · 5 minute window
Largest workload	analytics-worker: 4 GiB requested, 4.6 GiB measured
Maintenance blocker	Its disruption budget permits no eviction
Placement constraint	Its volume is tied to fi-hel1

Review the workload before selecting a capacity change

Check its usage trend and requests, restore replica availability, then compare adding compatible fi-hel1 capacity with resizing its group.

- Review compatible capacity
- Return to workloads

STAGING-EU

Prepare general-03 for maintenance

Pause placement, move workloads, verify the result, then return the machine to service

Machines	Node groups	Control plane	Bastion & VMs
CHECK	OBSERVATION	RESULT	
Node and ownership	Worker in general · fresh matching identity	Verified	
Disruption allowance	analytics-worker permits 0 evictions	Blocked	
Replacement placement	Requires fi-hel1 capacity and volume attachment	Review needed	
Pod policy	DaemonSets excluded; local data and unmanaged pods must be checked	Not complete	

Cordon prevents new placement. It does not move existing pods or make this machine safe to shut down.

Cordon and drain

Inspect blockers

Review cordon only

Back to machine

STAGING-EU

Pause new scheduling

general-03 · Cordon only · Existing pods keep running

Machines	Node groups	Control plane	Bastion & VMs
Target	general-03 · node UID and current revision revalidated		
Change	Scheduling available → cordoned		
Existing workloads	18 pods continue running on this node		
Cloud machine	No restart, resize or provider deletion		
Return to service	A separate uncordon action enables new scheduling		

This does not complete maintenance

Resolve the drain blockers and verify workloads before a restart or infrastructure change.

View sample receipt

Keep scheduling enabled

STAGING-EU

general-03 is cordoned

Scheduling observed at 10:35 UTC · Existing workloads remain on the node

- Machines
- Node groups
- Control plane
- Bastion & VMs

Kubernetes Ready Observed 10:35 UTC	Scheduling Cordoned New placement paused	Pods on node 18 No pods evicted	Provider Running No infrastructure change
--	---	--	--

Accepted change	Cordon only
Maintenance readiness	Not established; workloads remain
Recorded action	op-2111 · Alex · 10:34 UTC
Next decision	Resolve blockers or review returning scheduling to service

- Review return to service
- Inspect blockers

Receipt requires fresh node observation. Return to service is an independent reviewed uncordon; it does not rebalance existing workloads. Standalone drain remains a separate planned capability.

STAGING-EU

Return general-03 to scheduling

Uncordon only · Confirm maintenance is finished or abandoned

Machines	Node groups	Control plane	Bastion & VMs
Current node	Ready · cordoned · fresh observation at 10:36 UTC		
Pressure	MemoryPressure remains True; investigate the workload		
Action	Allow new pods to be scheduled		
Existing workloads	No automatic rebalance and no restart		
Before confirmation	Verify no maintenance or infrastructure operation still owns the node		

Pressure is unresolved

Returning scheduling does not fix memory pressure. Keep the node cordoned if maintenance or diagnosis is still in progress.

View later observed state

Keep cordoned

A separate action rechecks access, identity, freshness and current scheduling state. Existing pods are not automatically redistributed. Pressure is surfaced for the operator rather than disguised as a successful maintenance receipt.

STAGING-EU

Scheduling is enabled

Later sample · Uncordon observed at 10:38 UTC

- Machines
- Node groups
- Control plane
- Bastion & VMs

Scheduling

Available

unschedulable = false

Kubernetes

Ready

Observed 10:38 UTC

Existing pods

18

No automatic rebalance

Memory pressure

Still reported

Diagnosis remains open

Scheduling changed; the pressure investigation remains open

The operation completed its requested scheduling change. Review workload usage separately.

Inspect workloads

Action	op-2112 · uncordon general-03
Related action	op-2111 · cordon only
Observed result	Scheduling is enabled; no repair or drain was performed

- Back to machines
- Open node groups

STAGING-EU

inference-01

GPU worker · Separate GPU-capable provider example · Workload capacity and device health

- Machines
- Node groups
- Control plane
- Bastion & VMs

Provider GPUs 1 From the instance catalogue	Advertised GPUs 1 nvidia.com/gpu · device plugin	GPU requests 1 One inference pod	Measured utilisation Unavailable No DCGM source connected
--	---	---	--

Node readiness	Ready · observed 10:40 UTC
Placement	workload=inference · matching NoSchedule toleration
Workload	model-server · requests 1 GPU · Running
Driver / device health	Report only the fields the node or exporter supplies
Next decision	Inspect requests before adding expensive GPU capacity

- Review GPU group
- Back to machines

STAGING-EU

What happens if a zone is lost?

staging-eu · 3 control planes across 3 zones · UpCloud mesh topology

Machines	Node groups	Control plane	Bastion & VMs
ZONE	CONTROL PLANE	WORKERS	DEPENDENCIES
fi-hel1 Primary zone	1 / 1 Ready	3 workers	Ingress and zonal volumes
fi-hel2	1 / 1 Ready	2 workers · 1 joining	WireGuard mesh
se-sto1	1 / 1 Ready	1 worker	WireGuard mesh

Control-plane quorum and application availability differ

The controller spread can retain quorum after one zone loss. Primary-zone ingress, volumes and application replicas need their own resilience review.

Inspect control plane

Review workload placement

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Control-plane topology may preserve etcd quorum, but node placement alone does not prove application resilience. In UpCloud the load balancer and CSI storage remain primary-zone concerns.

STAGING-EU

Control plane

3 controllers · kubeadm · Stacked etcd · Ankra-managed machines

- Machines
- Node groups
- Control plane
- Bastion & VMs

Controllers

3 / 3 Ready

Observed 10:40 UTC

Topology

Stacked etcd

One controller per zone

Kubernetes

v1.33.2

All observed at target

Instance size

4 vCPU / 8 GiB

Same type on all controllers

MACHINE	ZONE	READINESS	VERSION
control-plane-01	fi-hel1	Ready	v1.33.2
control-plane-02	fi-hel2	Ready	v1.33.2
control-plane-03	se-sto1	Ready	v1.33.2

- Resize controllers
- Change controller count
- Kubernetes upgrade

Capabilities are read independently: live rolling resize allowed · count change requires stopped cluster

- Review rolling resize
- Inspect zone resilience

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Count and instance-size capabilities are independent. A running three-controller cluster may support rolling resize while changing count still requires a stopped cluster. Never share one permission/disabled reason for both.

STAGING-EU

Resize the control plane

3 controllers · Live rolling change supported by this cluster

Machines	Node groups	Control plane	Bastion & VMs
Size change	4 vCPU / 8 GiB → 8 vCPU / 16 GiB		
Sequence	One controller at a time; observe Ready before continuing		
Disruption	Drain respects workload disruption budgets		
Quorum	Do not start while another controller is unavailable		
Provider constraints	Check target type, availability and irreversible resize effects		
Estimated delta	+ €72 / month across 3 machines · sample price		

The next controller waits for a verified result

If a drain or readiness check fails, preserve the machine state and operation evidence. Do not continue automatically.

Follow sample operation

Back to control plane

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Supported rolling resize requires a running three-controller topology and fresh backend capability. Each node is drained and resized in sequence; a PDB block stops progress. Do not imply rollback or downgrade.

STAGING-EU

Change controller count

1 or 3 controllers · This change requires the cluster to be stopped

Machines	Node groups	Control plane	Bastion & VMs
Current / proposed	3 → 1 controllers · illustrative reduction		
Availability impact	One controller is a single point of failure		
Required state	Cluster stopped; current cluster is running		
Standalone VMs	Remain running when the Kubernetes cluster is stopped		
Apply timing	Definition is used when the cluster starts again		

Plan the outage before changing topology

Stopping Kubernetes affects its workloads. This screen does not stop the cluster as a side effect.

Change count

Back to control plane

STAGING-EU

Review a Kubernetes upgrade

Upgrade the cluster version · Do not confuse this with resizing instance types

Machines	Node groups	Control plane	Bastion & VMs
Observed version		v1.33.2 across control planes and workers	
Target		Select an eligible next version from the platform catalogue	
Before execution		Validate compatibility and take the supported etcd snapshot	
Order		Control planes first, then workers; one node at a time	
Stop condition		PDB refusal or node not Ready at target version	
Version policy		No downgrade or skipped minor version	

Rehearse outside production first

Use a reviewed target and monitor the durable operation. Nodes shows which machines reached the version and which are still waiting.

View operation history

Back to control plane

STAGING-EU

Bastion & VMs

Infrastructure that is not a Kubernetes node

- Machines
- Node groups
- Control plane
- Bastion & VMs

Bastion

MACHINE	PROVIDER	ACCESS PATH	LAST CHECK
bastion-01 2 vCPU · 4 GiB · fi-hel1	Running	Needs diagnosis	10:30 UTC

Diagnose access path

Standalone VMs

Private machines for databases, build services or other workloads outside Kubernetes.

MACHINE	SIZE / NETWORK	STATE	KUBERNETES
build-cache 4 vCPU · 8 GiB	fi-hel1 · private only	Running	Not a node

Add a VM

Inspect build-cache

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. This first-class Ankra use case was absent from the initial design. Bastion access-path health is different from provider power; standalone VMs have independent lifecycle and remain up when the cluster stops.

STAGING-EU

Diagnose the access path

bastion-01 · Provider running · Private SSH reachability is a separate observation

Machines	Node groups	Control plane	<u>Bastion & VMs</u>
HOP / CHECK		RECORDED RESULT	OBSERVED
Provider power		Running	10:30:12 UTC
Ankra → bastion SSH		Reachable	10:33:02 UTC
Bastion → general-04		Timed out	10:33:18 UTC
Cloud-init on worker		Not readable yet	10:33:18 UTC

The timeout is beyond the bastion

A bastion restart is not supported by this evidence alone. Keep the provisioning operation and inspect the worker boot and private route.

Back to worker

Review bastion resize

STAGING-EU

Review bastion resize

bastion-01 · Access and egress impact must be explicit

Machines	Node groups	Control plane	Bastion & VMs
Before → after	2 vCPU / 4 GiB → 4 vCPU / 8 GiB		
Machine role	Bastion; not a Kubernetes node		
Interruption	SSH hops become unavailable while the instance changes		
Network dependency	Review private-network egress and in-flight provisioning		
Concurrent operation	general-04 is still bootstrapping through this access path		
Expected result	Provider resize complete, then fresh access-path health check		

Wait for the active provisioning operation

Do not resize the access path while a worker bootstrap depends on it. Review again when that operation is finished.

Resize bastion

Inspect provisioning

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Resize is provider-capability gated and distinct from worker-group resize. The blast radius includes SSH access and, where routed through this bastion, private-network egress. No Kubernetes drain applies to this VM.

STAGING-EU

Add a standalone VM

A machine on the cluster network, independent of Kubernetes scheduling

- Machines
- Node groups
- Control plane
- Bastion & VMs

Name

build-cache-02

Provider size

4 vCPU · 8 GiB · catalogue option

Location

fi-hel1 · same as bastion

Network exposure

Private only

Access	Cluster SSH keys, through the bastion
Kubernetes	No kubelet, node group or pod scheduling
Cluster stop	This VM remains running and billable
Estimated cost	€24 / month · sample catalogue price

- Review VM creation
- Back to VMs

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Matches the Ankra standalone VM contract: name, provider type, optional location and public_ip=false by default. It uses cluster SSH keys and remains running when the cluster stops.

STAGING-EU

Review VM creation

build-cache-02 · Private network · No public address

Machines	Node groups	Control plane	Bastion & VMs
Machine	build-cache-02 · selected provider size		
Zone and network	fi-hel1 · cluster private network		
Public address	None; access through bastion-01		
Keys	Existing cluster SSH key references		
Lifecycle	Runs independently of Kubernetes cluster stop/start		
Estimated cost	+ €24 / month · sample catalogue estimate		

This provisions a machine, not a managed application

Install and maintain the software on it separately, or use an Ankra managed service when that matches the workload.

View later VM result

Edit VM plan

STAGING-EU

build-cache

Standalone VM · Running · Private network · fi-hel1

- Machines
- Node groups
- Control plane
- Bastion & VMs

Provider state

Running

Observed 10:42 UTC

Capacity

4 vCPU / 8 GiB

Provider instance

Network

Private only

Reach through bastion-01

Estimated run rate

€24 / month

Sample list price

Ownership	Ankra standalone VM
Kubernetes	Not a node; no pods or disruption budgets
Power dependency	Remains running when staging-eu stops
Size change	Review provider restart and irreversible size constraints
Removal	Destroys this VM; review data and storage retention first

- Review removal
- Back to VMs

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Standalone VMs support add, resize and remove. Removal destroys the provider VM, not merely an inventory row; data protection is an explicit decision. Kubernetes cordon/drain controls do not belong here.

STAGING-EU

Remove build-cache

Delete the provider VM · This is not a Kubernetes node action

Machines	Node groups	Control plane	Bastion & VMs
Target	build-cache · provider vm-020 · fi-hel1		
Effect	Remove the VM from the provider		
Software and data	Anything stored on its instance disk may be lost		
Attached storage	Retention has not been established; inspect before proceeding		
Dependencies	Confirm no build service still relies on this private address		

Confirm machine name

build-cache

Data retention is unresolved in this sample. Keep removal unavailable until the review is complete.

Remove VM

Keep the VM

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Removal requires explicit scope and data review. Do not assert attached-disk retention without provider evidence, and do not substitute deleting a Kubernetes Node record for infrastructure removal.

STAGING-EU

Node pools

Provider-managed Kubernetes example · The provider owns the control plane

- Machines
- Node pools**

Control plane

Provider managed

No controller VM controls

Observed workers

5

From provider and Kubernetes

Ready workers

5 / 5

Observed 10:40 UTC

Pool ownership

Mixed

Ankra and external autoscaler

POOL	OBSERVED COUNT	OWNER	ALLOWED NEXT STEP
workers Fixed count	2	Ankra	Review pool configuration
burst UKS external autoscaler	3	External autoscaler	Inspect ownership

Provider capabilities decide the controls

UKS autoscaling is owned externally. Ankra must not offer its own autoscaling policy for this pool.

- Inspect external pool
- Compare imported cluster

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Managed kinds expose Overview and Node pools only. Pool scaling and autoscaling follow provider capabilities. No fabricated controller machines, bastion or self-managed resize actions appear.

STAGING-EU

workers

Provider-managed pool · Fixed count · Ankra owns explicit pool changes

Machines	Node pools
Provider example	UKS · provider-managed control plane
Pool ownership	Ankra
Observed / desired	2 / 2 workers · both Ready
Proposed count	3 workers
Autoscaling	Not managed by Ankra for UKS
Execution	Provider pool operation followed by observed worker readiness

Scaling changes the pool, not the control plane

Read provider quota and permitted pool sizes, then follow the provider result and Kubernetes observations independently.

Review pool scale-up

Back to pools

STAGING-EU

Review workers pool scale-up

UKS workers · 2 → 3 nodes · Ankra-owned pool

Machines	Node pools
Target	UKS cluster / workers pool
Change	2 → 3 workers · same pool size
Control plane	Unchanged; owned by the provider
Current ownership	Ankra; externally managed is false
Estimated delta	Not available from the current catalogue read
After submission	Track provider operation, then the new node joining and becoming Ready

Pricing is not available in this sample

Do not display a zero-cost scale-up. Obtain the provider estimate before approving a cost-sensitive change.

Return to pool workspace

Keep 2 workers

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. Managed-pool scale reviews provider scope and count. An absent price is unavailable rather than zero. Once accepted, retain the provider operation and observed node state.

STAGING-EU

burst is externally managed

UKS · The external autoscaler owns this pool

Machines	Node pools
Owner	External autoscaler · not Ankra
Observed count	3 workers · read from the provider at 10:40 UTC
Ankra writes	Scale, edit, replace and delete are unavailable
Before handback	Coordinate and stop the external controller from changing this pool
Handback result	Adopt the provider live count as the Ankra-owned count

Changing this flag does not stop the external autoscaler

Hand back ownership only after coordinating the controller. Otherwise two systems may compete over the count.

Review ownership handback

Back to pools

Inspect imported view

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. While externally_managed is true, Ankra refuses count changes, other setting changes, replacement and deletion. Handback is a separate operation that adopts the provider live count.

STAGING-EU

Review pool ownership handback

burst · External controller → Ankra

Machines	Node pools
External controller	Operator confirms it no longer changes this pool
Fresh provider count	3 workers at the time of this sample review
Change	Clear external ownership and adopt the live count
Retained infrastructure	No scale operation is included in handback
After handback	A separate reviewed edit may change pool settings

Re-read the count when the change executes

If ownership or access changed, stop and review again. Do not apply a count from a stale form.

View pool workspace

Leave external ownership

Reference: Ankra infrastructure Nodes on portal develop 35cfefc21; live Next Ink shell. Fictional observations, prices and names. The backend handback adopts observed provider count and must not reuse an old draft. This review makes the external-controller prerequisite explicit; it does not claim to disable it.

STAGING-EU

Nodes

Imported Kubernetes example · Ankra observes nodes; another system provisions machines

Machines

Observed nodes

3

From Kubernetes

Ready nodes

3 / 3

Observed 10:40 UTC

Allocatable memory

46 GiB

Not measured usage

Machine ownership

External

Provisioning outside Ankra

NODE	ROLE	SCHEDULING	KUBELET
imported-control-01	Control plane	Available	v1.33.2
imported-worker-01	Worker	Available	v1.33.2
imported-worker-02	Worker	Available	v1.33.2

Capacity changes belong to the infrastructure owner

Use Ankra to inspect workload placement and node conditions. Add or resize machines in the system that owns this cluster.

STAGING-EU

Machine data is only partly available

Provider inventory read succeeded · Kubernetes agent is disconnected

- Machines
- Node groups
- Control plane
- Bastion & VMs

Current Kubernetes health is unknown

The provider still reports these machines as running. The last Kubernetes observation was 18 minutes ago.

MACHINE	PROVIDER / OBSERVED	LAST KUBERNETES	CURRENT READINESS
general-03	Running · 10:48 UTC	Ready · 10:30 UTC	Unknown
control-plane-01	Running · 10:48 UTC	Ready · 10:30 UTC	Unknown
bastion-01	Running · 10:48 UTC	Not a Kubernetes node	Not applicable

Do not offer a drain, restart or scale-down based on stale workload evidence. Diagnose the access path first.

Inspect access path

View sample recovery

STAGING-EU

The cluster cannot be reached

Last recorded nodes · 10:30 UTC · Agent disconnected

Current readiness is unknown

These are the last recorded observations. The agent is offline; current conditions may have changed.

NODE	LAST READINESS	LAST SCHEDULING	OBSERVED
worker-03	Ready	Available	10:30 UTC
worker-04	Unknown	Available	10:30 UTC

Maintenance actions are unavailable until fresh evidence returns.

View reconnected example

View access-change state

STAGING-EU

No operations match this view

Failed · Last hour · Northstar / staging-eu



Nothing matches these filters.

Clear the status or time range to see other recorded operations.

Clear filters

First-use variant

When there are no records at all: “Nothing has run against this cluster yet.” Do not imply a successful health check.

- N

Northstar

Admin
- Q

Find

⌘K
- <

Clusters
- staging-eu
- 🏠

Overview
- 🔲

Add-ons
- 📁

Stacks
- 📄

Operations
- 🔲

Pods
- 📁

Nodes
- >_

Kubectl access
- ⚙️

Settings

STAGING-EU

Operations could not be read

Northstar / staging-eu · Request failed

The platform could not load operation history

Keep cluster context visible. Retry the read; do not run a deployment as a diagnostic probe.

No operation status has been inferred.

A successful retry restores the list. A denied request moves to the access-changed state.

View successful retry

View access change

Treat service errors independently from empty results. Loading is a neutral skeleton, never fabricated counts.

Your access has changed

We can no longer display this cluster in the current session.

Protected node and operation data has been cleared.

Reload to establish your identity, organisation and current permissions.

[View reloaded example](#)

Storyboard navigation only. No session or permissions are changed.