

Modelling Microservices

Petter Måhlén
February 2016



About me

- Building software professionally since 1995
- Infrastructure at Spotify since 2013
 - Service Discovery (Nameless)
 - Service Framework (<http://spotify.github.io/apollo/>)
 - System-Z
 - Currently doing Data infrastructure

About this talk

1. Why model microservices?

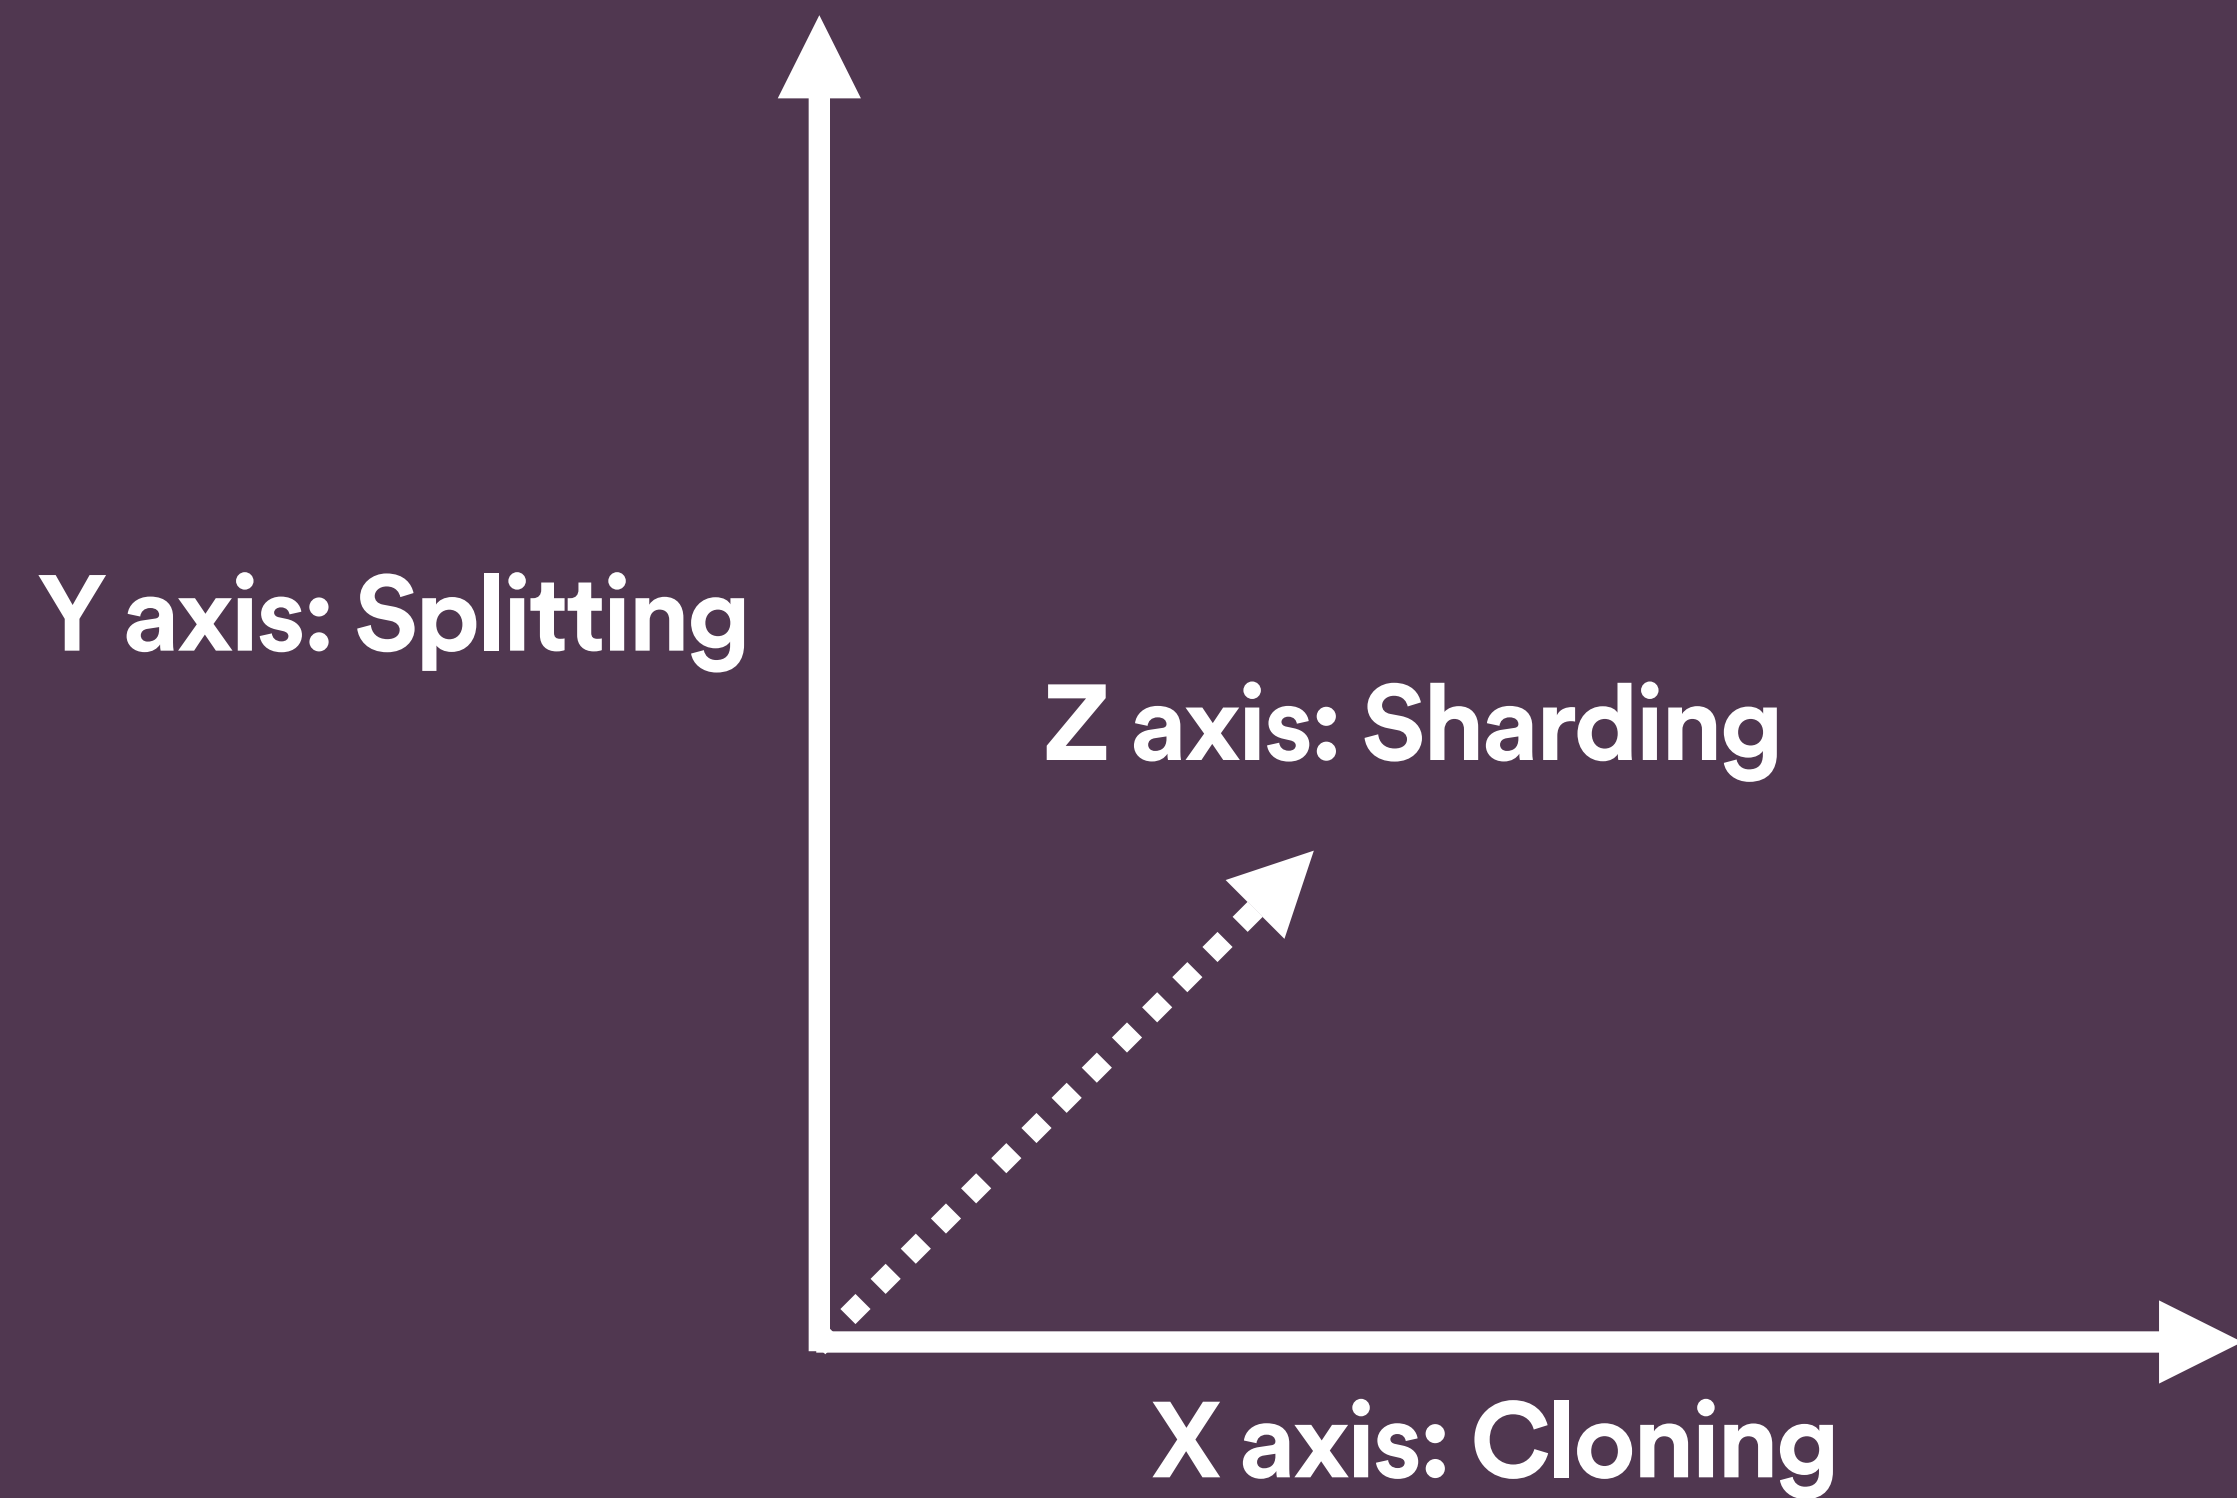
2. Our solution: System-Z

3. Design

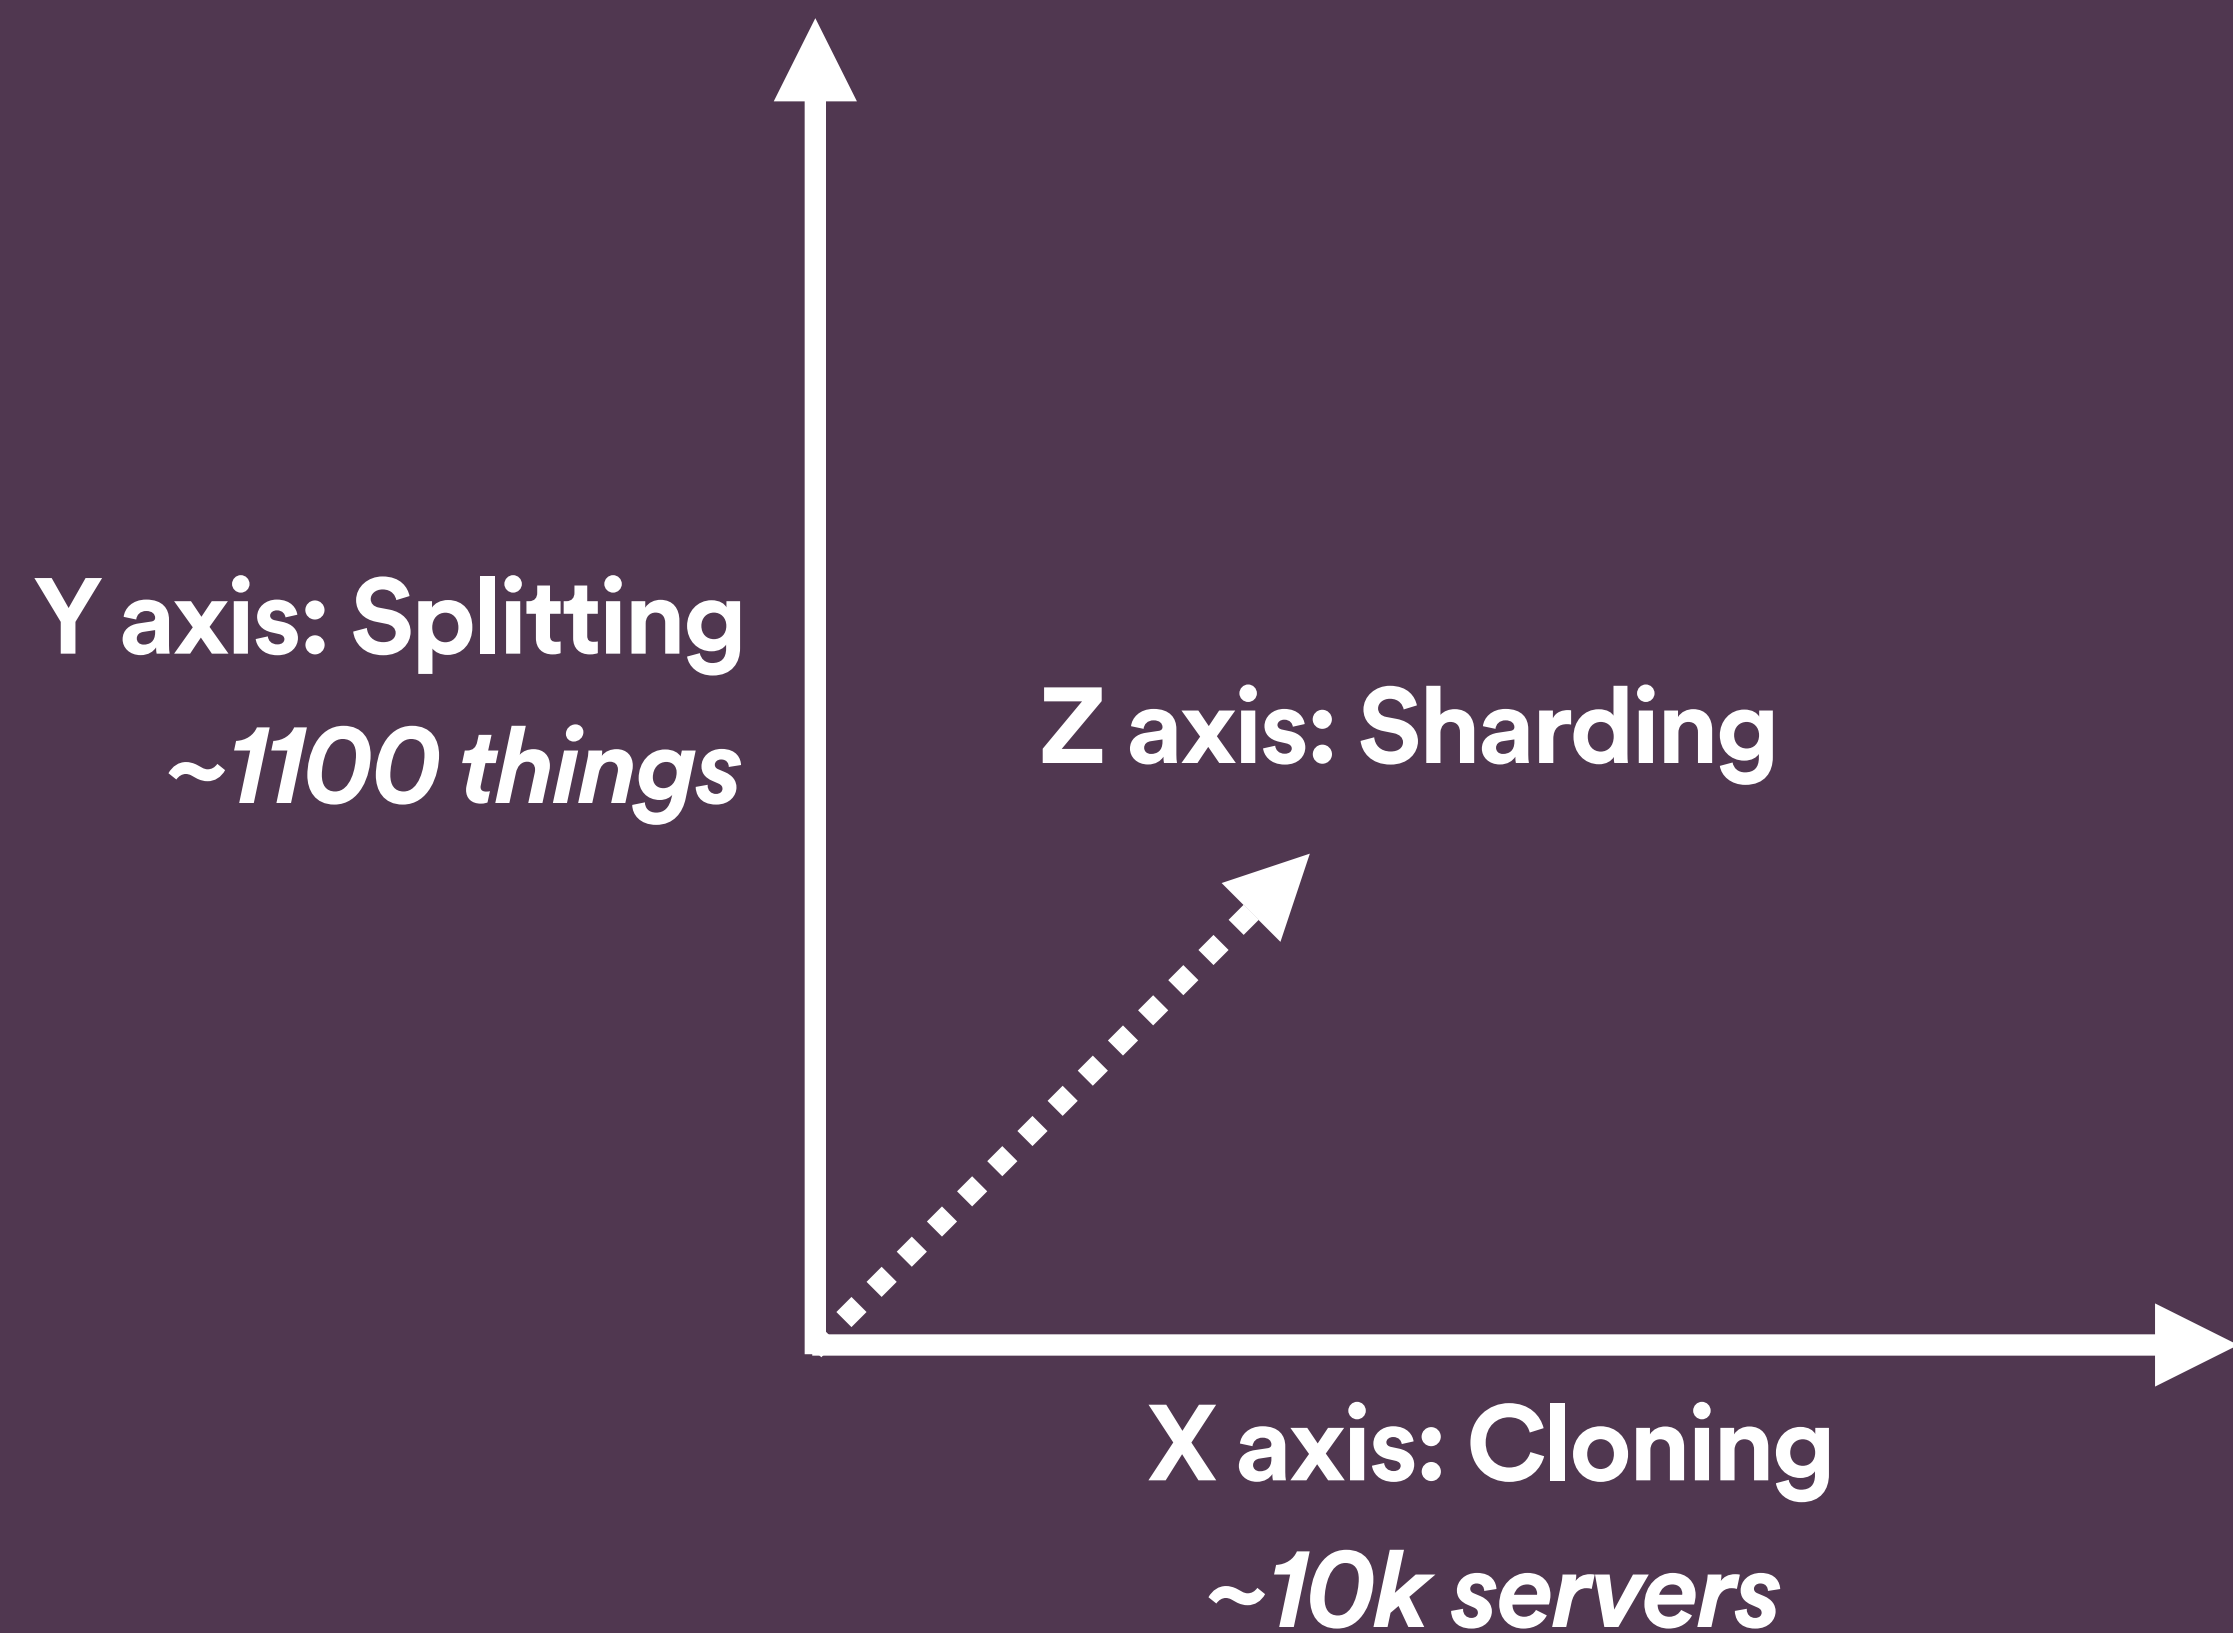
4. Impact so far

=> Ideas about running microservices at scale

Why model microservices?

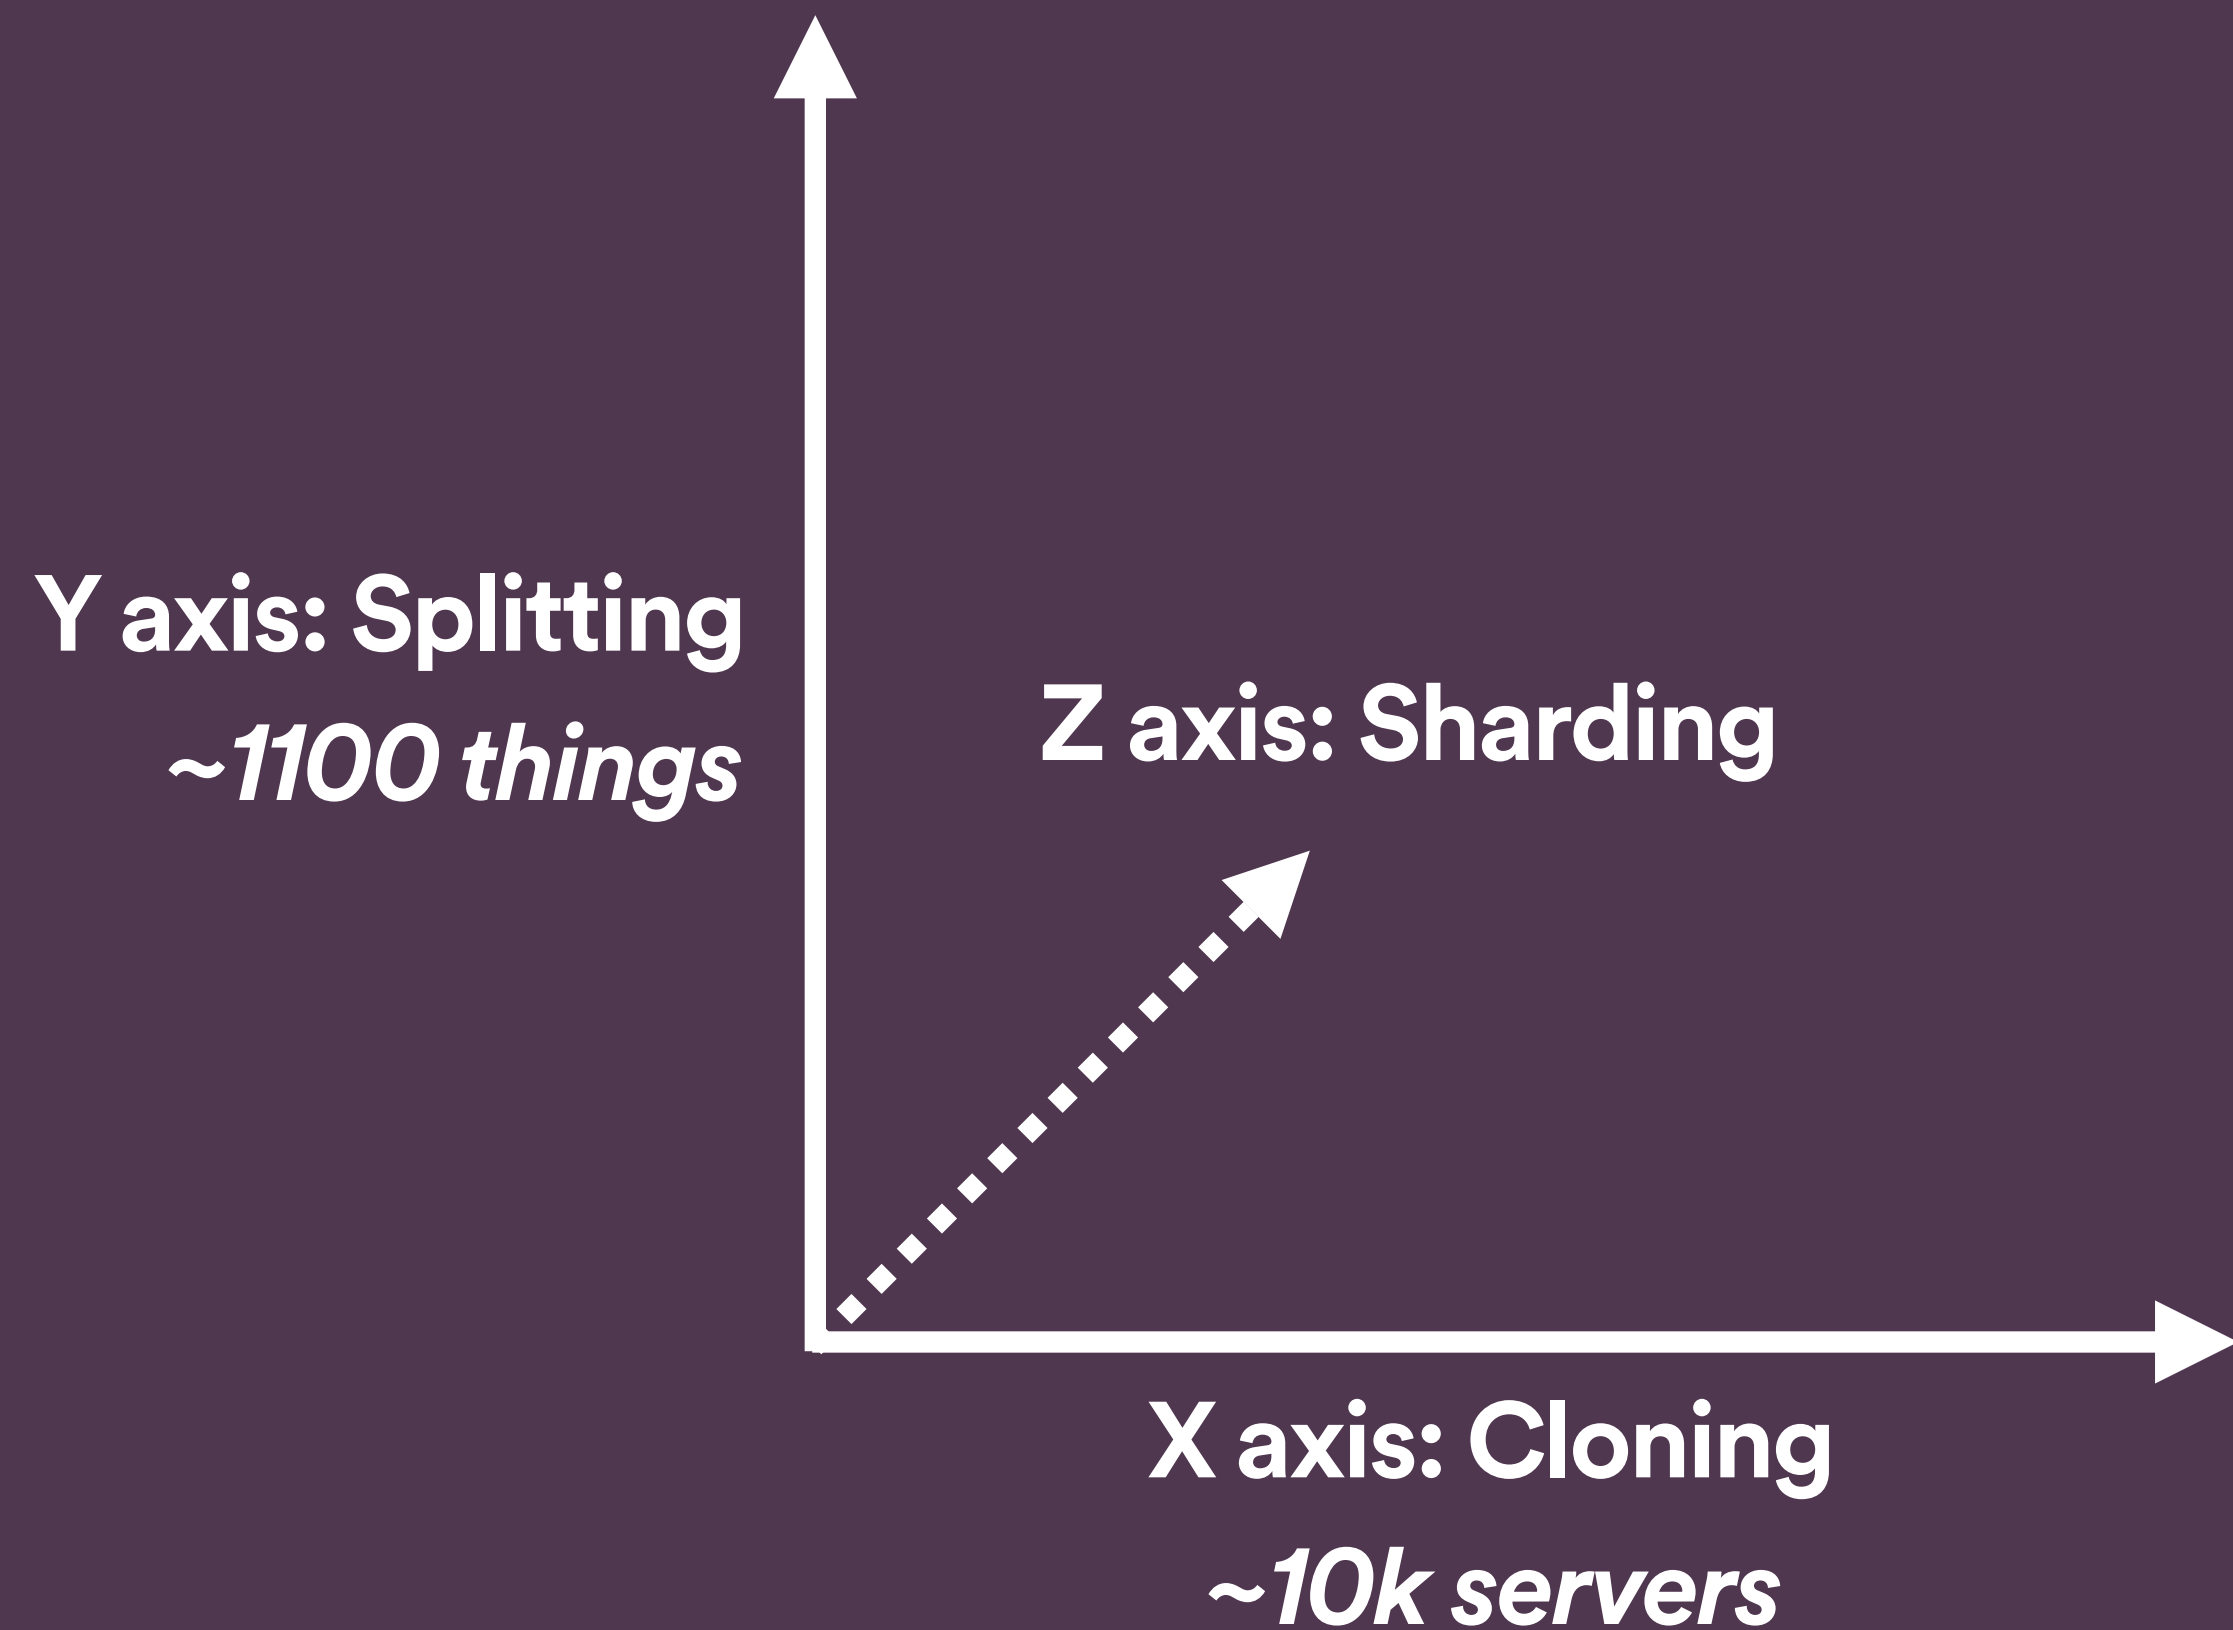


Why model microservices?



<http://artofscalability.com/>


Why model microservices?




plus ~100 teams writing code

<http://artofscalability.com/>


Infrastructure squads



Client platform squads

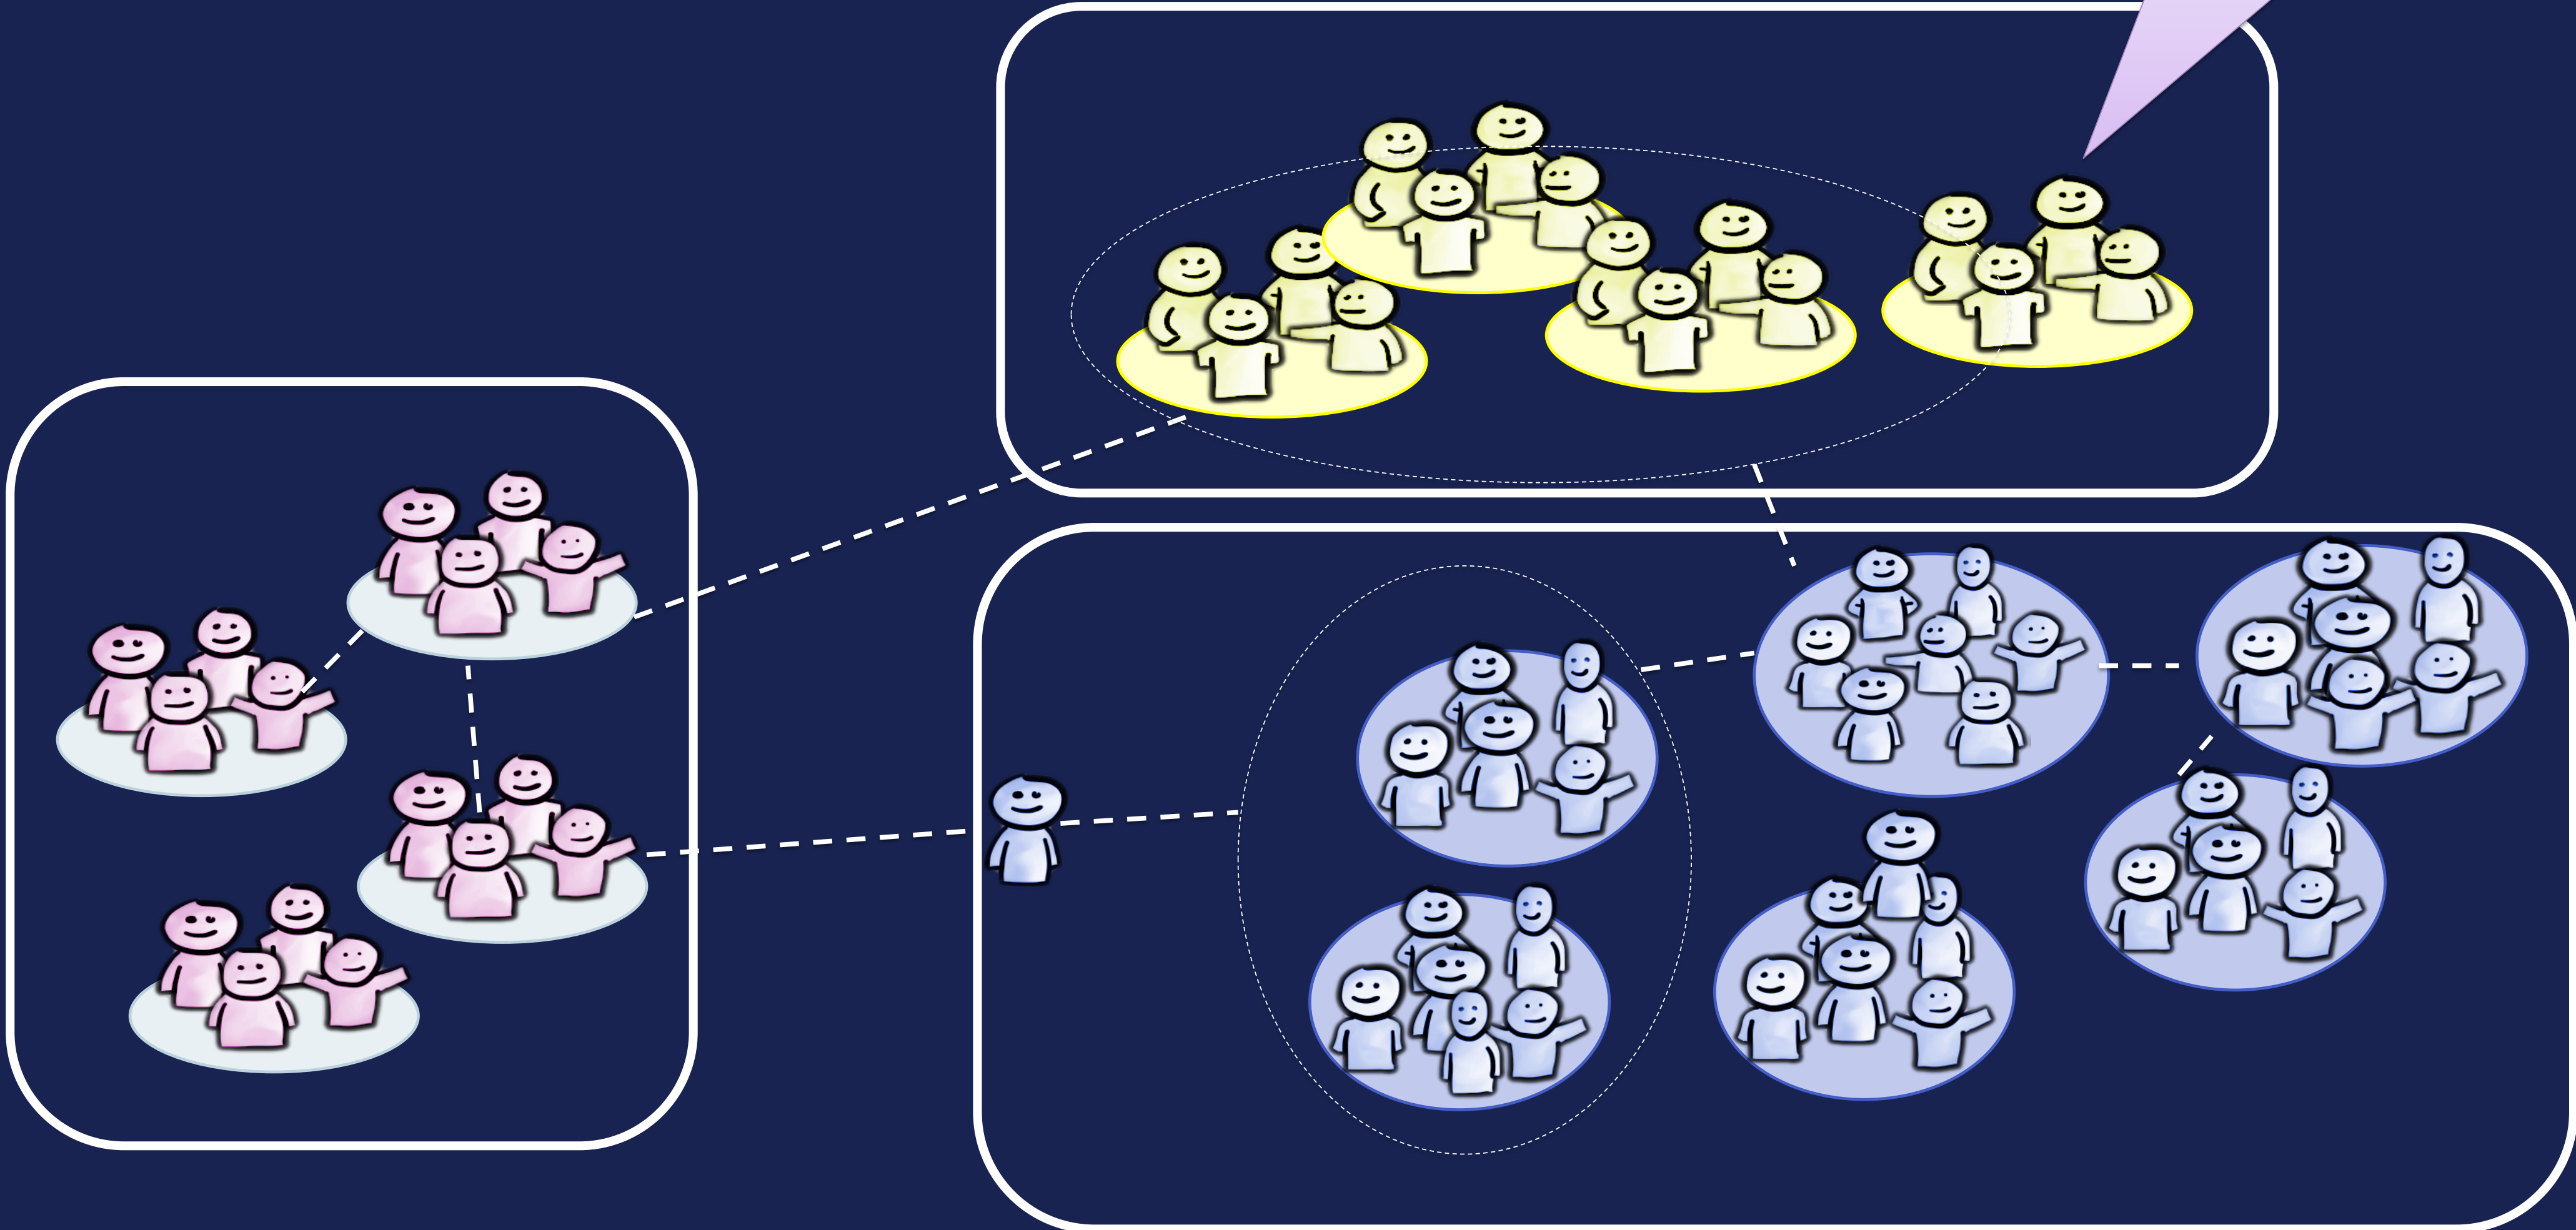


Feature squads



Organic structure

If you need to know exactly who is making decisions, you are in the wrong place



Problems to Solve


Discovering and understanding:

- What things are “out there”
- Deployments and configurations
- The system as a whole, how do things fit together?
- How to get more information: ownership
- What’s broken and how to fix it.

Our solution



What came before



Emil

What came before




Emil



ServiceDB

What came before



Emil



ServiceDB



System-Z

Enter System-Z

- A system for system metadata
- Z? Wat?

The screenshot shows the Spotify System-Z dashboard. At the top, there is a search bar with the text "Go to component...". Below the search bar, a green banner displays the text: "There are **1134** registered components in System-Z, of which **107** are waiting to be migrated from serviceDB. Migration is **90.56%** complete!". The main heading is "System-Z" with a subtitle: "A tool for interacting with and managing Spotify systems and components. [More info.](#)".

The main content area is titled "Systems & components I own" and contains a table with the following columns: SYSTEM, COMPONENTS, and STATUS.

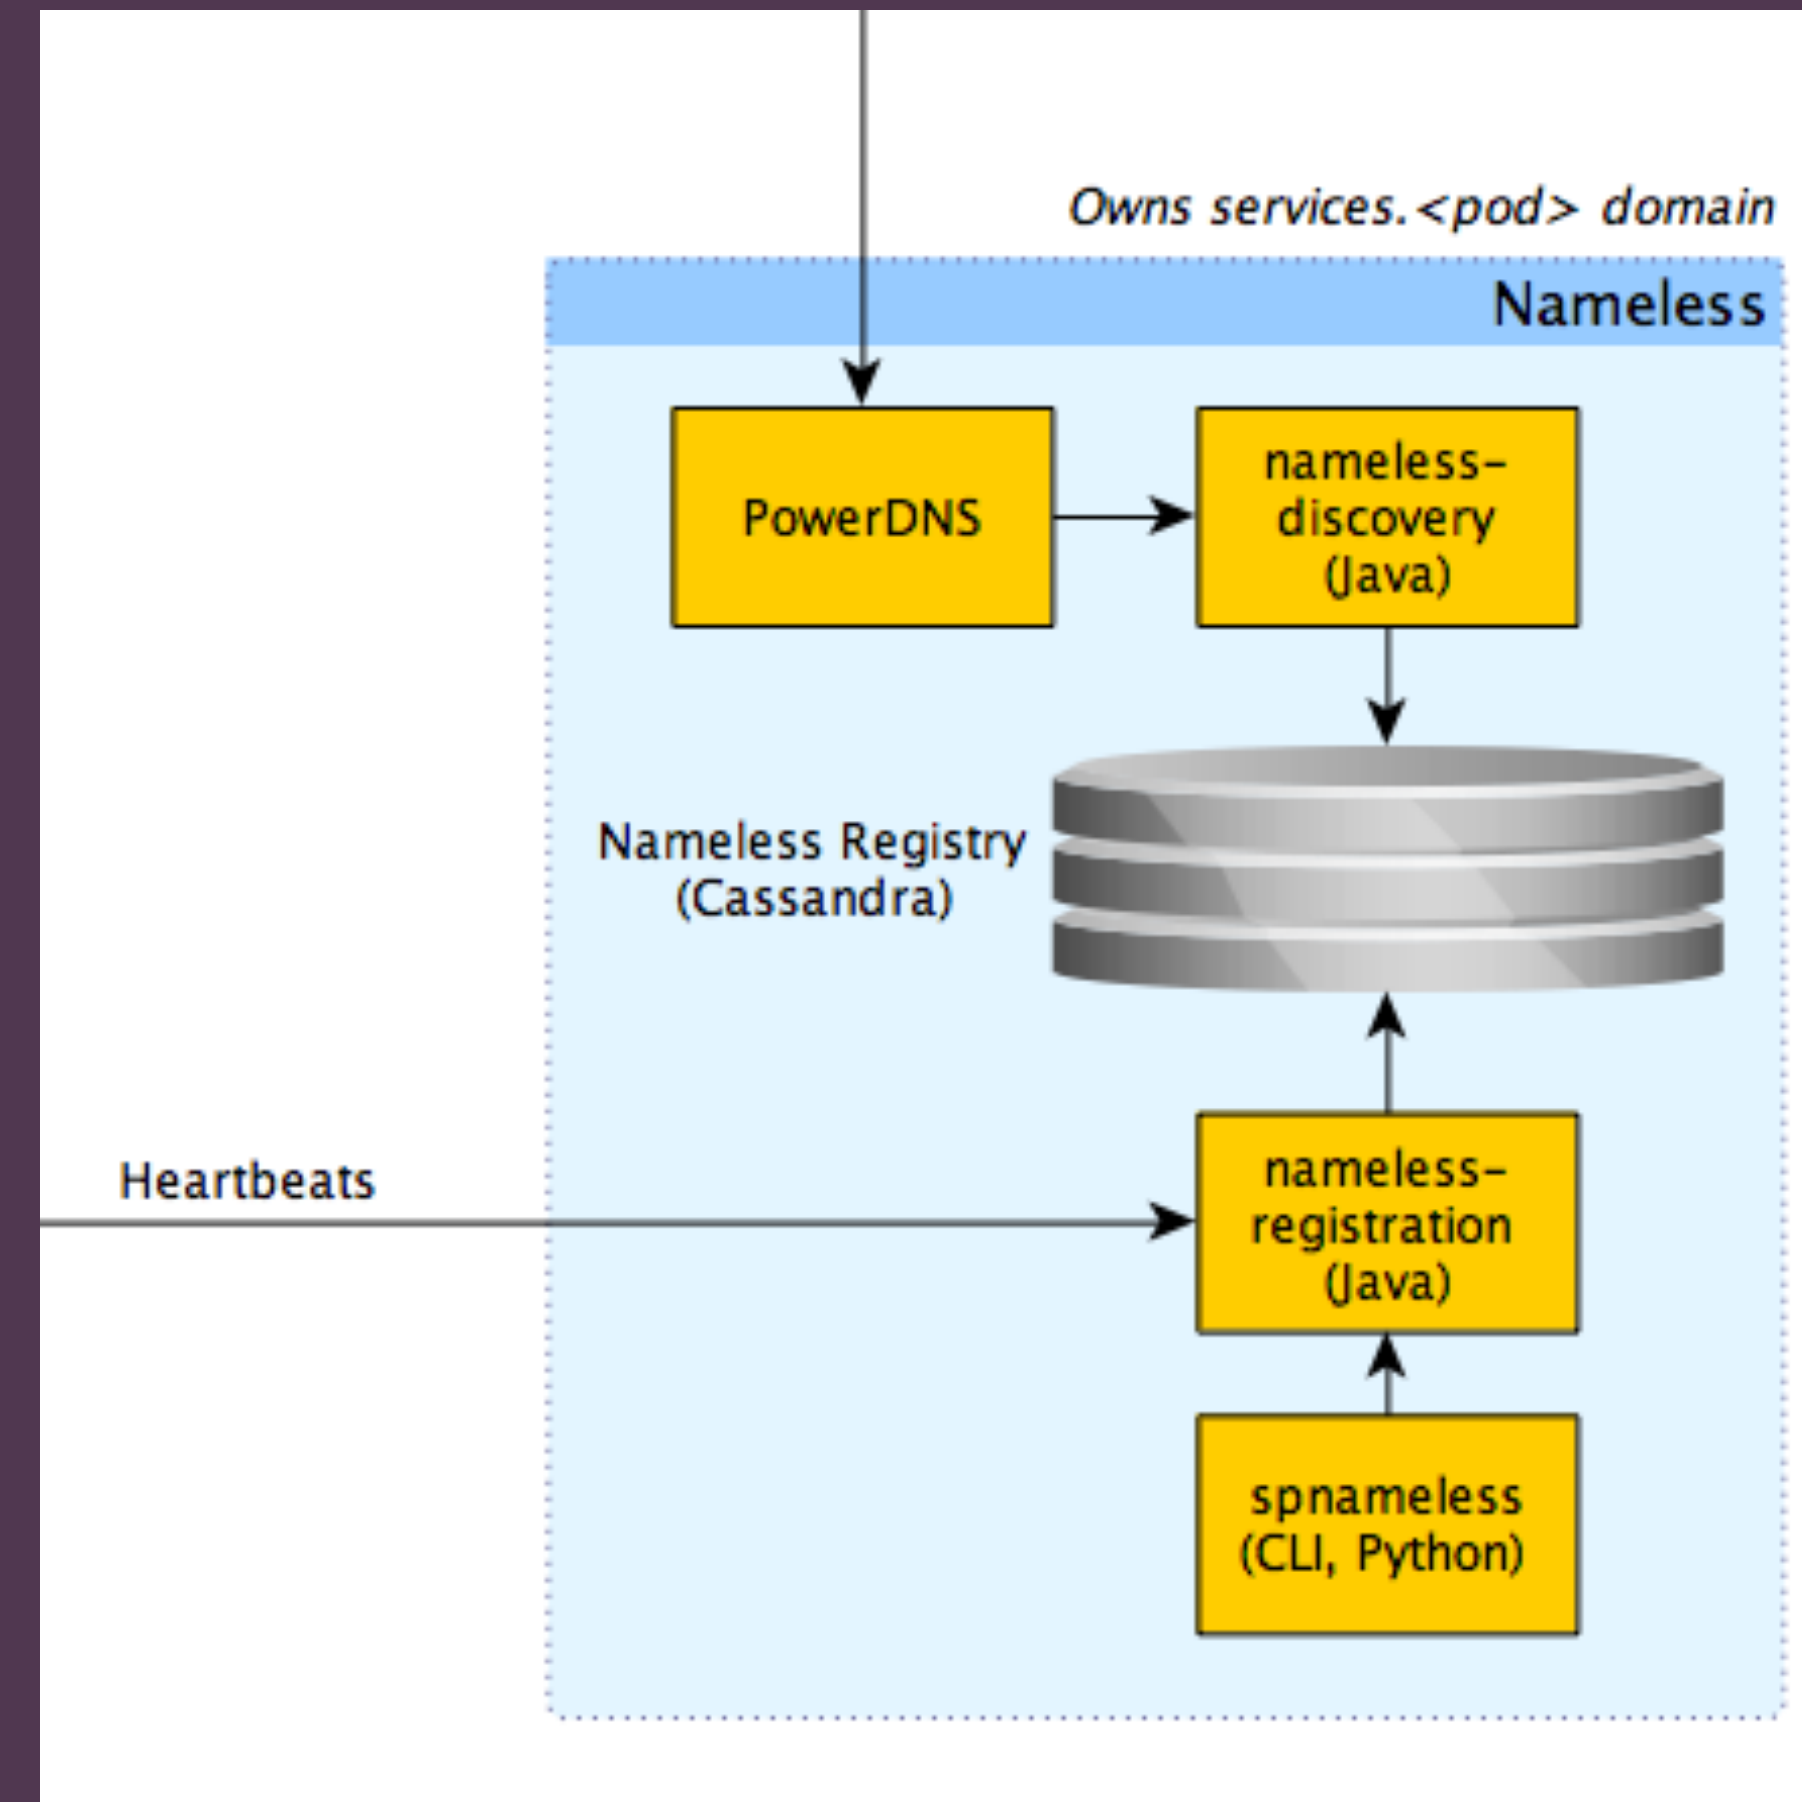
SYSTEM	COMPONENTS	STATUS
apollo	apollo	
deprecated-datainfra-pipelines	shrekcore	✓
gabo	fortress-of-solitude	✓
	gabelastic	✓
	gabo-harvester	
	gabobroker	✓
	gaboeventservice	✓
	gabofiletailer	✓
	gaboheliosagent	
	gabomirror	✓
	gabomonitor	✓
	gaboperftest	✓
	gabozookeeper	✓
	psconsumerexp	✓
	pubsub-drain	
pubsub2hdfs	✓	

On the right side of the dashboard, there are three sections: "CREATE / ADD" with options "Create a new Component" and "Register an existing service/component"; "LEARN" with links "What is System-Z?" and "How do I define components in YAML?"; and "GIVE US FEEDBACK" with options "Suggest a feature", "Email us: tools@spotify.com", and "Find us on Slack: #tools".

Some terminology

- **Component** - some thing
 - A microservice, data store, data pipeline, or library
- **System** - some related things

These terms, like most things Z, are intentionally vague



Main features

- **General component information**
- **Dependency Tracking**
- **Managing deployments**
- **Ownership and alerting**



searchview

OVERVIEW **ACTIVITY *a*** DEPLOYMENT *a* PIPELINE HISTORY *a* ROUTING CONFIGURATION API

Apollofied aggregation search view service

DOCUMENTATION

<https://ghe.spotify.net/kiwi/searchview/blob/master/README.md>

INFO

owner	kiwi
system	search
visibility	private
service_discovery	searchview
slack_channel	kiwi
tags	kiwi apollo search
pagerdutykey	
role	searchview
artifact_id	searchview-alone
maintainers	backend3

EDIT

REFRESH

What fields can I edit?

No ongoing incidents

pd



Escalation policy

DEPLOYMENT [HERMES]

searchview 0.2.74 76 instances

DEPENDENCIES

charts-api-backend	declared + runtime ●
chartview	declared + runtime ●
collection (collection2)	declared + runtime ●
facebook	declared + runtime ●
identity (identity2)	declared + runtime ●
mergedprofile (socialgraph3)	declared + runtime ●
metadata (metadata-proxy)	declared + runtime ●
playlist -> playlist-proxy playlist4	declared + runtime ●
playlist-mosaic	declared + runtime ●
popcount (popcount2)	declared + runtime ●
search	declared + runtime ●



searchview

- OVERVIEW
- ACTIVITY *a*
- DEPLOYMENT *a*
- PIPELINE HISTORY *a*
- ROUTING
- CONFIGURATION
- API

Apollofied aggregation search view service

DOCUMENTATION

<https://ghe.spotify.net/kiwi/searchview/blob/master/README.md>

INFO

owner	kiwi
system	search
visibility	private
service_discovery	searchview
slack_channel	kiwi
tags	kiwi apollo search
pagerdutykey	
role	searchview
artifact_id	searchview-alone
maintainers	backend3

EDIT

REFRESH

What fields can I edit?

No ongoing incidents pd

Escalation policy

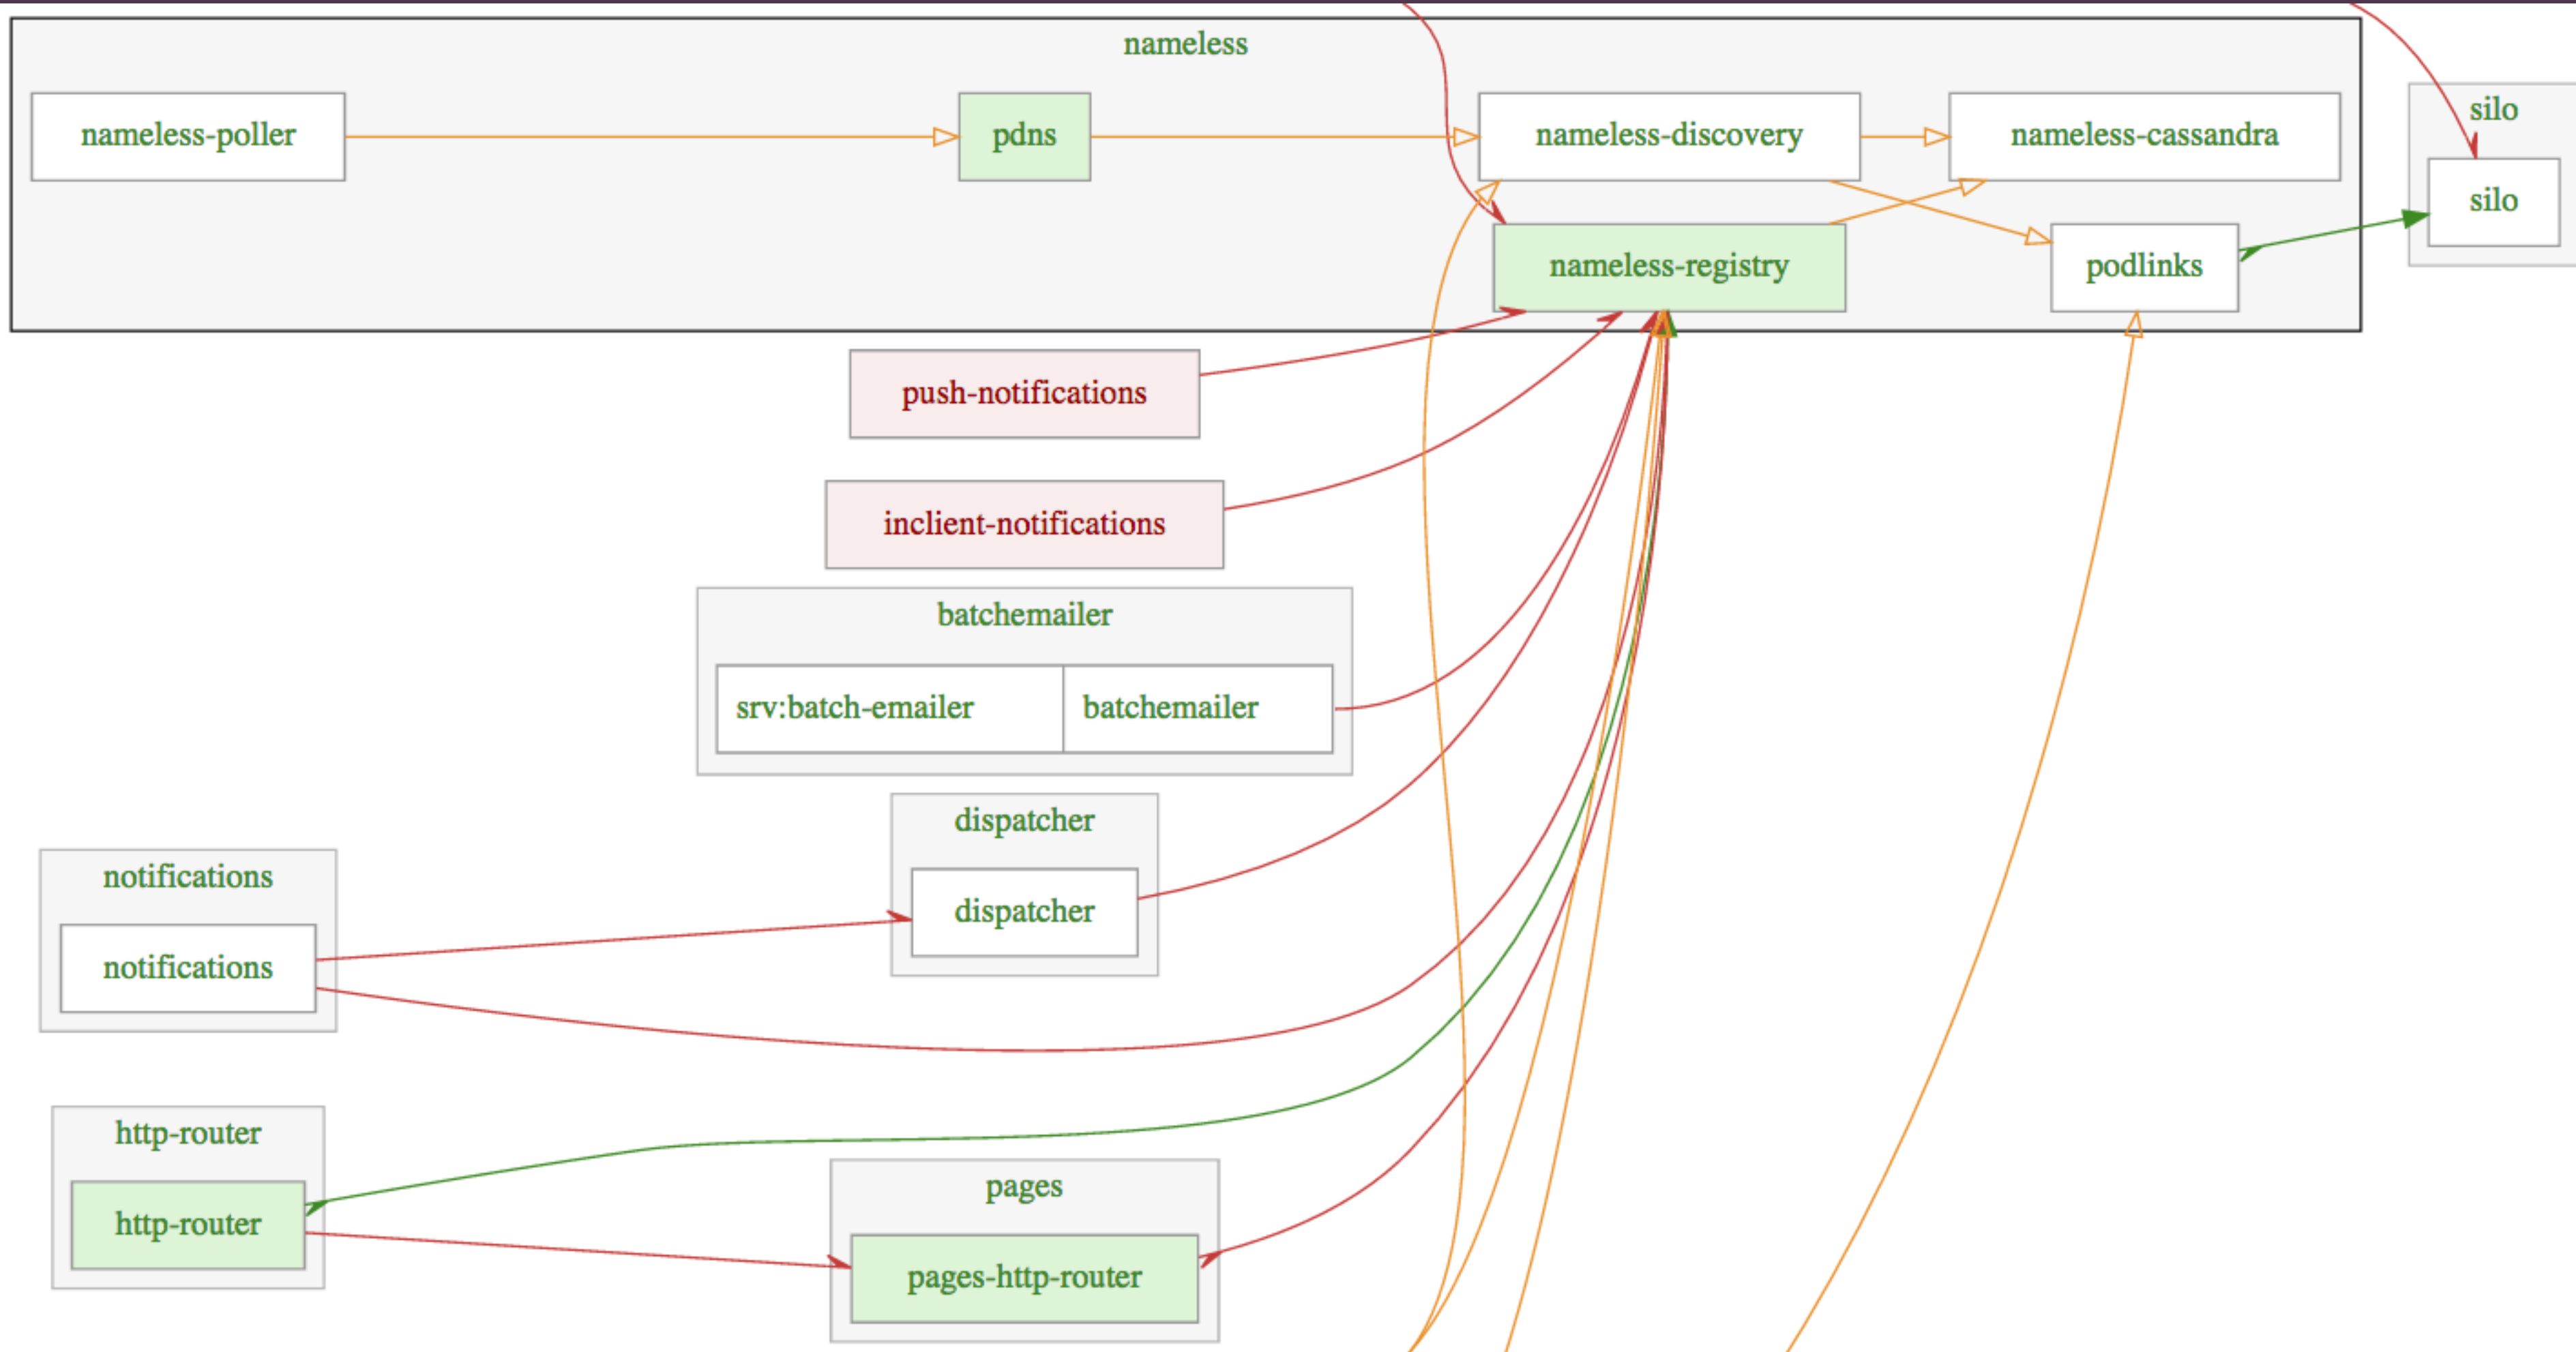
DEPLOYMENT [HERMES]

searchview 0.2.74

76 instances

DEPENDENCIES

charts-api-backend	declared + runtime ●
chartview	declared + runtime ●
collection (collection2)	declared + runtime ●
facebook	declared + runtime ●
identity (identity2)	declared + runtime ●
mergedprofile (socialgraph3)	declared + runtime ●
metadata (metadata-proxy)	declared + runtime ●
playlist -> playlist-proxy playlist4	declared + runtime ●
playlist-mosaic	declared + runtime ●
popcount (popcount2)	declared + runtime ●
search	declared + runtime ●





searchview

- OVERVIEW
- ACTIVITY *a*
- DEPLOYMENT *a*
- PIPELINE HISTORY *a*
- ROUTING
- CONFIGURATION
- API

Apollofied aggregation search view service

DOCUMENTATION

<https://ghe.spotify.net/kiwi/searchview/blob/master/README.md>

INFO

owner	kiwi
system	search
visibility	private
service_discovery	searchview
slack_channel	kiwi
tags	kiwi apollo search
pagerdutykey	
role	searchview
artifact_id	searchview-alone
maintainers	backend3

EDIT

REFRESH

What fields can I edit?

No ongoing incidents

pd



Escalation policy

DEPLOYMENT [HERMES]

searchview 0.2.74

76 instances

DEPENDENCIES

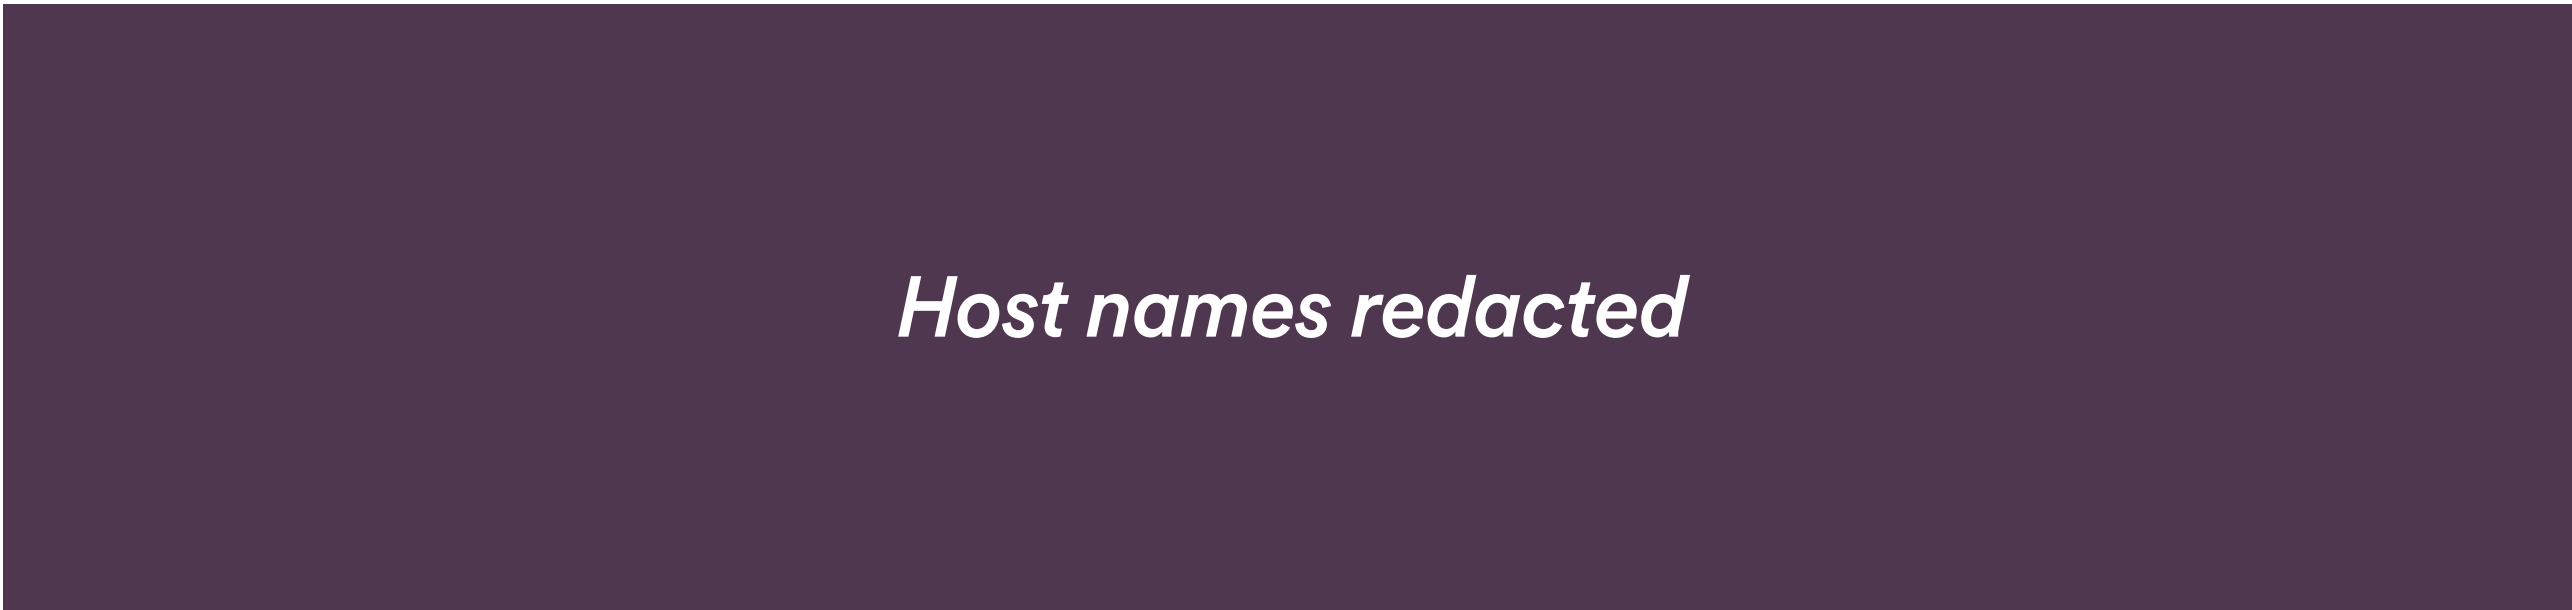
charts-api-backend	declared + runtime
chartview	declared + runtime
collection (collection2)	declared + runtime
facebook	declared + runtime
identity (identity2)	declared + runtime
mergedprofile (socialgraph3)	declared + runtime
metadata (metadata-proxy)	declared + runtime
playlist -> playlist-proxy playlist4	declared + runtime
playlist-mosaic	declared + runtime
popcount (popcount2)	declared + runtime
search	declared + runtime

gaboeventservice

OVERVIEW ACTIVITY **DEPLOYMENT** PIPELINE HISTORY ROUTING CONFIGURATION API

Filters Only show deployed versions Only consider 10 last builds

MACHINES



VERSION

- 0.1.32-SNAPSHOT-c585f70-20160112-0833
(created 2016-01-12 09:35 by jenkins)
- 0.1.30-SNAPSHOT-572b690-20151221-1357
(created 2015-12-21 14:59 by jenkins)
- 0.1.29-SNAPSHOT-0f9029d-20151216-1510
(created 2015-12-16 16:11 by jenkins)
- 0.1.28-SNAPSHOT-f6b8ee7-20151216-1410
(created 2015-12-16 15:12 by jenkins)
- 0.1.27-SNAPSHOT-f66a3e0-20151130-2055
(created 2015-11-30 21:57 by jenkins)
- 0.1.26-SNAPSHOT-9cfa070-20151124-1832
(created 2015-11-24 19:33 by jenkins)
- 0.1.25-SNAPSHOT-8be1ae9-20151124-1715
(created 2015-11-24 18:25 by jenkins)

<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
...
...
...
...

Deploy
Deploy and undeploy others



searchview

OVERVIEW ACTIVITY *a* DEPLOYMENT *a* PIPELINE HISTORY *a* **ROUTING** CONFIGURATION API

Manage routing between pods in our infrastructure. If you need any support with this please talk to [tools@](#) or [#tools](#).

— This view is powered by [podlinks](#)

Configuration for service discovery name **searchview**

searchview

HM

HTTP

NEW SET +



	From	Link	Status
✓	A	→ B	
✓	B		18 running instances

Change history

Time	From	Update	Protocol	User
2 months ago	D	E → A	hm	alexey
2 months ago	D	F → E	hm	alexey
2 months ago	D	F	hm	alexey



searchview

- OVERVIEW
- ACTIVITY *a*
- DEPLOYMENT *a*
- PIPELINE HISTORY *a*
- ROUTING
- CONFIGURATION
- API

Apollofied aggregation search view service

DOCUMENTATION

<https://ghe.spotify.net/kiwi/searchview/blob/master/README.md>

INFO

owner	kiwi
system	search
visibility	private
service_discovery	searchview
slack_channel	kiwi
tags	kiwi apollo search
pagerdutykey	
role	searchview
artifact_id	searchview-alone
maintainers	backend3

EDIT

REFRESH

What fields can I edit?

No ongoing incidents

pd



Escalation policy

DEPLOYMENT [HERMES]

searchview 0.2.74

76 instances

DEPENDENCIES

charts-api-backend	declared + runtime
chartview	declared + runtime
collection (collection2)	declared + runtime
facebook	declared + runtime
identity (identity2)	declared + runtime
mergedprofile (socialgraph3)	declared + runtime
metadata (metadata-proxy)	declared + runtime
playlist -> playlist-proxy playlist4	declared + runtime
playlist-mosaic	declared + runtime
popcount (popcount2)	declared + runtime
search	declared + runtime



searchview

- OVERVIEW
- ACTIVITY *a*
- DEPLOYMENT *a*
- PIPELINE HISTORY *a*
- ROUTING
- CONFIGURATION
- API

Apollofied aggregation search view service

DOCUMENTATION

<https://ghe.spotify.net/kiwi/searchview/blob/master/README.md>

INFO

owner	kiwi
system	search
visibility	private
service_discovery	searchview
slack_channel	kiwi
tags	kiwi apollo search
pagerdutykey	<input type="text"/>
role	searchview
artifact_id	searchview-alone
maintainers	backend3

EDIT

REFRESH

What fields can I edit?

No ongoing incidents

pd



Escalation policy

DEPLOYMENT [HERMES]

searchview 0.2.74 76 instances

DEPENDENCIES

- charts-api-backend declared + runtime ●
- chartview declared + runtime ●
- collection (collection2) declared + runtime ●
- facebook declared + runtime ●
- identity (identity2) declared + runtime ●
- mergedprofile (socialgraph3) declared + runtime ●
- metadata (metadata-proxy) declared + runtime ●
- playlist -> declared + runtime ●
 - playlist-proxy
 - playlist4
- playlist-mosaic declared + runtime ●
- popcount (popcount2) declared + runtime ●
- search declared + runtime ●

And...

- **Create new components**
- **Compare live configs**
- **Check activity (logins, system updates, etc)**
- **Check Global stats (Apollo versions, etc)**
- **API documentation**
- **Provision hardware**
- **Set up service monitoring dashboards**
- **Configure data processing pipeline monitoring**
- **...**

Design

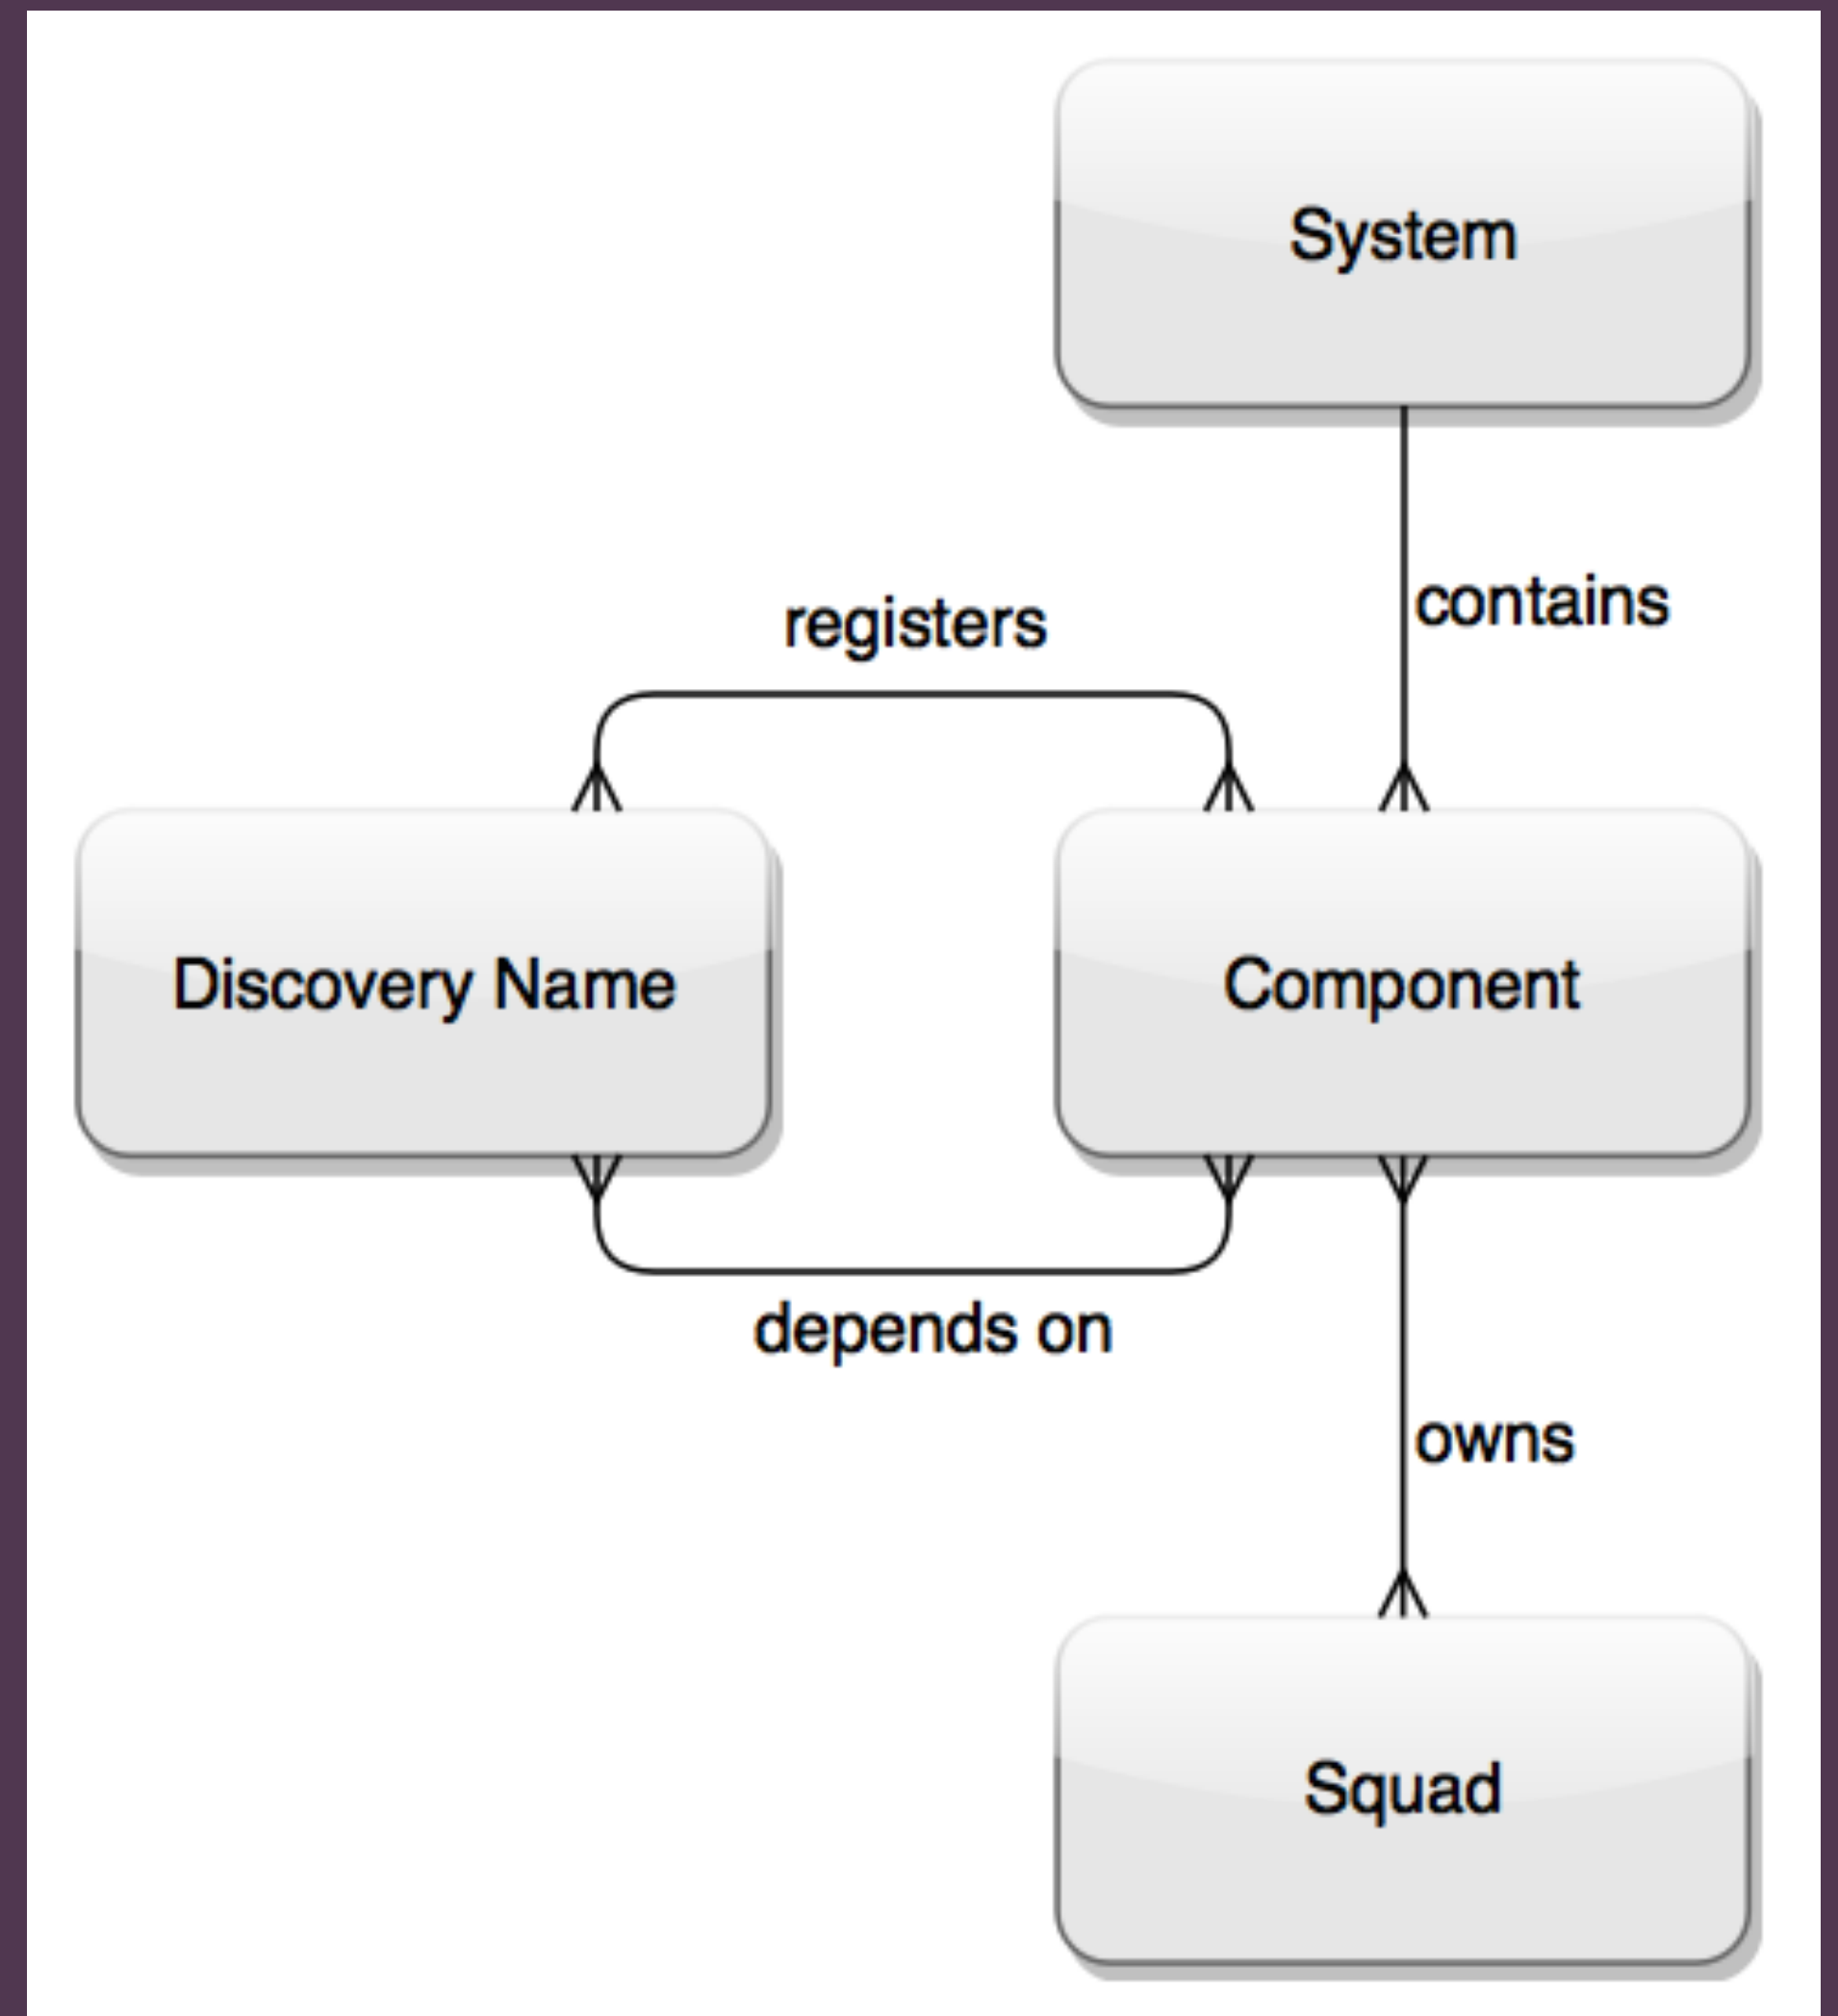


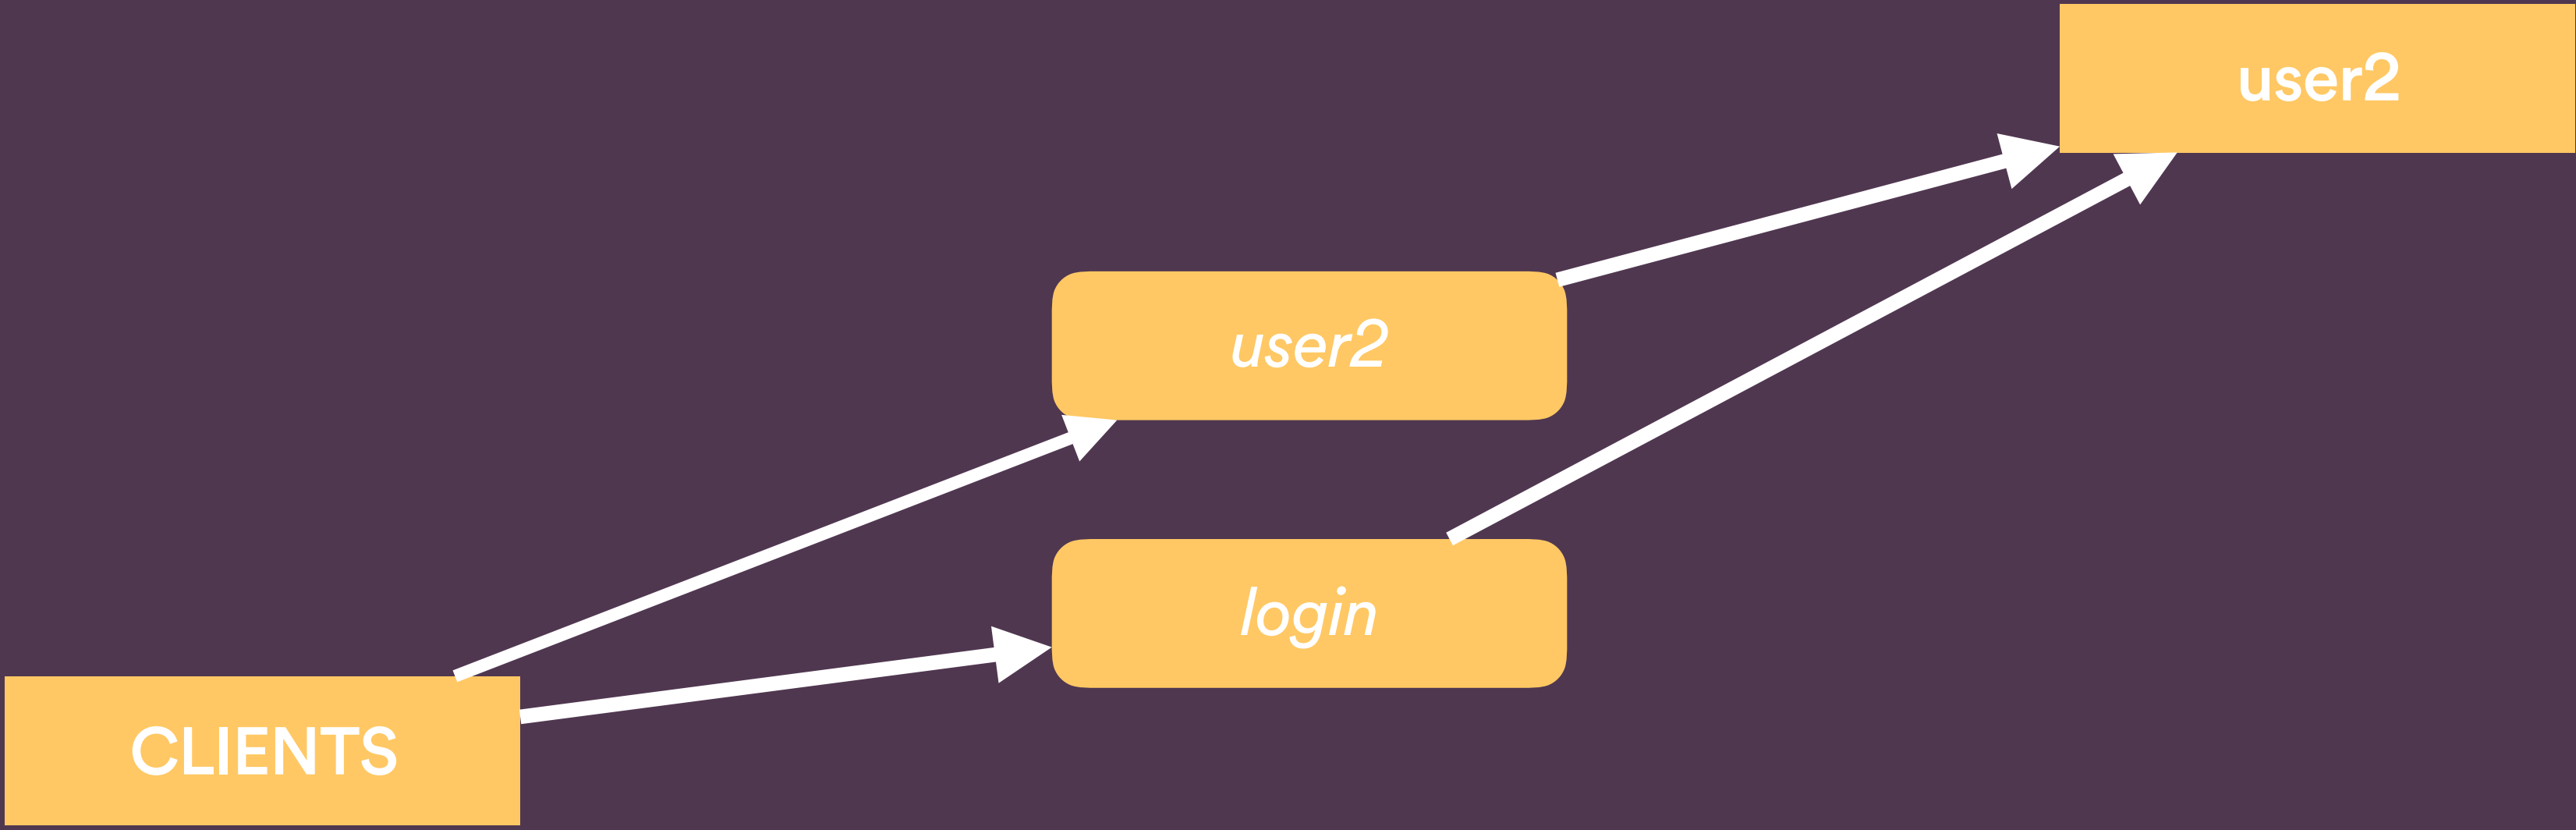
Core data model

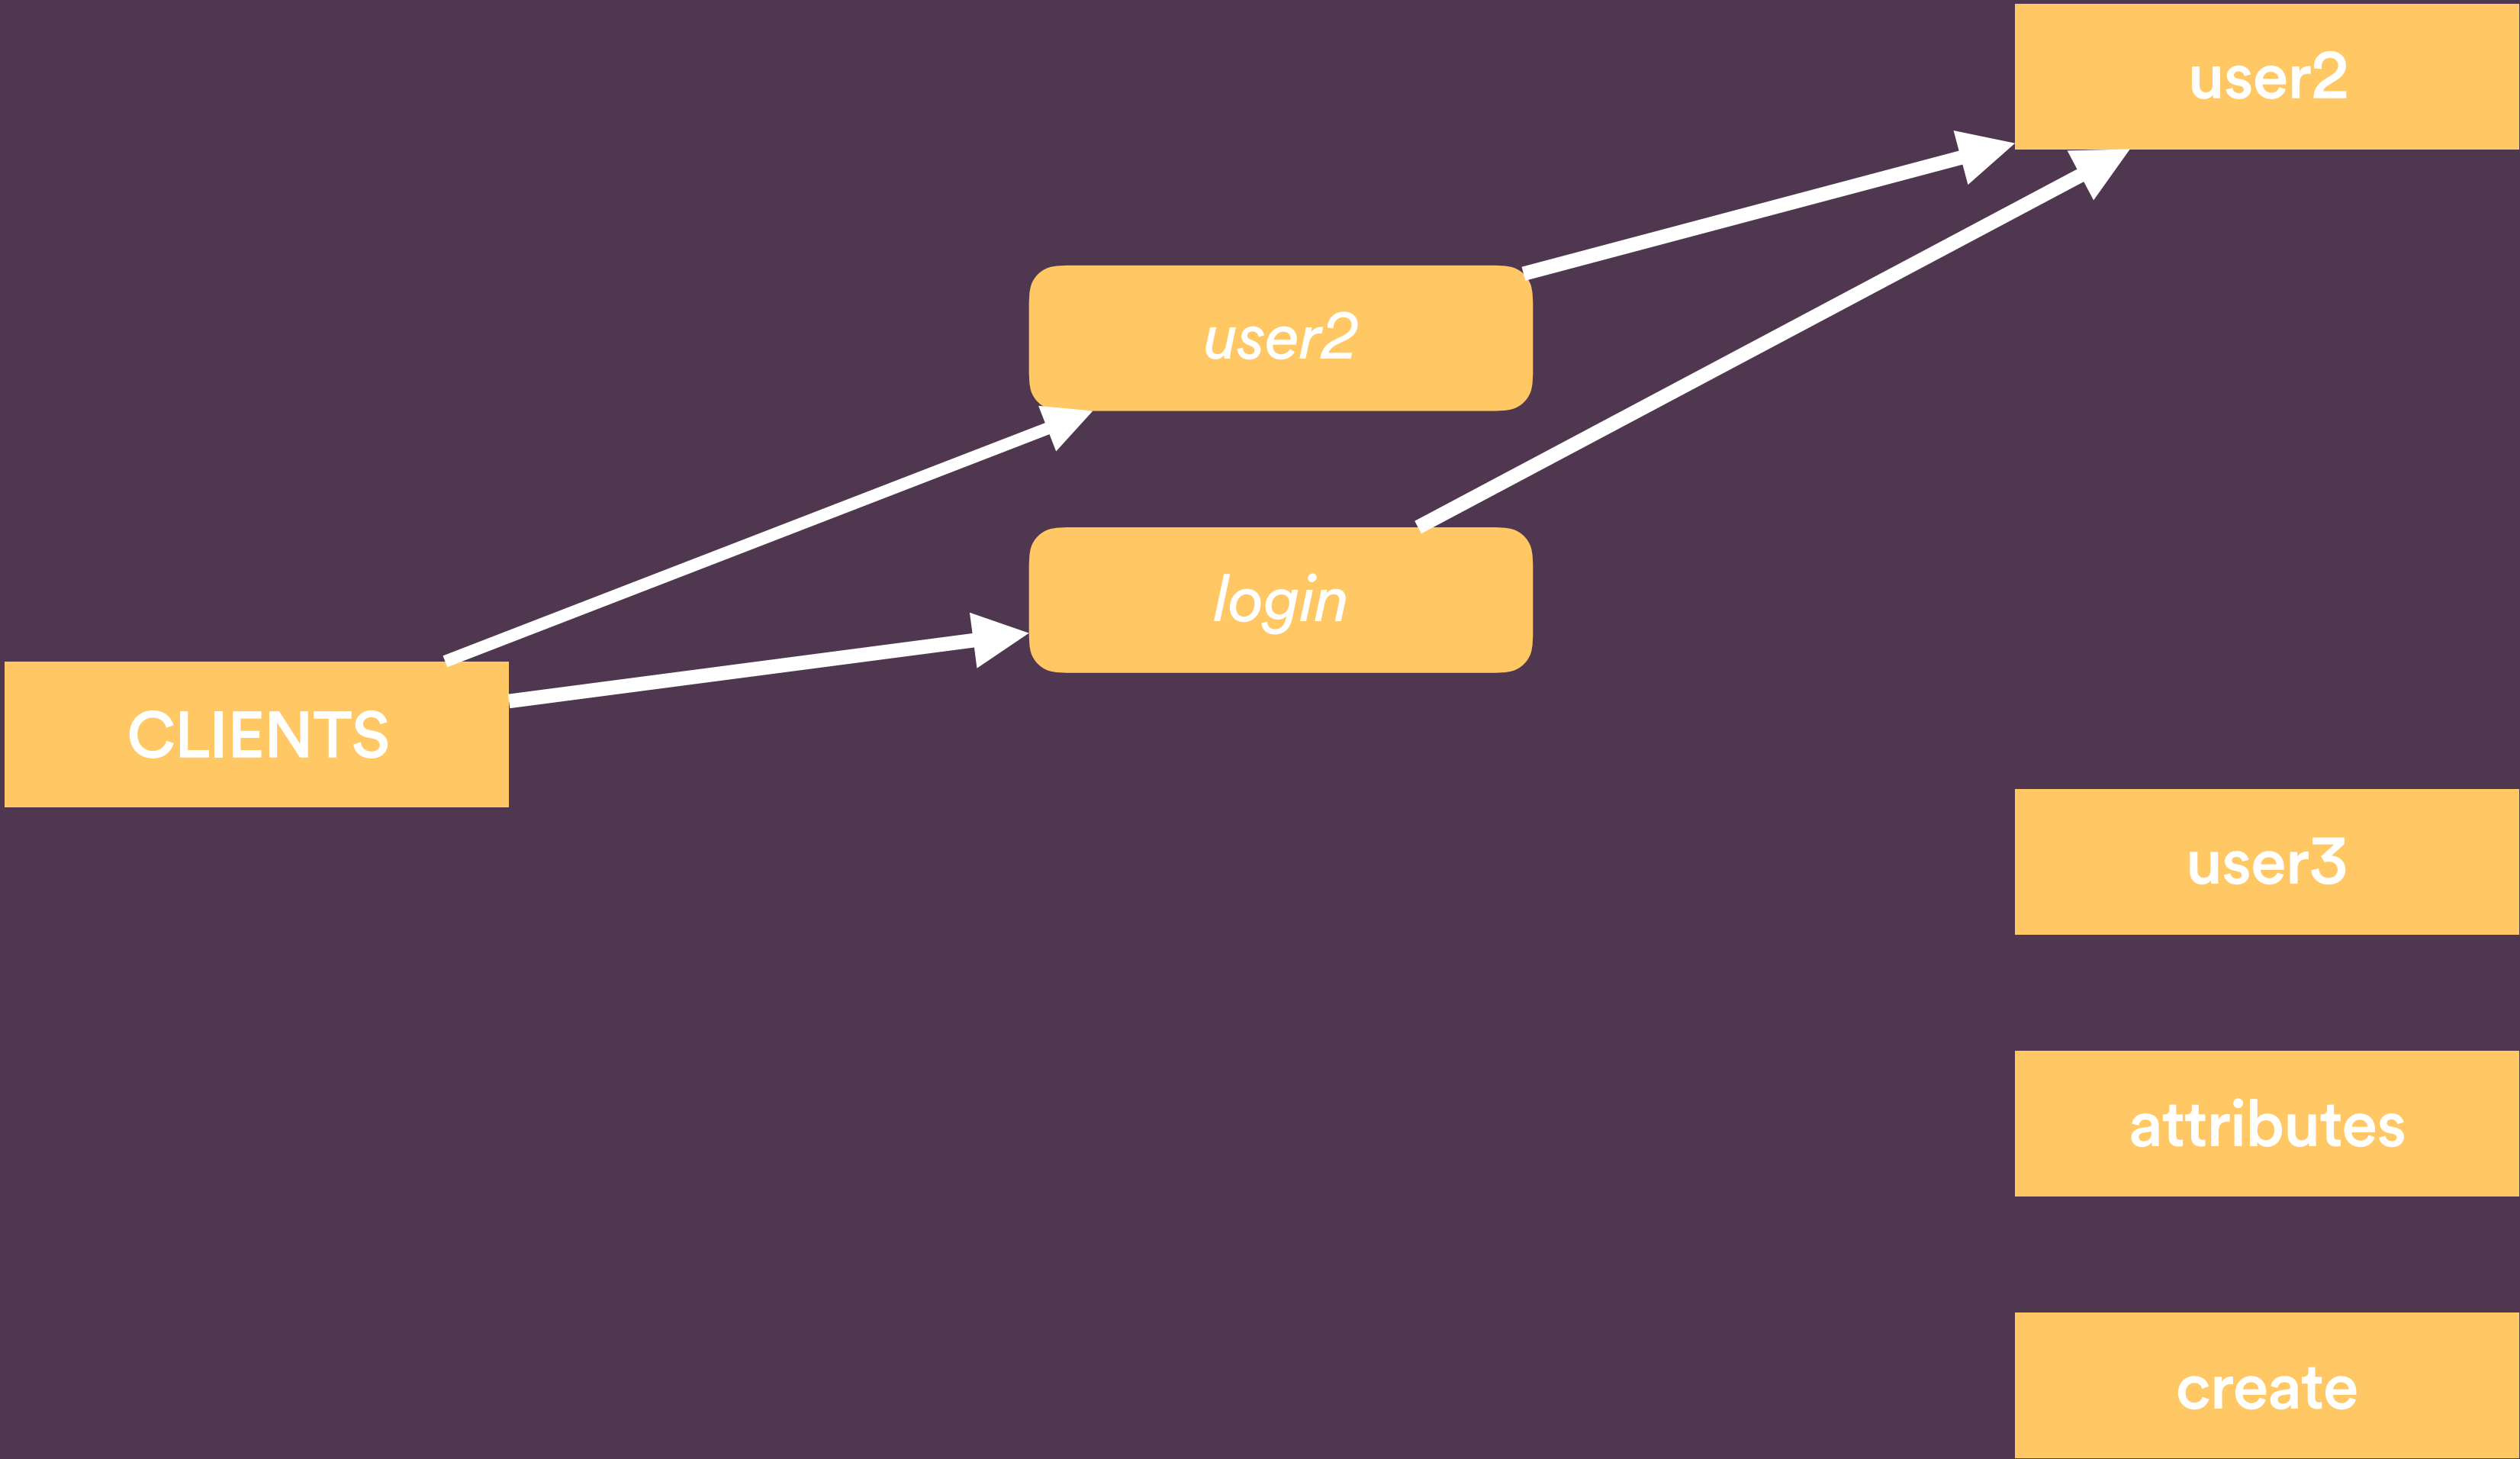
Many many-to-many relations

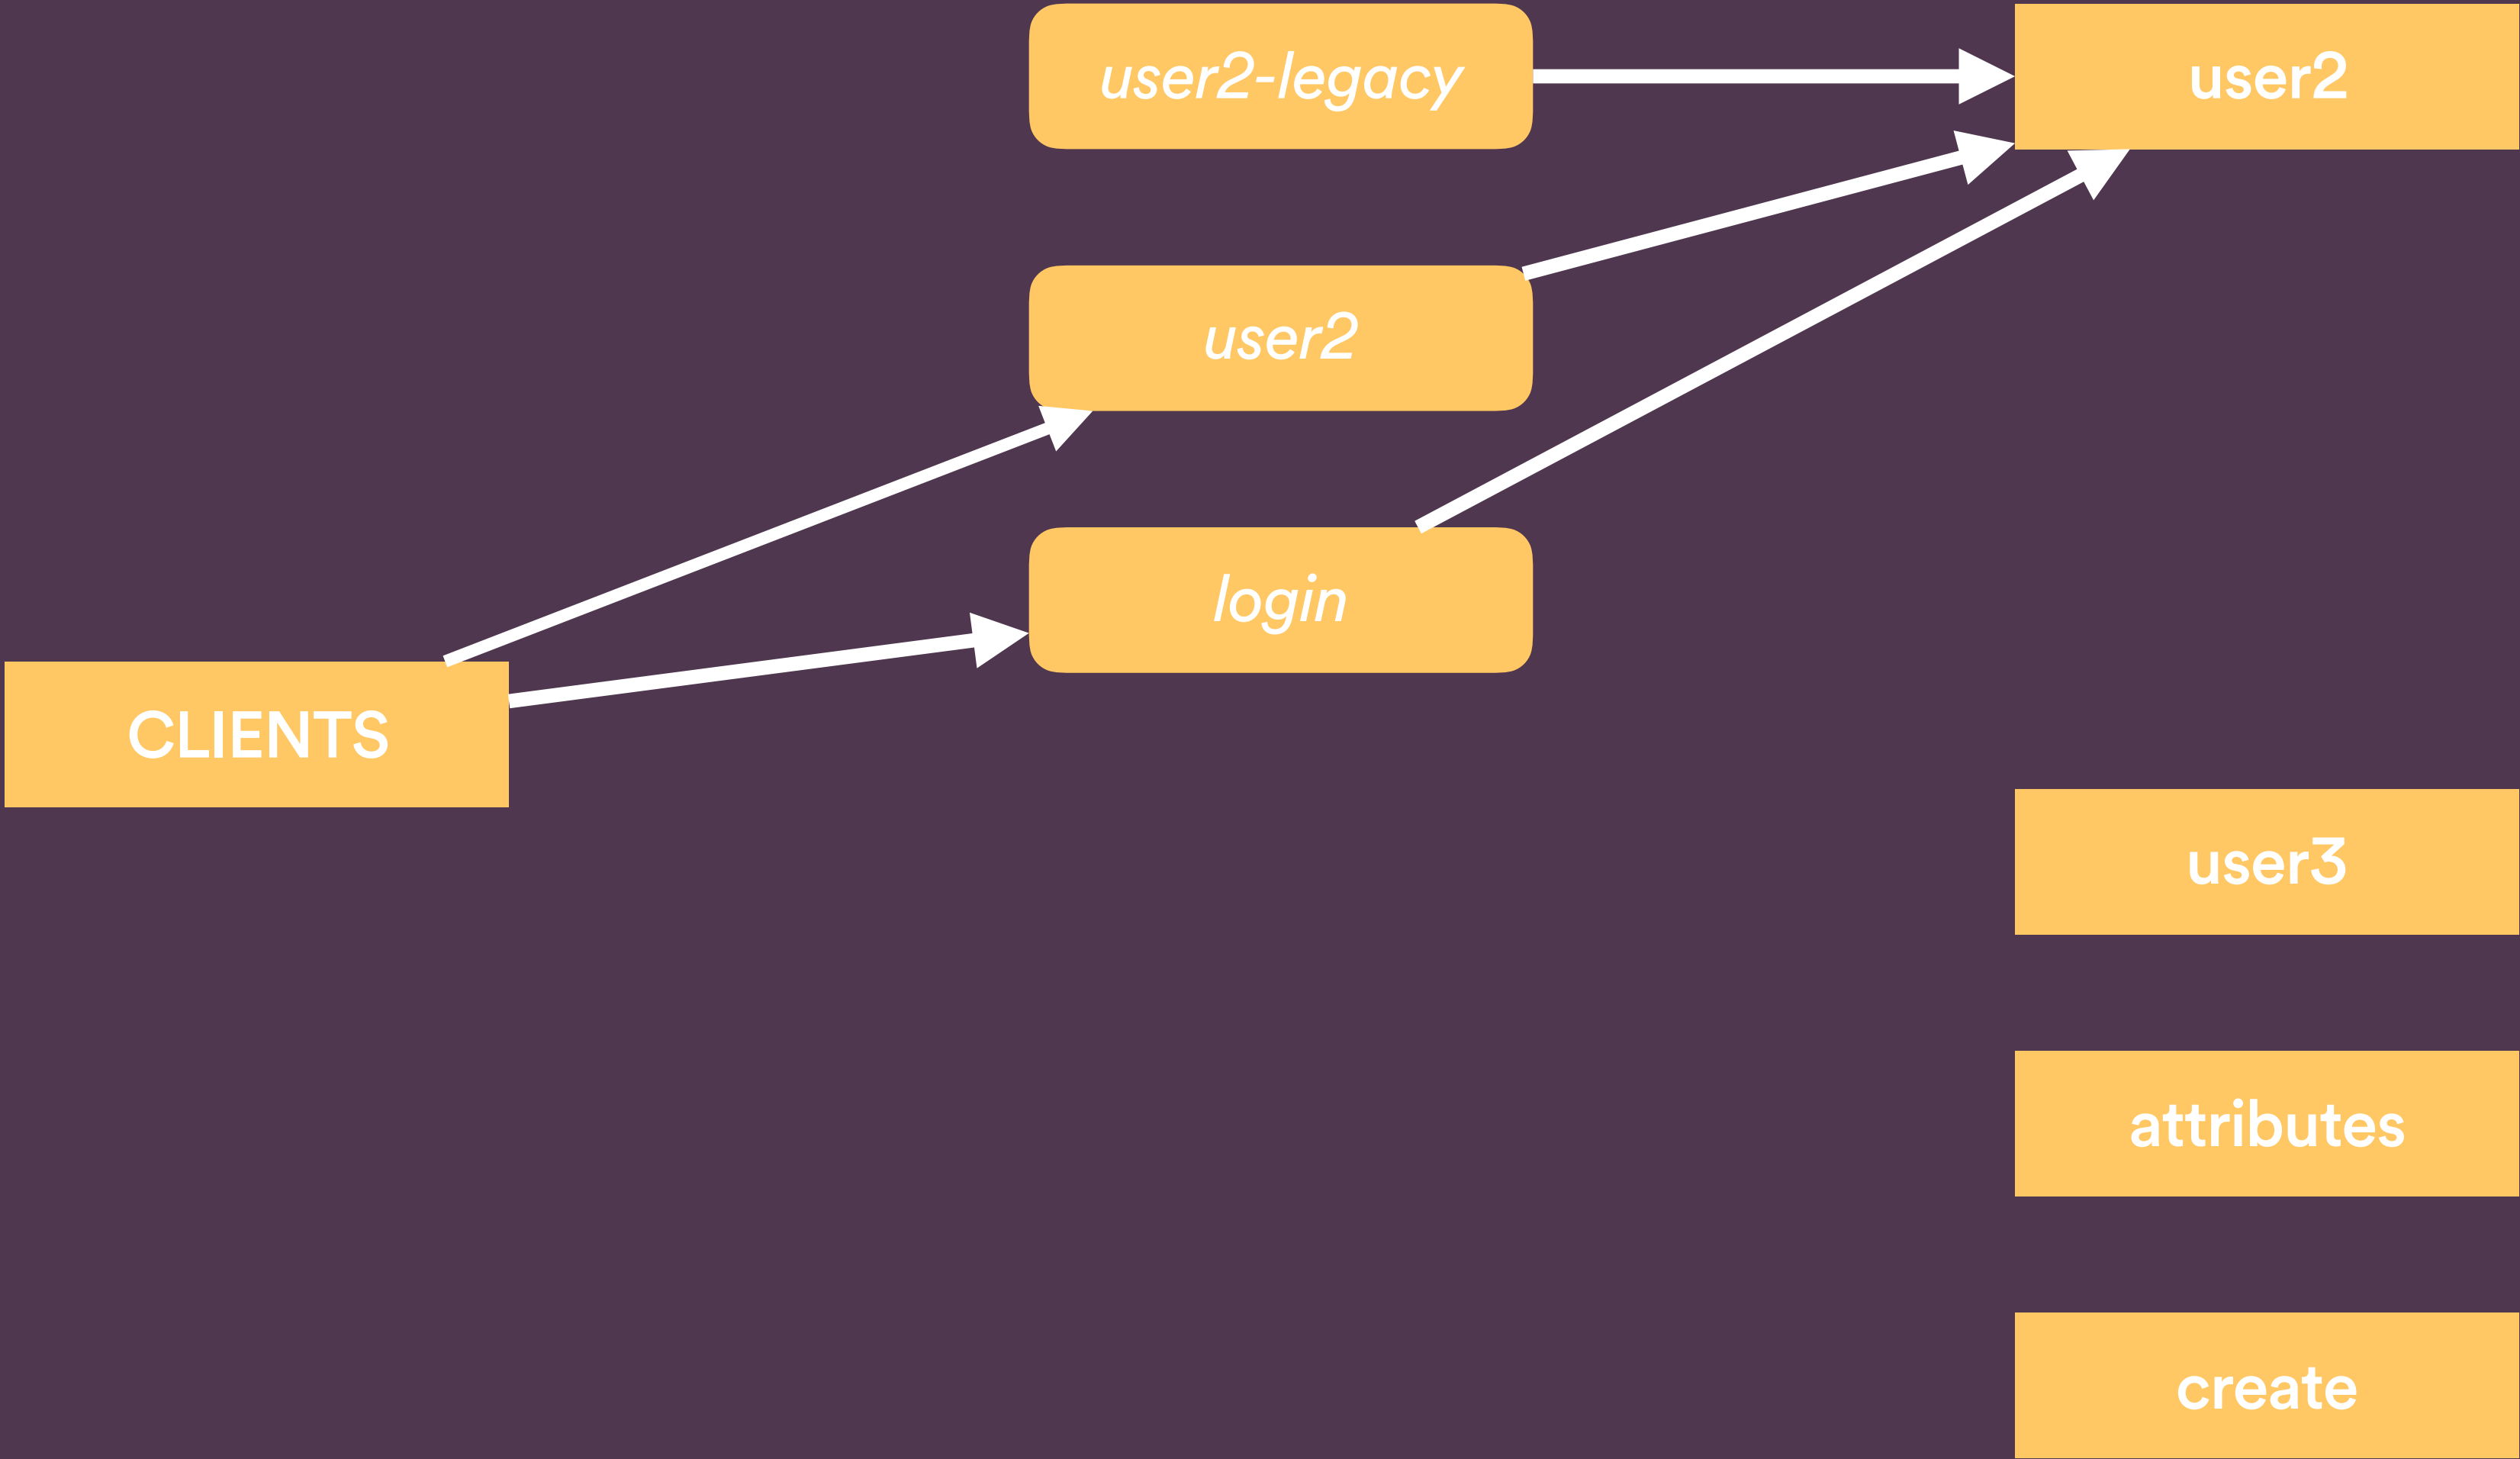
Features add specific data

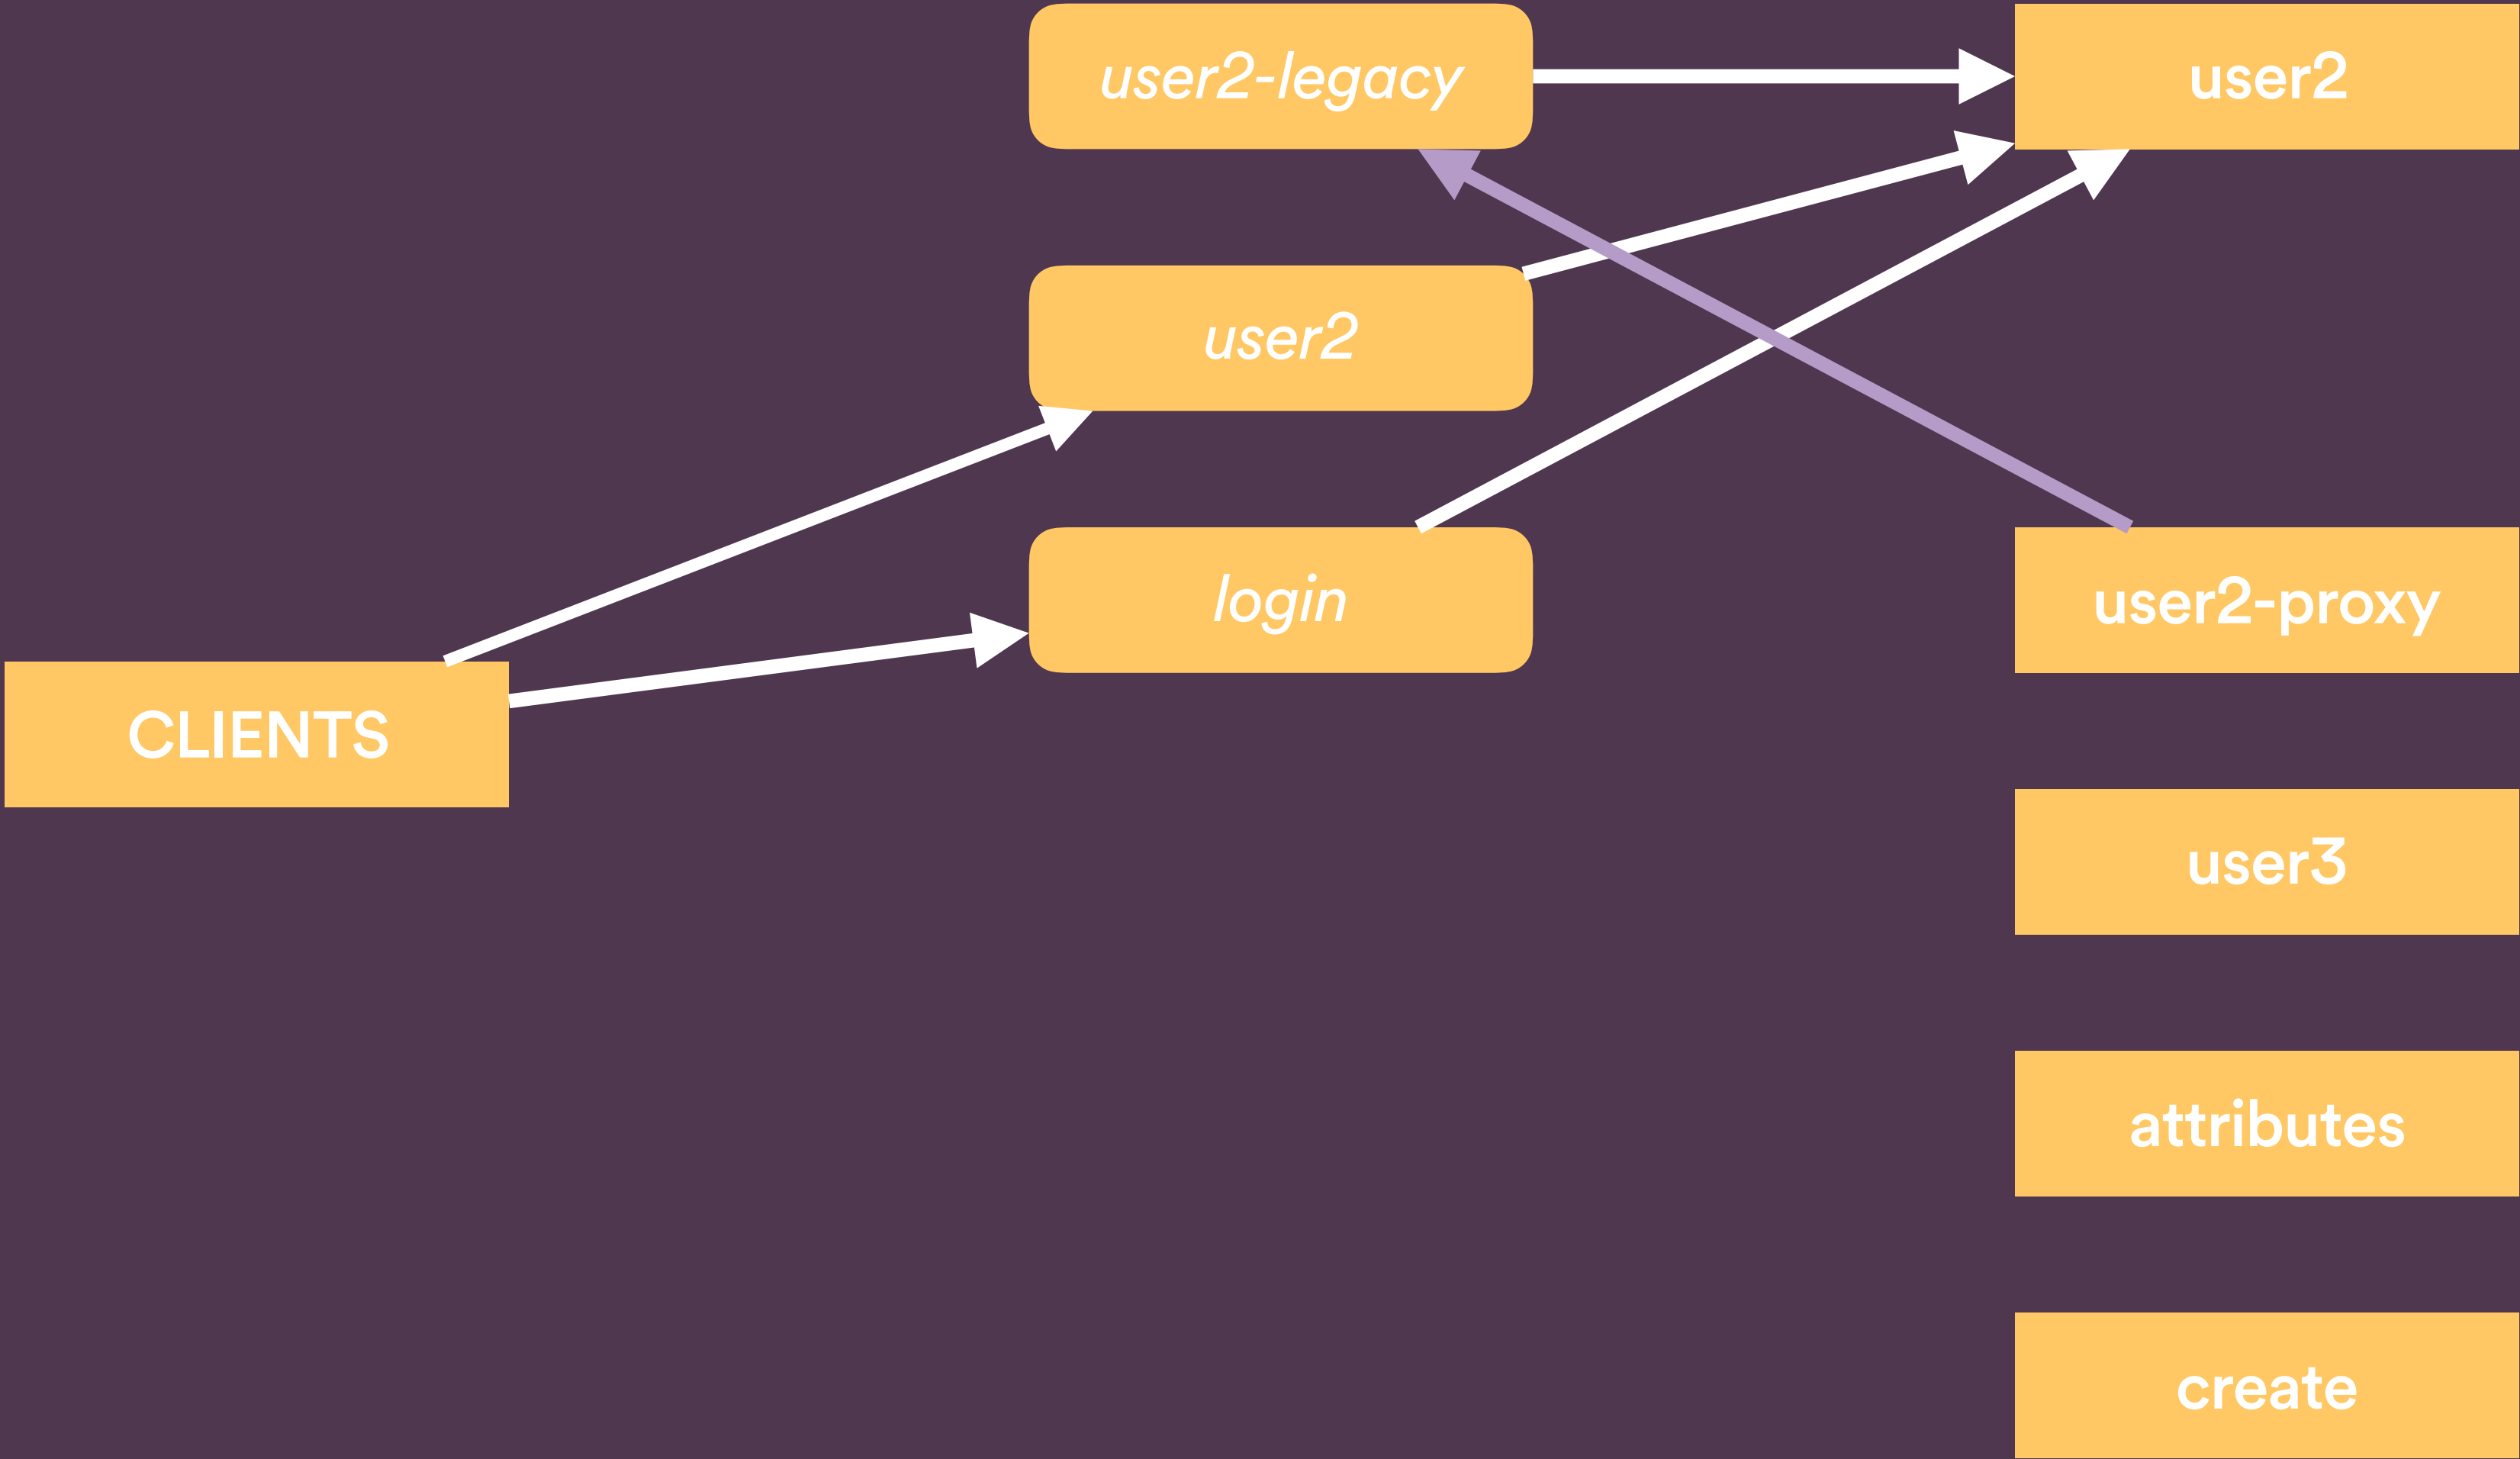
Discovery names as indirection




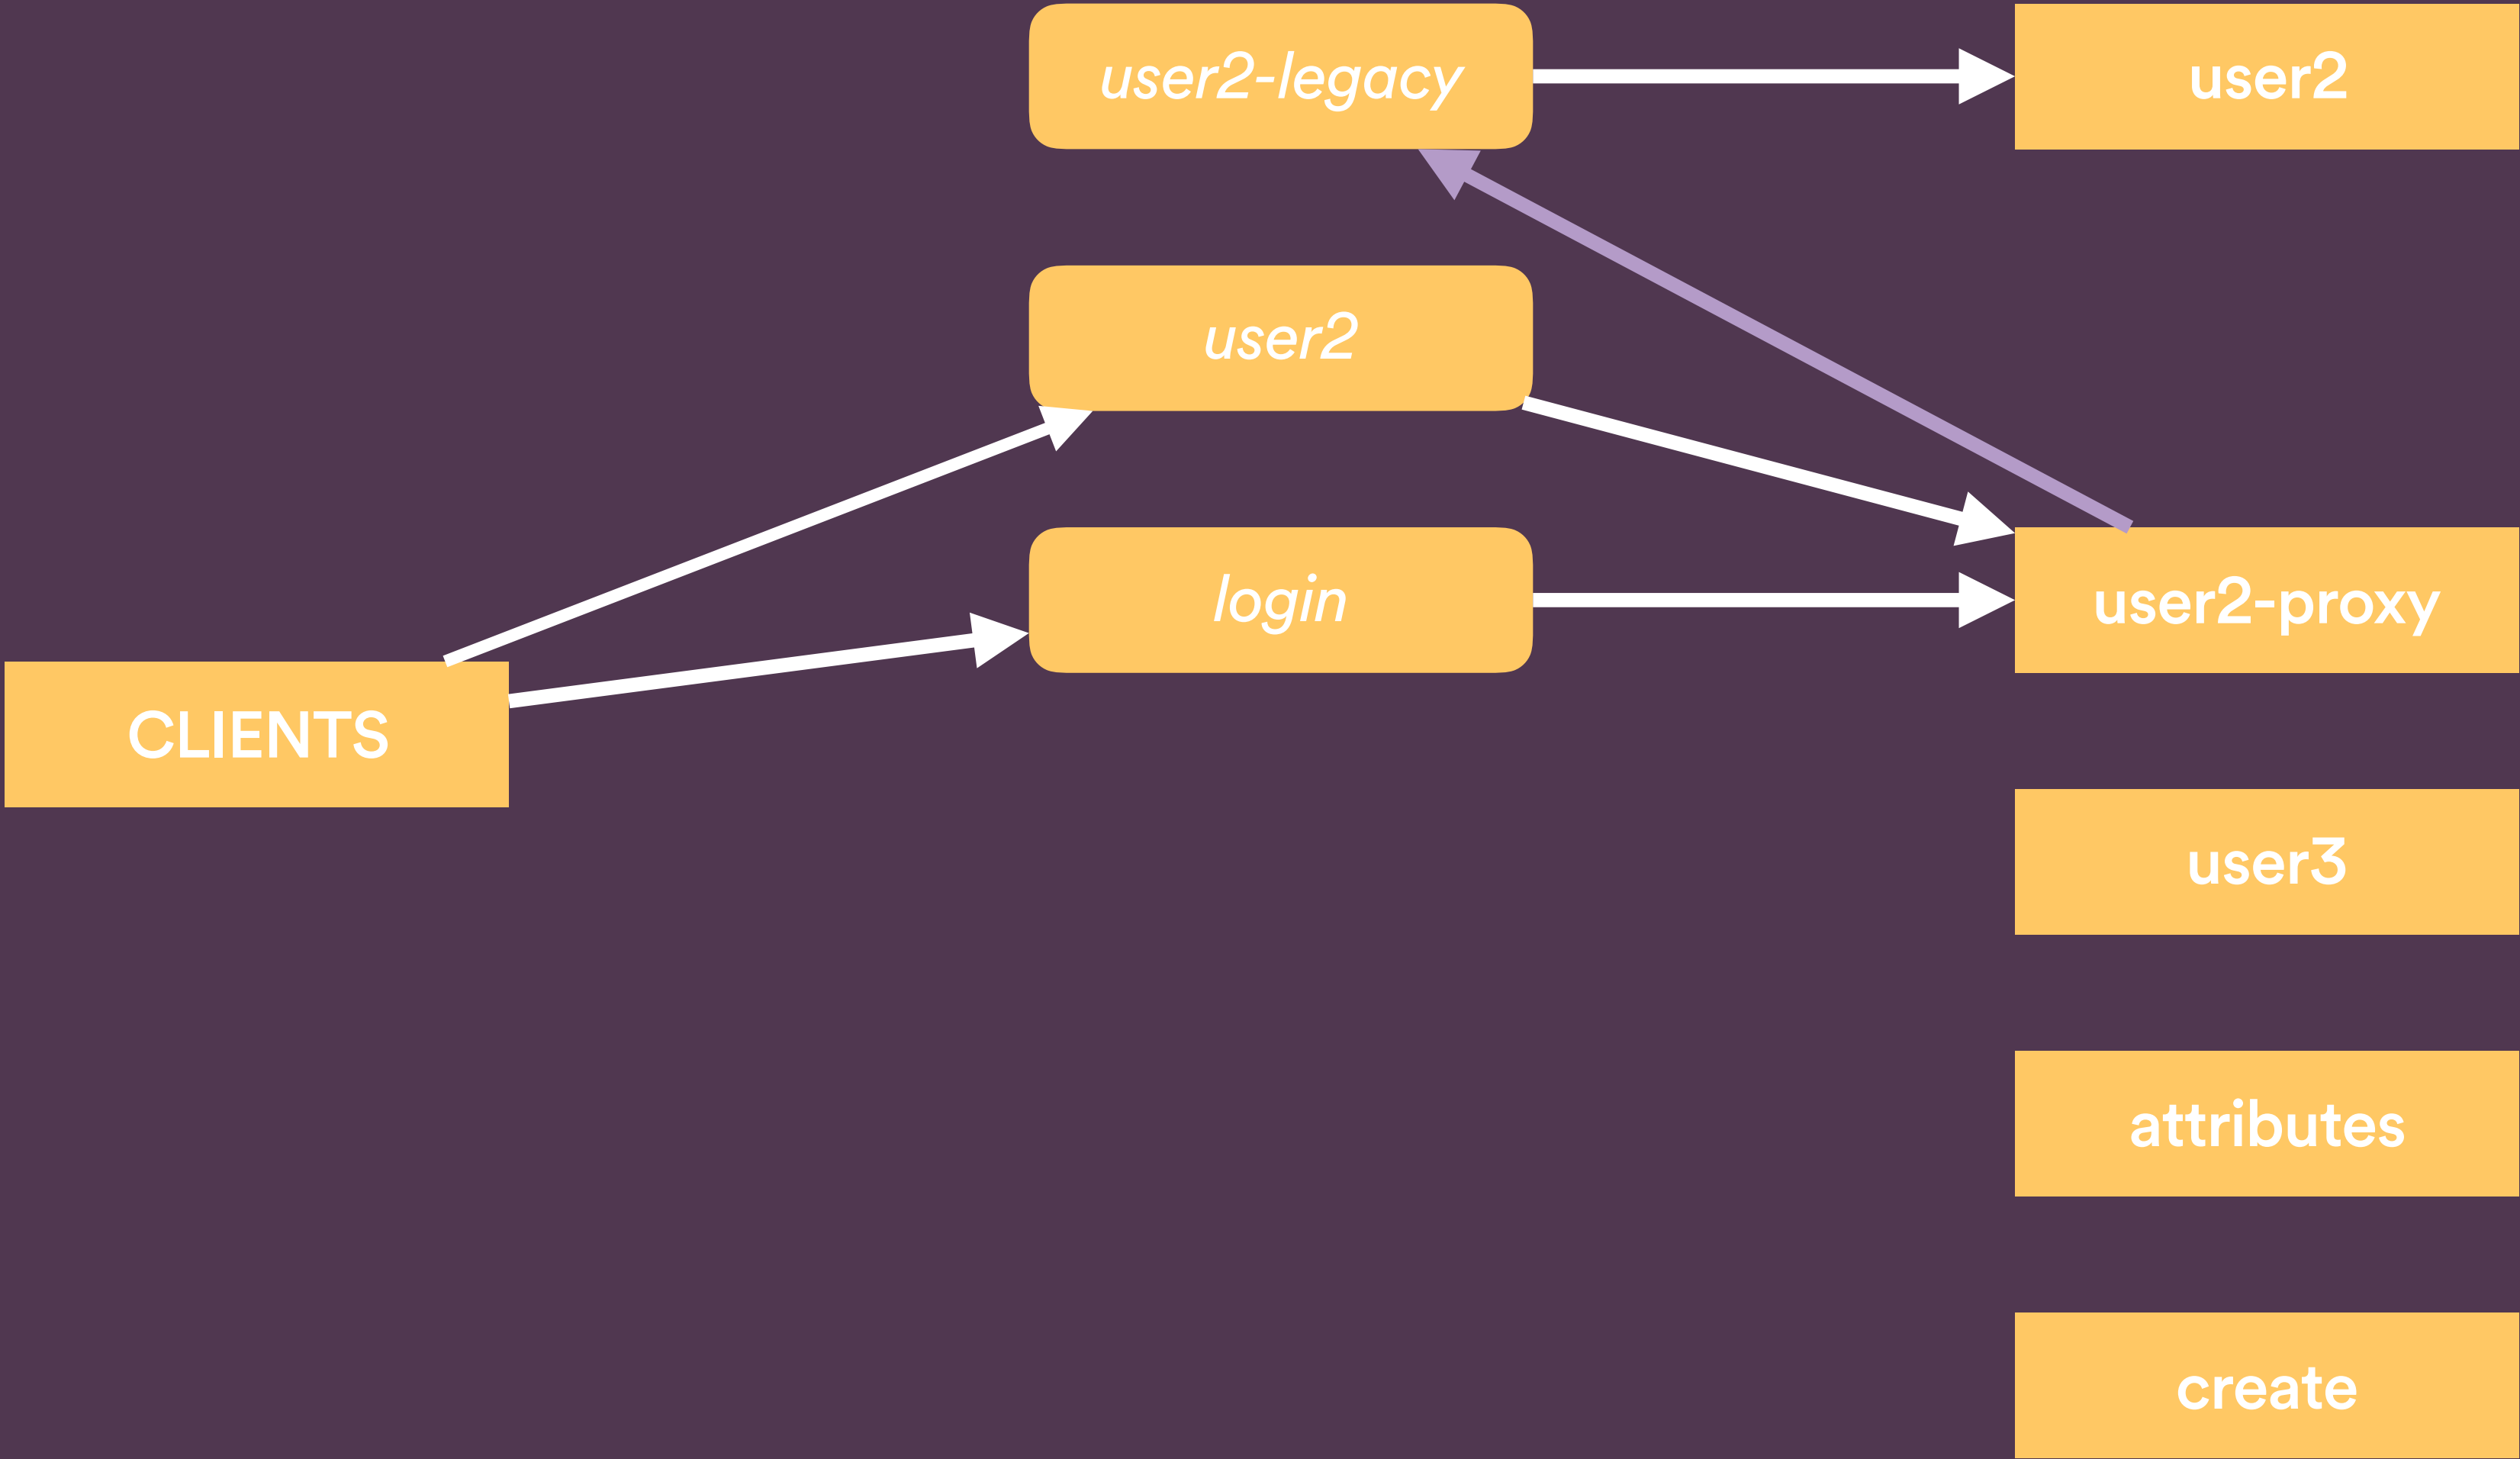


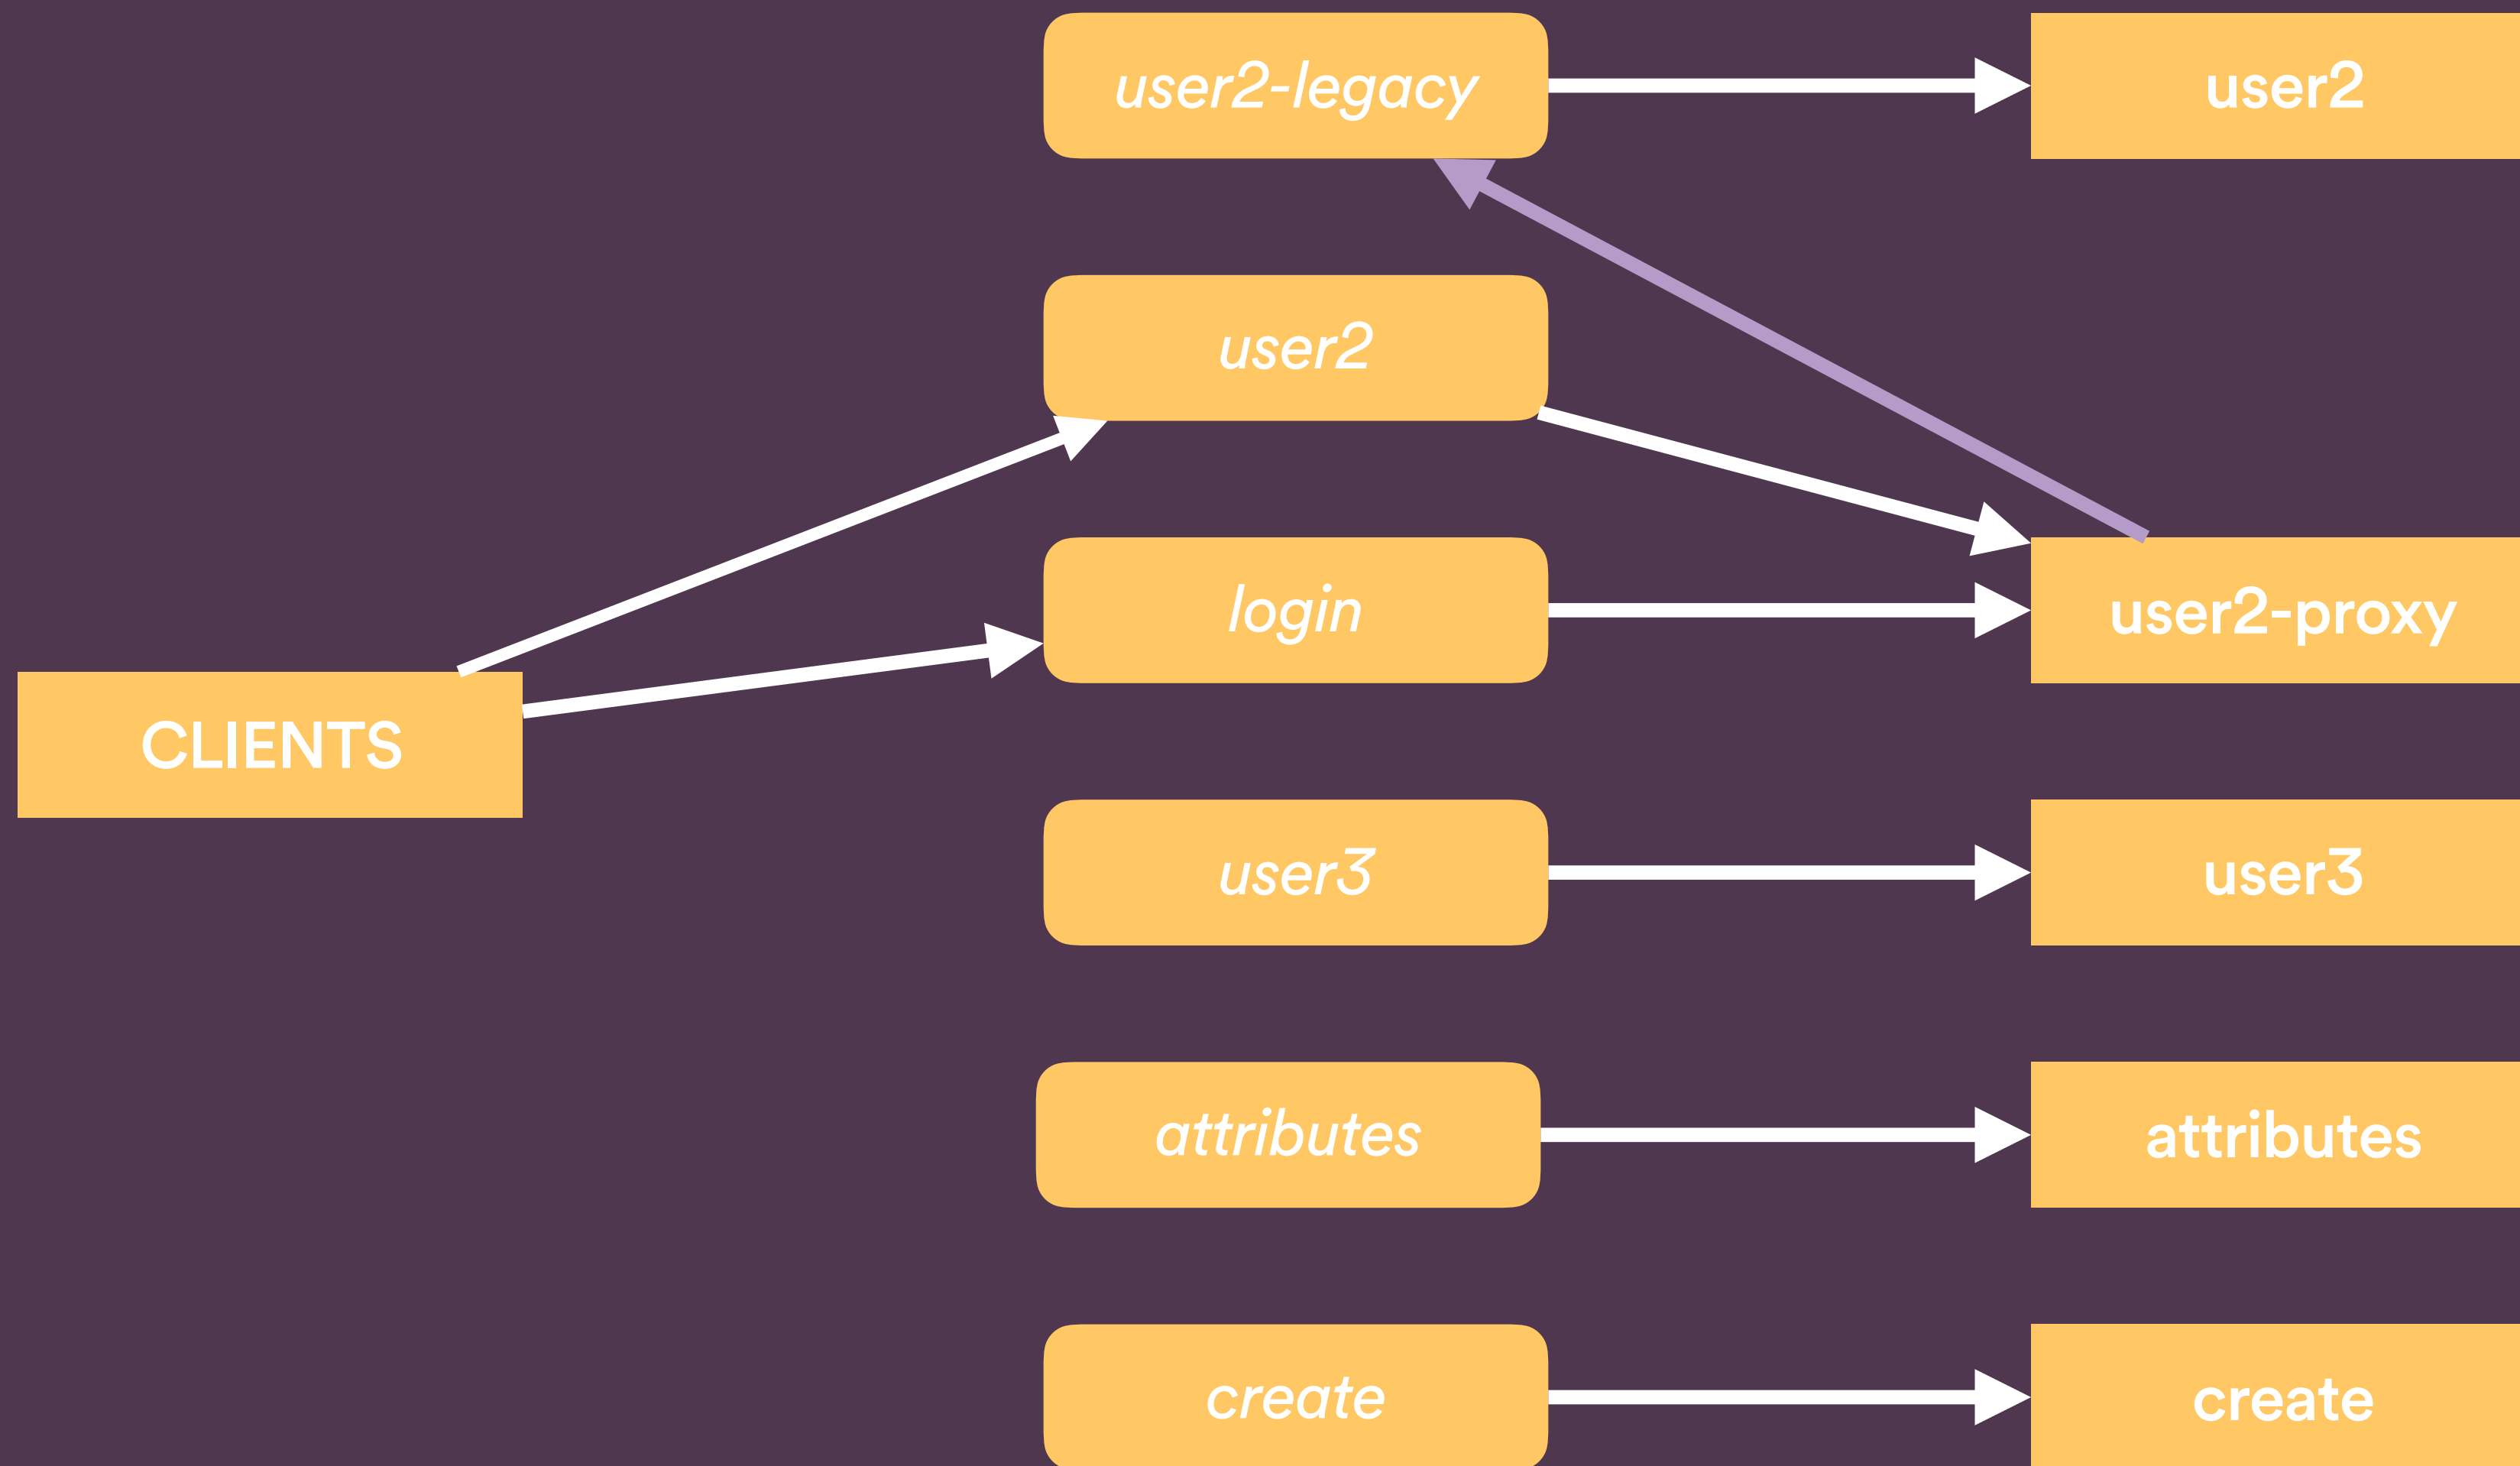


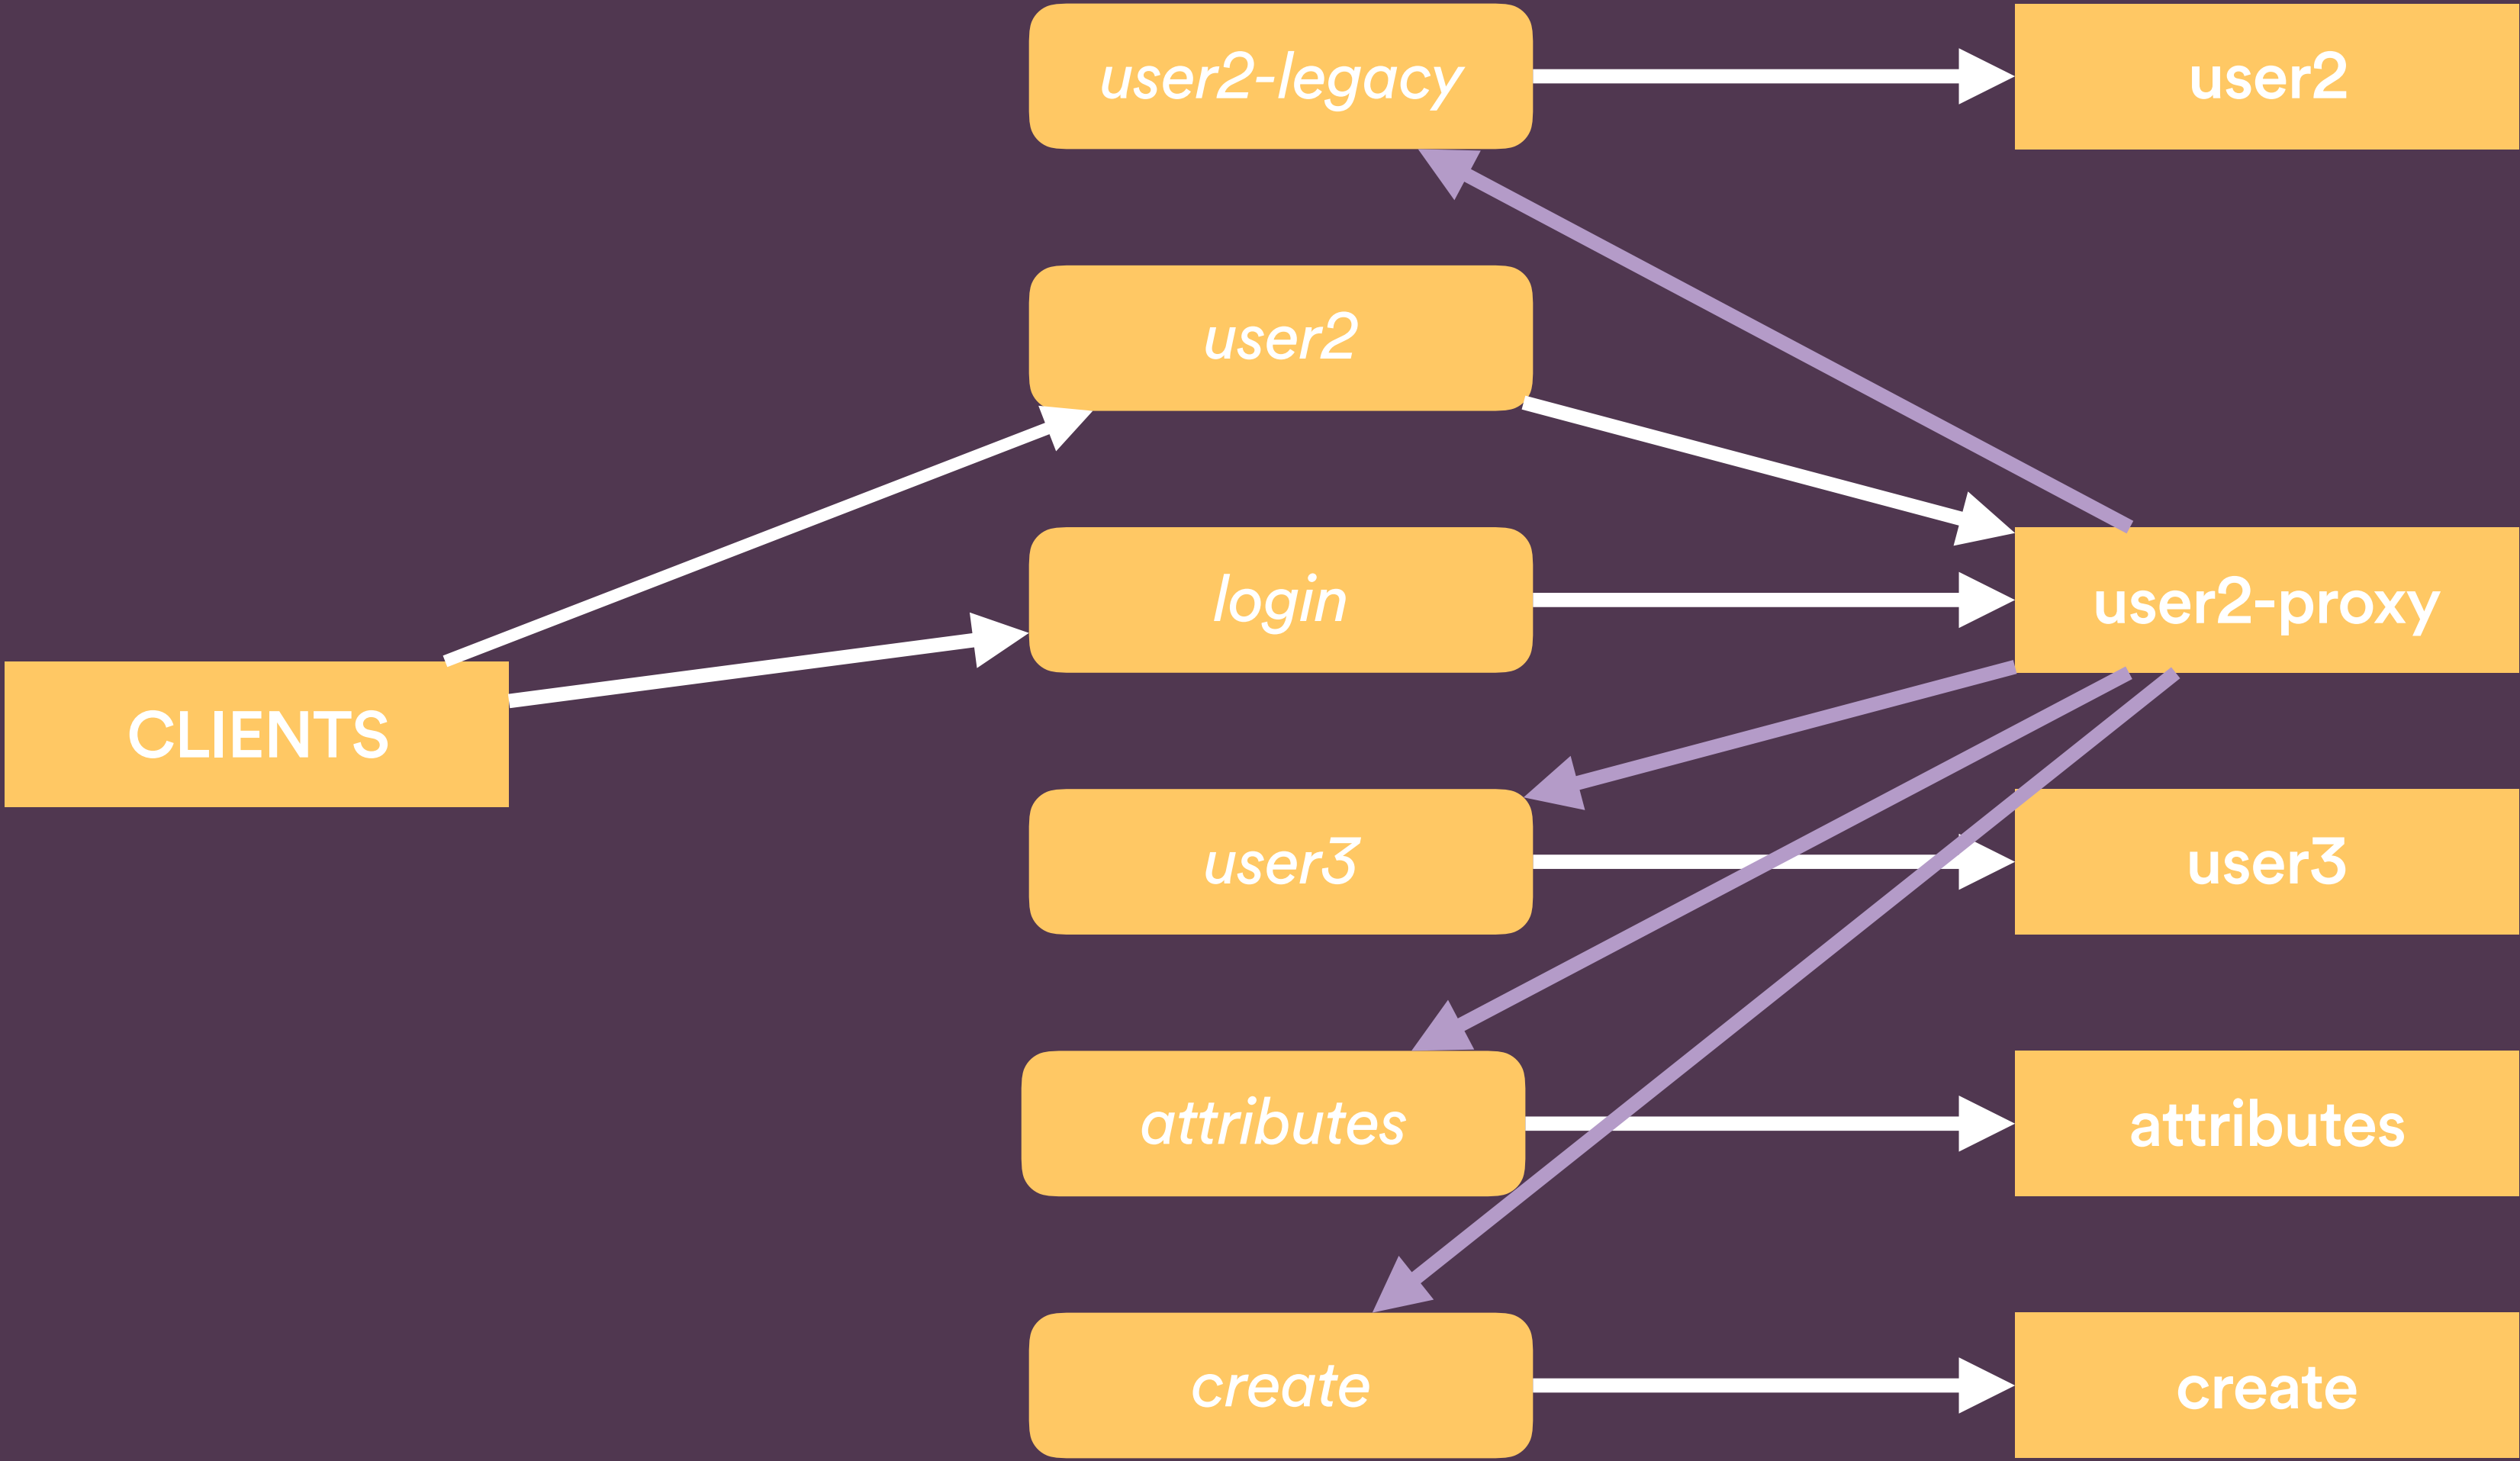


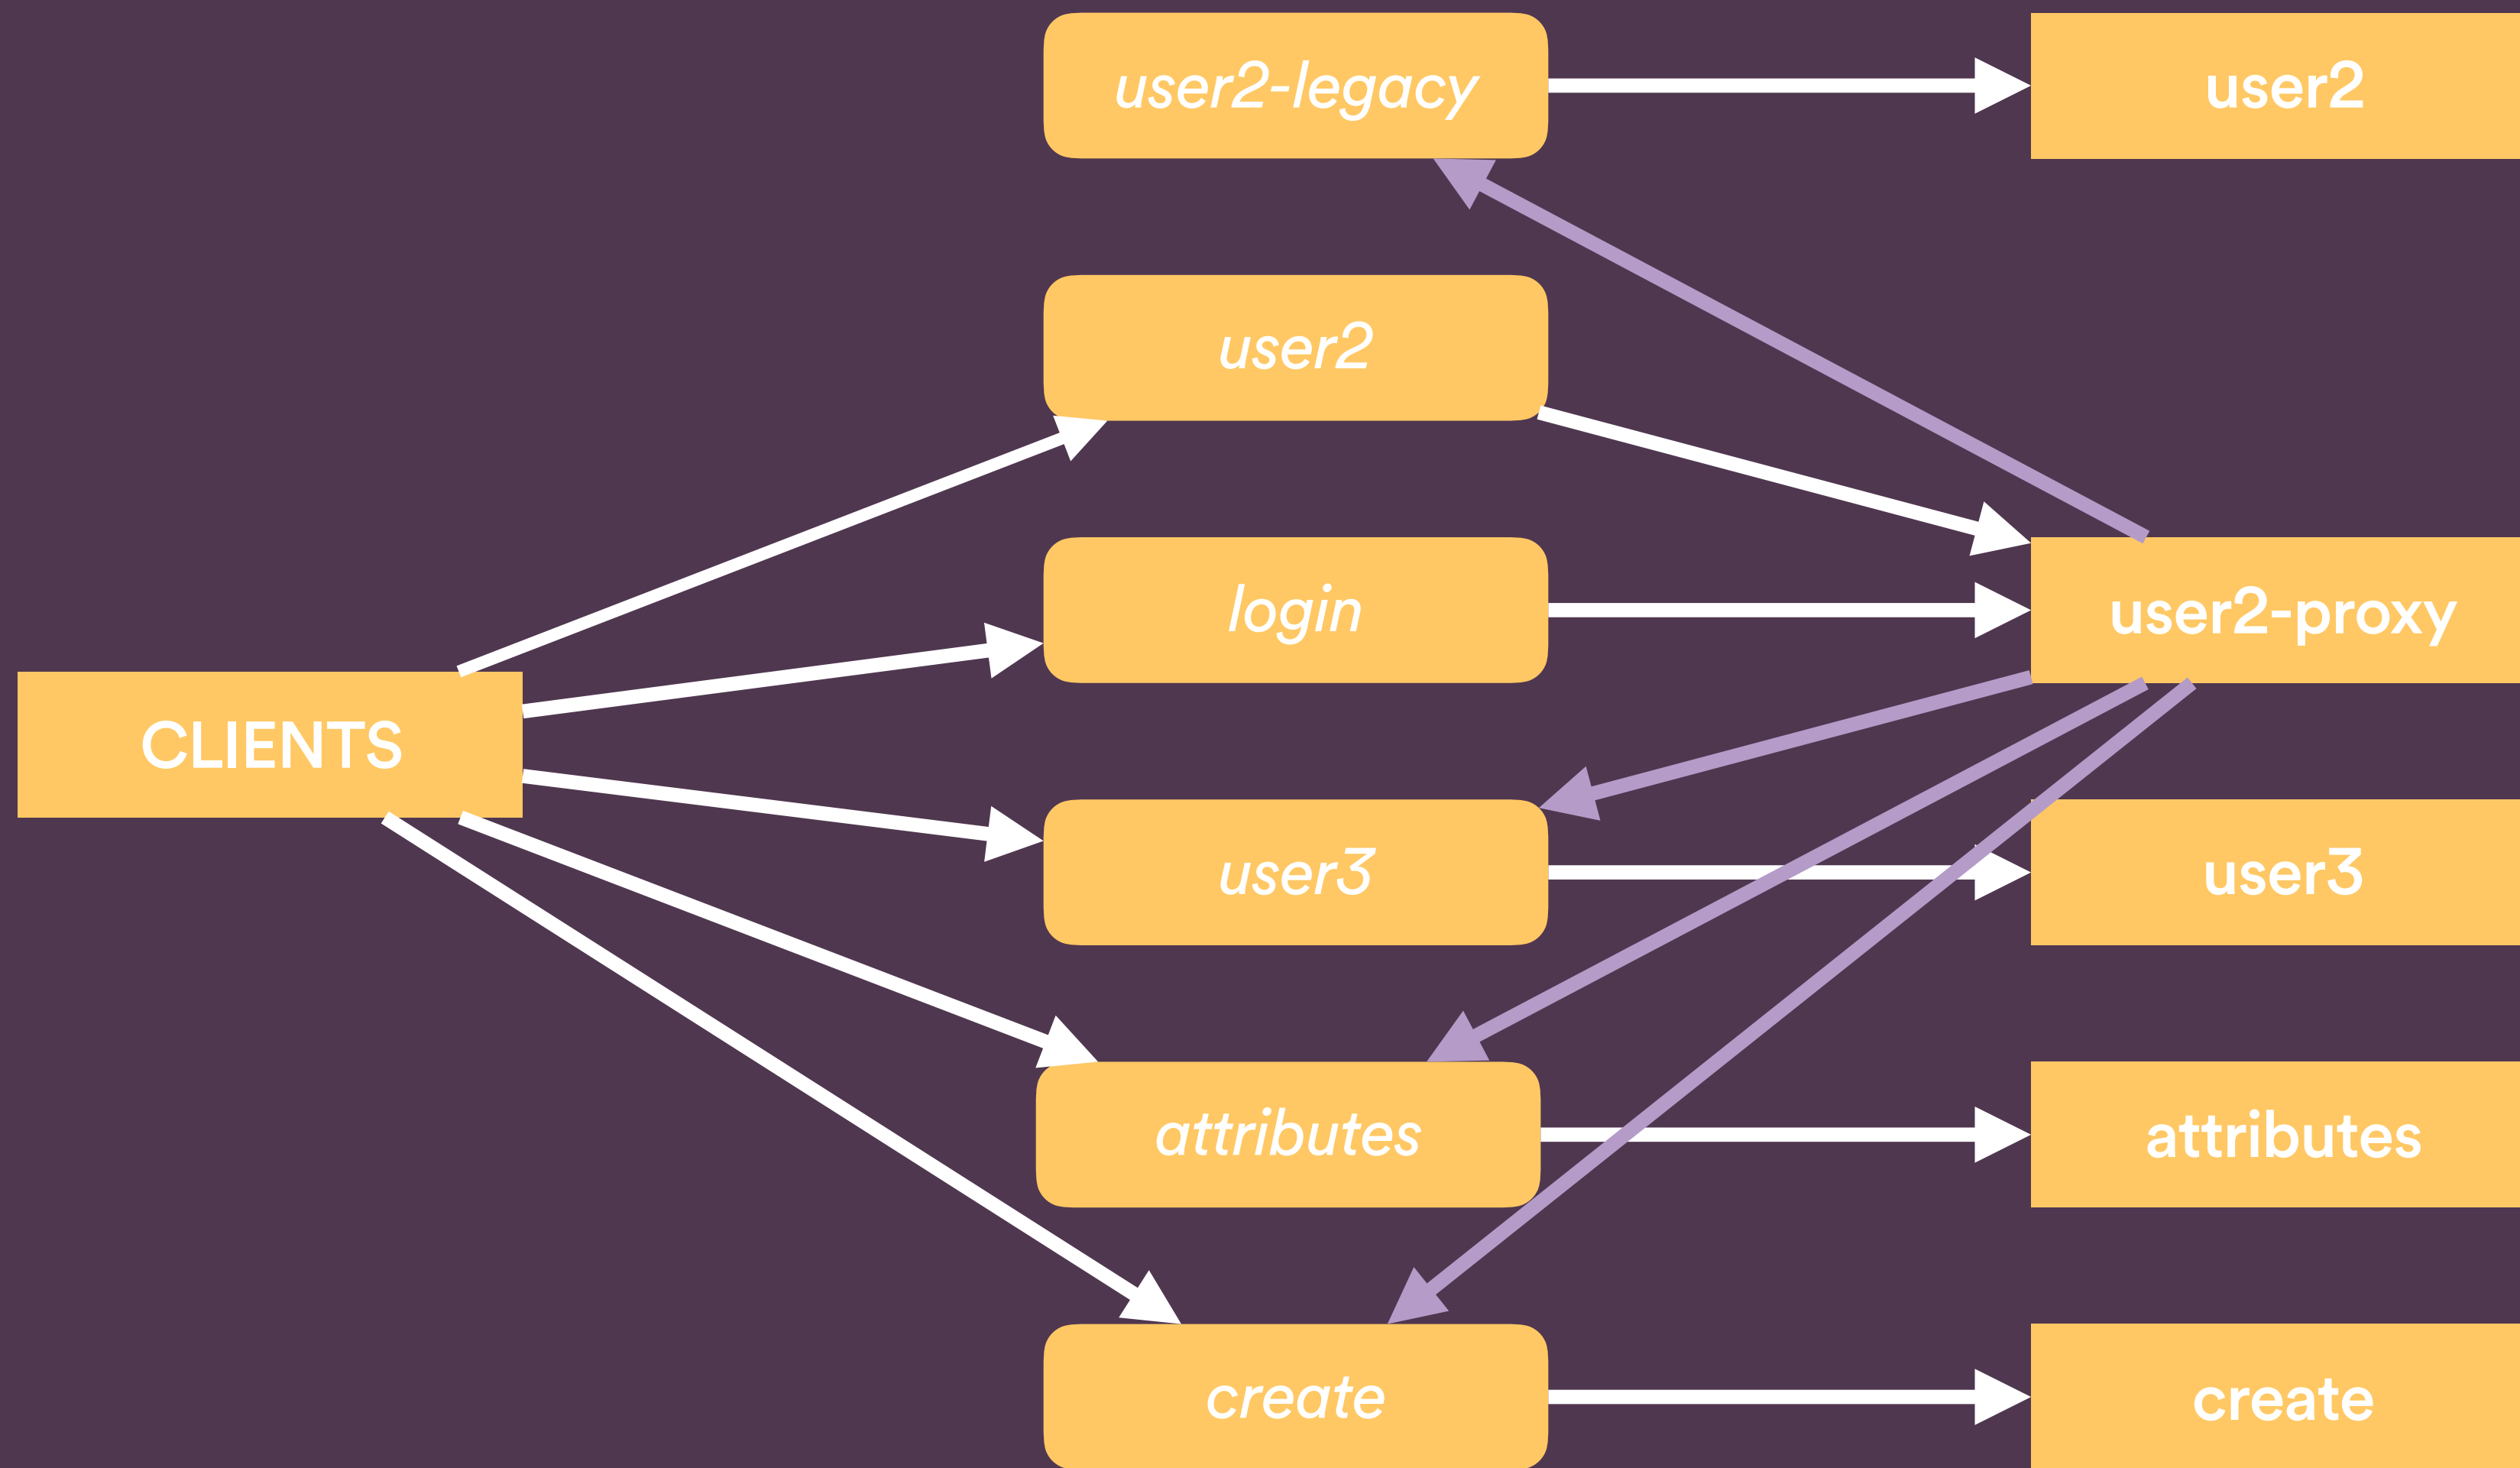


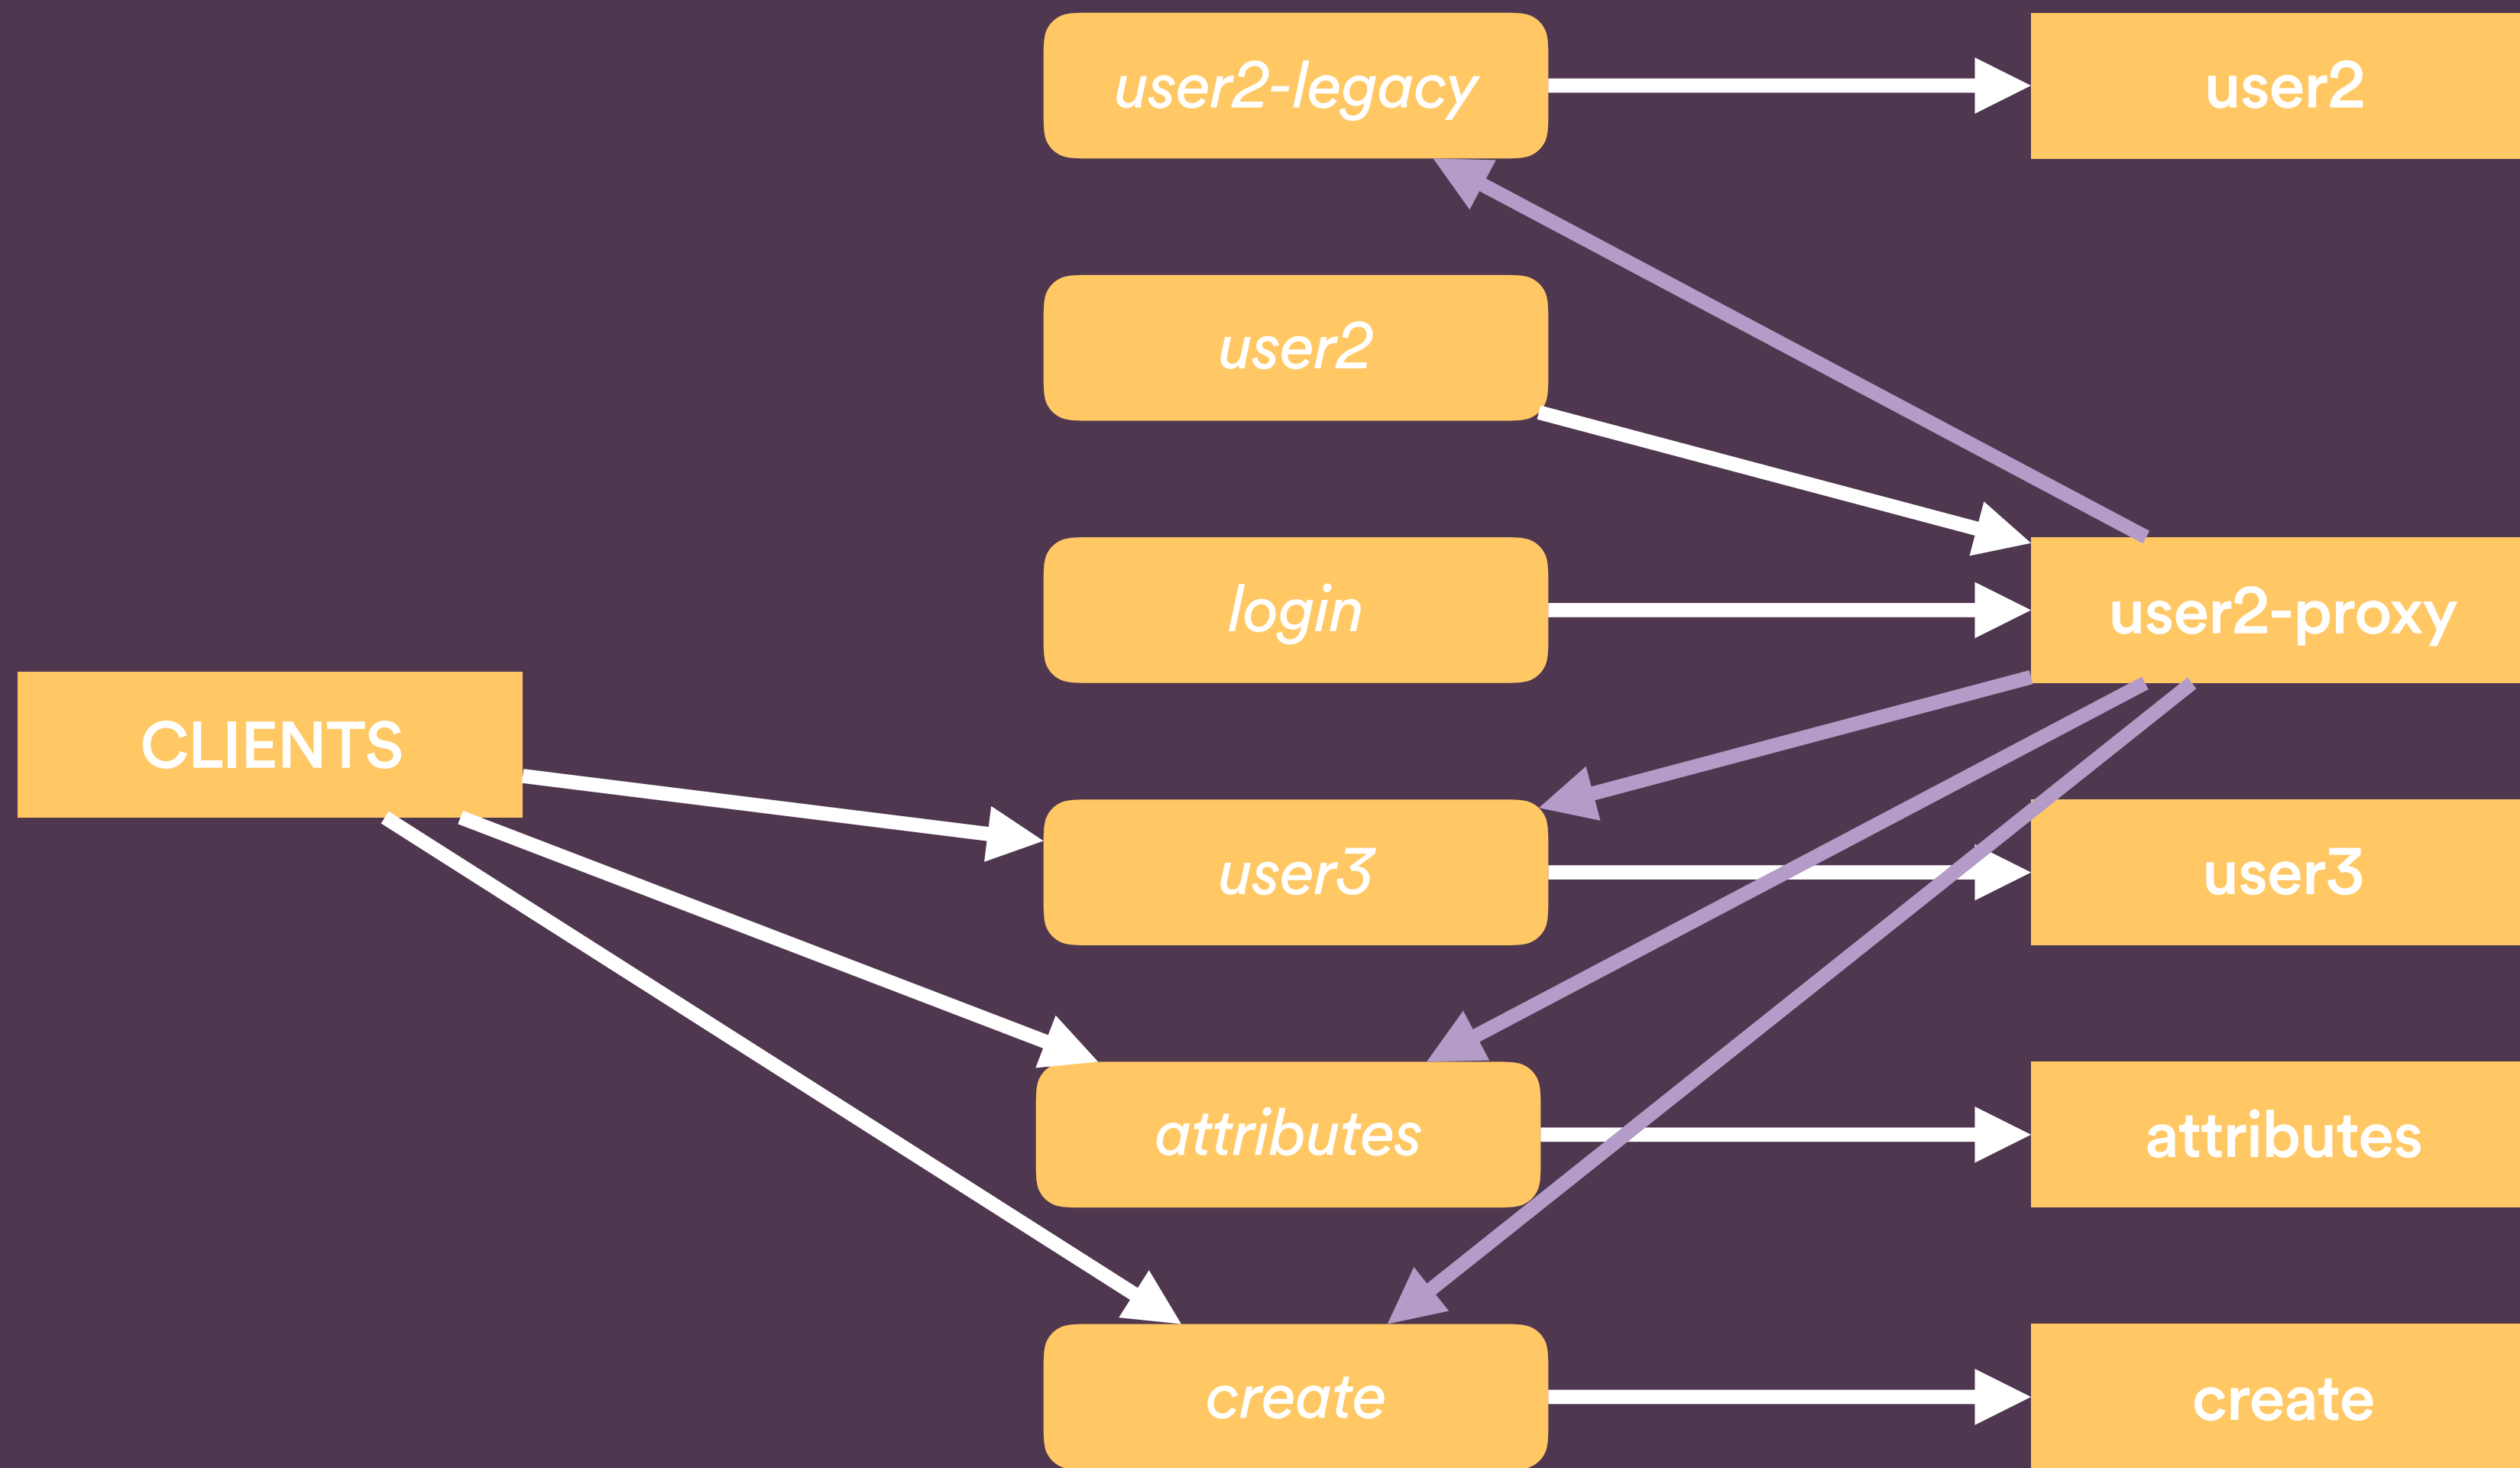


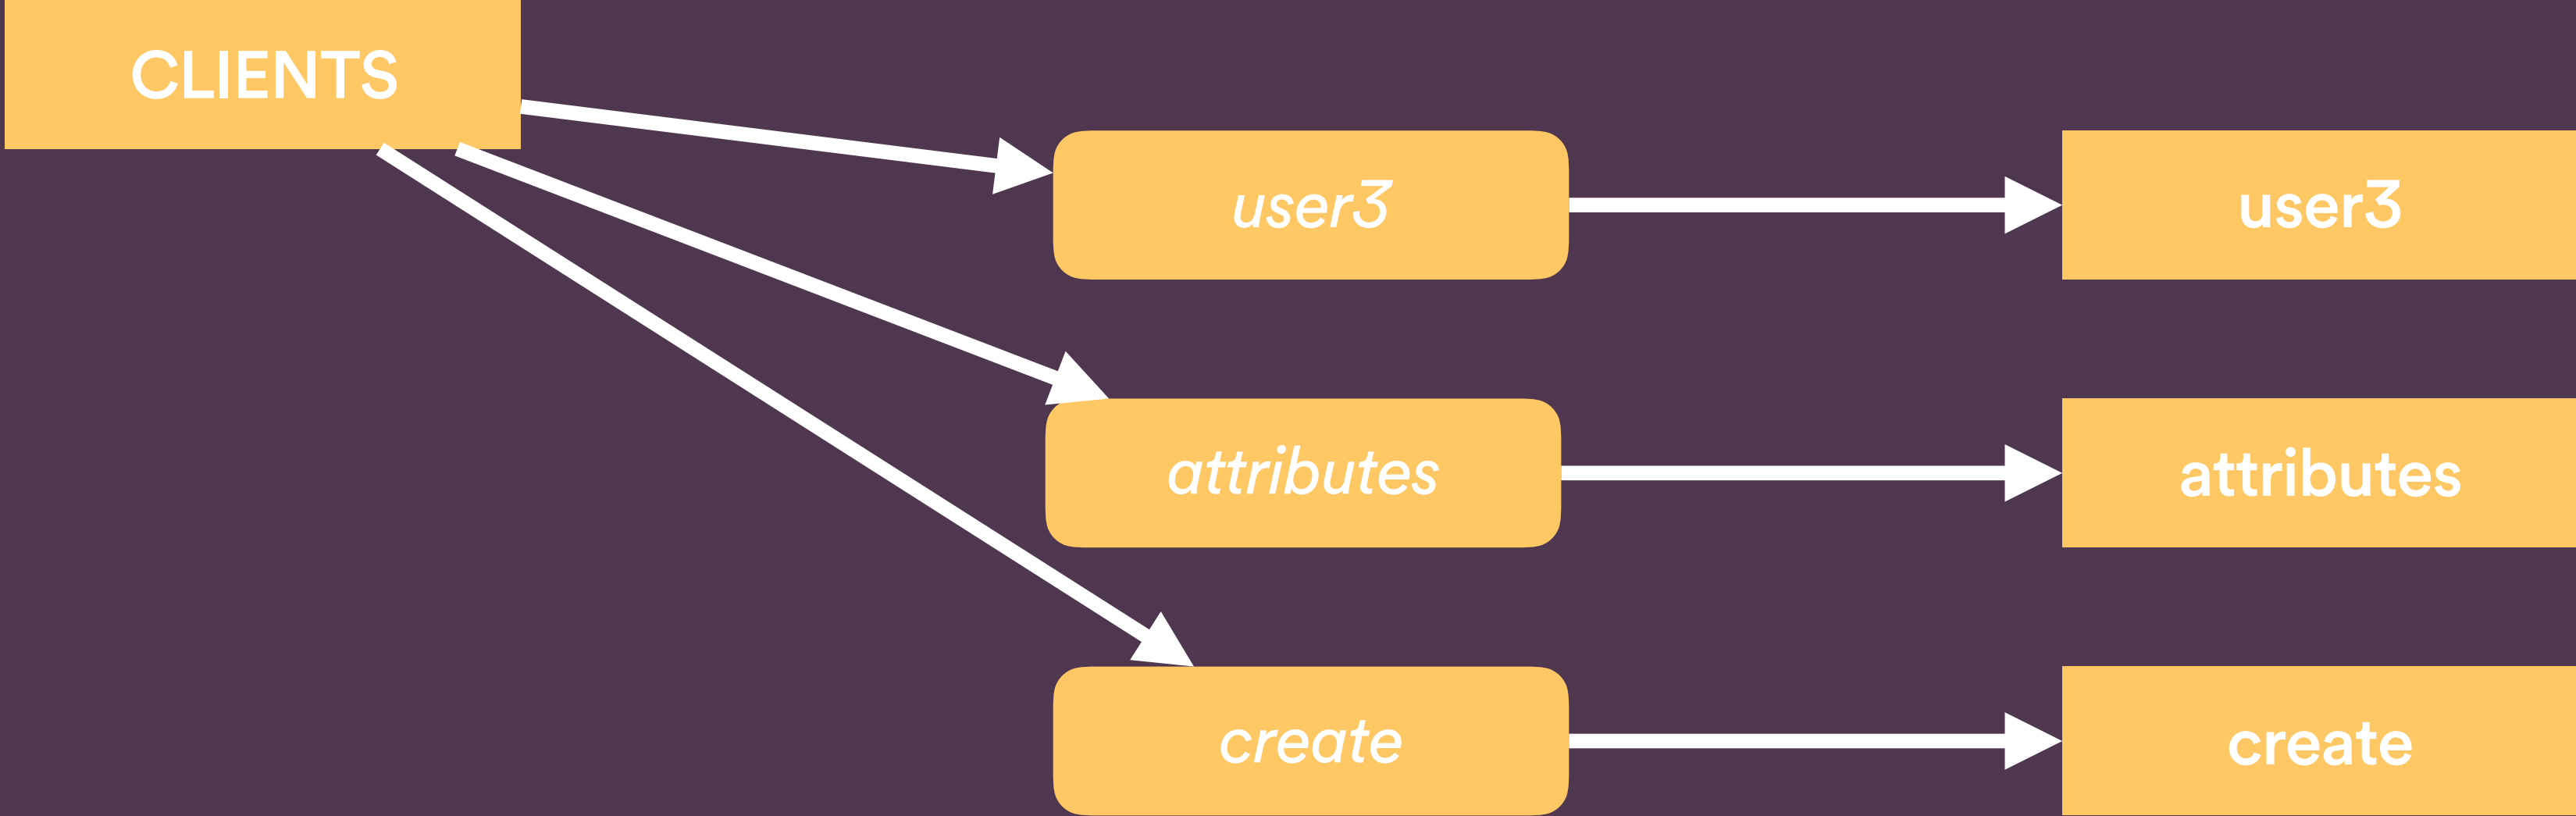




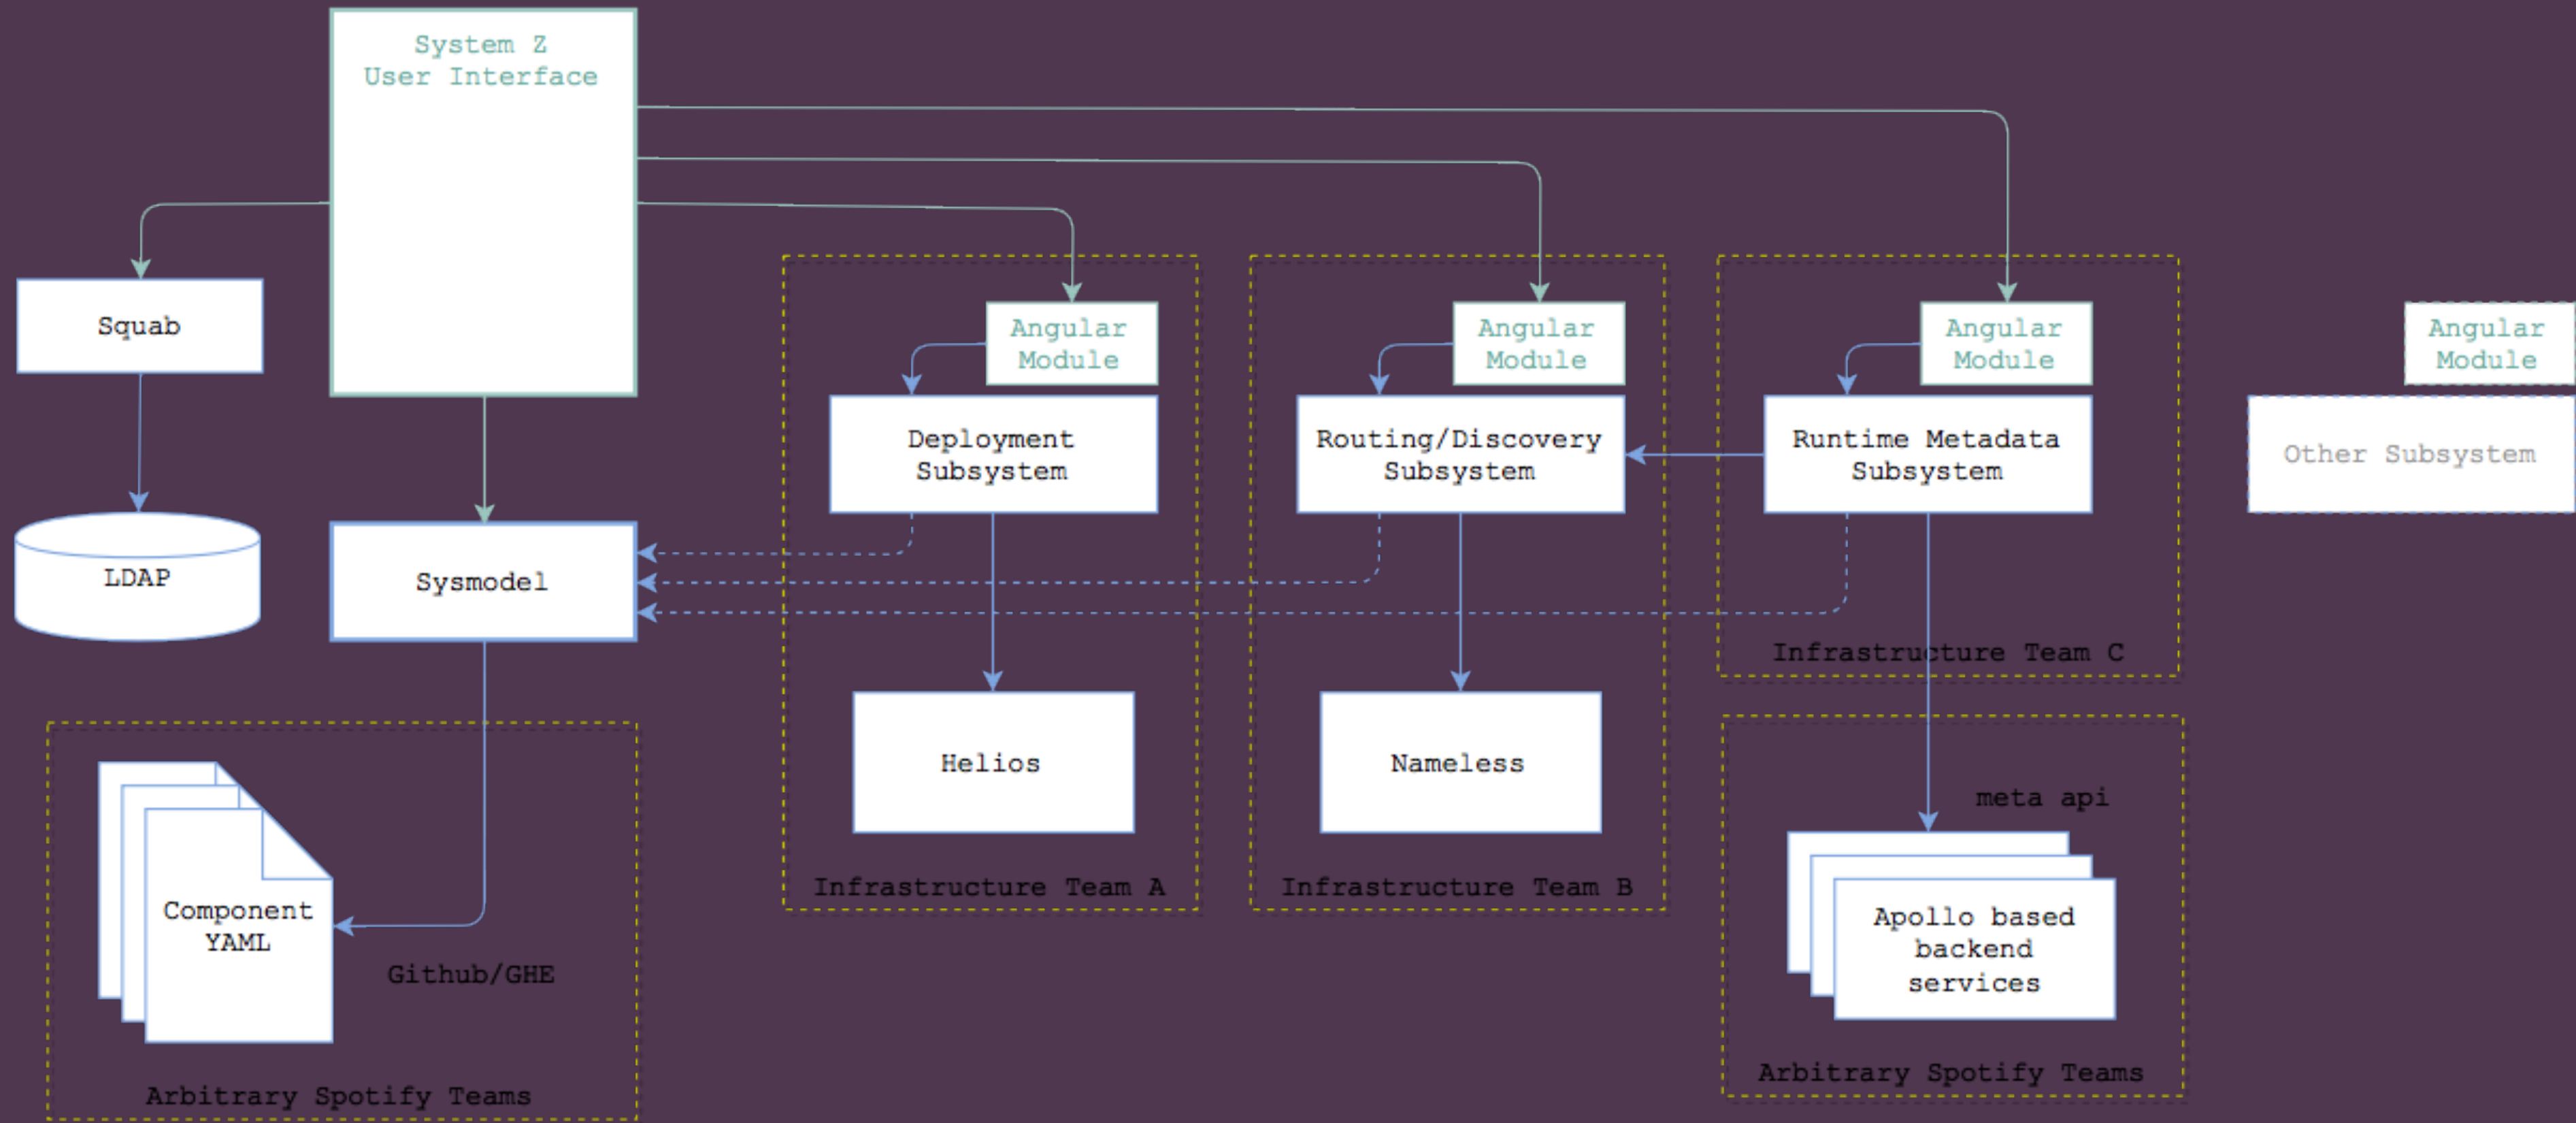









System Z Architecture



Legend



YAML files

Loose/multiple schemas

Files live with code

```
# Basic YAML syntax here: http://www.yaml.org/spec/1.2/spec.html
# Online parser here: http://yaml-online-parser.appspot.com/
# Short version:
#   Strings: single, double or no quotes. Multiline string? Read docs.
#   Lists: with [] or each entry indented on a line with a - .
#   Maps: with {key: val} or each pair indented on a line.
#       (also, x means "x", but no means false, because reasons.)

# Unique ID for this component (in System-Z):
id: gaboeventservice

# Brief description of this component:|
description: Hermes frontend for the Gabo pipeline

# The system this component belongs to:
system: gabo

# Main owner of this component (LDAP group or individual):
owner: 202-squad

# Intended usage of this component:
#   'private' - only for use by the owner
#   'public' - exposes an api for other squads (like metadata)
# Components that are under development, experimental, not
# supported, replaced, liable to change without warning, etc.
# should generally be marked as private.
visibility: public

# A list of other components (services, data stores, etc.) on which
# this component depends. Must be valid service_discovery names
# registered in the Nameless service.
dependencies: []

facts:

# A list of the names under which this component can be discovered
# on production machines. Must be valid service_discovery names
# registered in the Nameless service.
service_discovery:
| - gabo-event-service

# Links to the latest definitive documentation for this component.
# Markdown is supported, e.g. docs: "[Title](URL)", but you must
# use quotes, or it will fail to parse (as a list).
docs:
| - gabo-event-service
```


Dirty Data

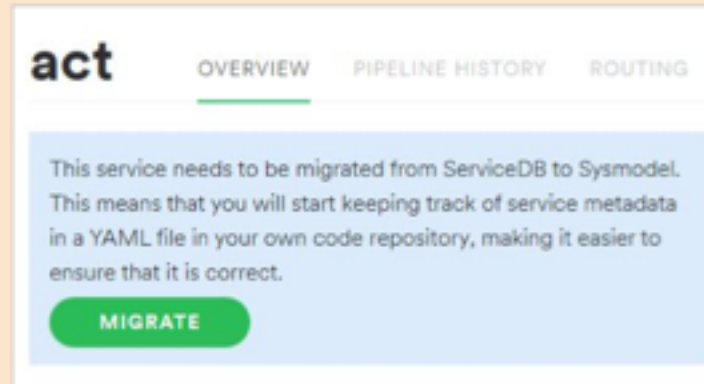
Organisational change => ownership confusion

Infrastructure evolution => runtime confusion

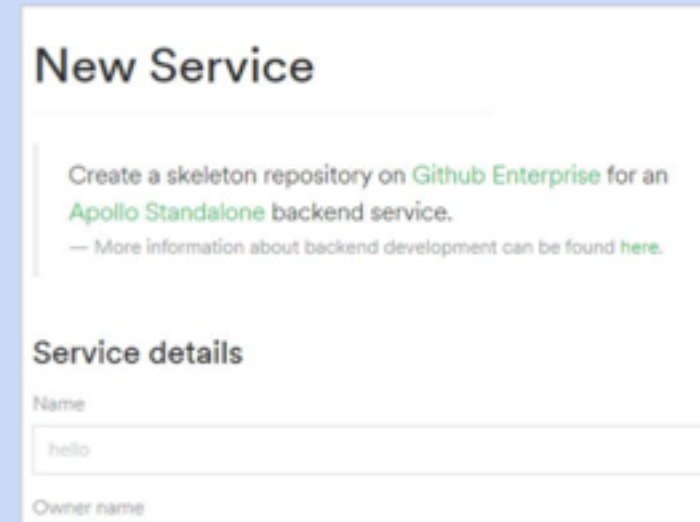
Owners don't benefit from metadata quality

Ten Things You Didn't Know About **System-Z** That Will Amaze You!

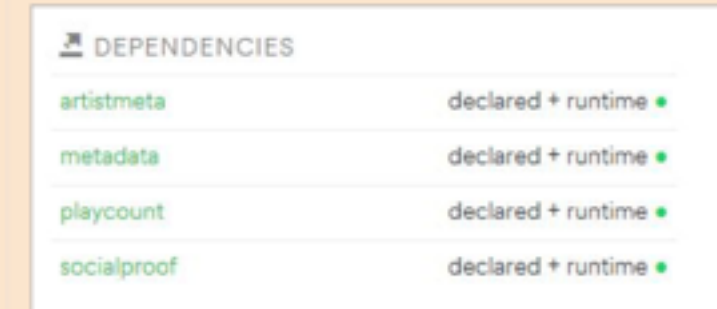
1. Migrating your service from the old "servicedb" takes less than two minutes and unlocks many of System-Z's useful features!



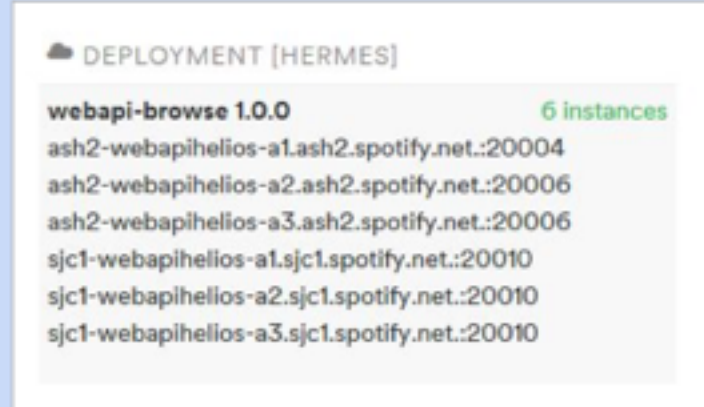
2. System-Z's **New Service** feature helps you create a skeleton Apollo backend service in seconds! You can get your build pipeline going on any existing Jenkins instance!



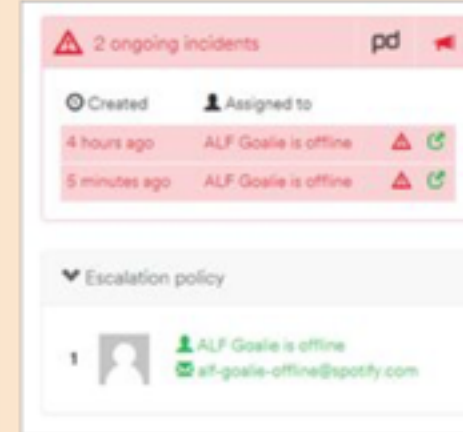
3. You can see an Apollo-Hermes service's dependencies (both declared and detected in runtime) and drill down to explore service relationships.



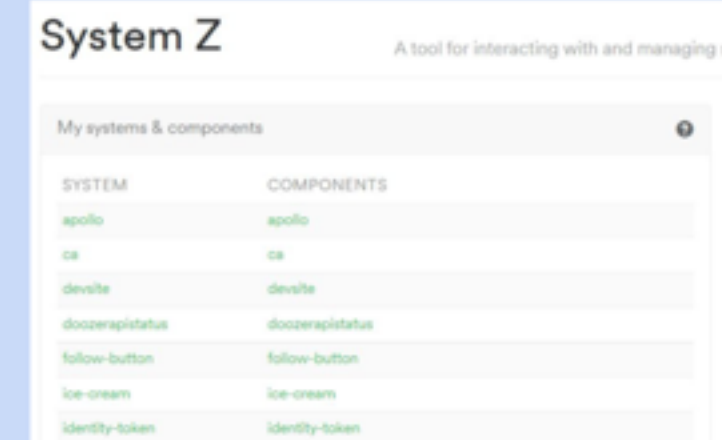
4. You can quickly see exactly where any service using Hermes is currently deployed!



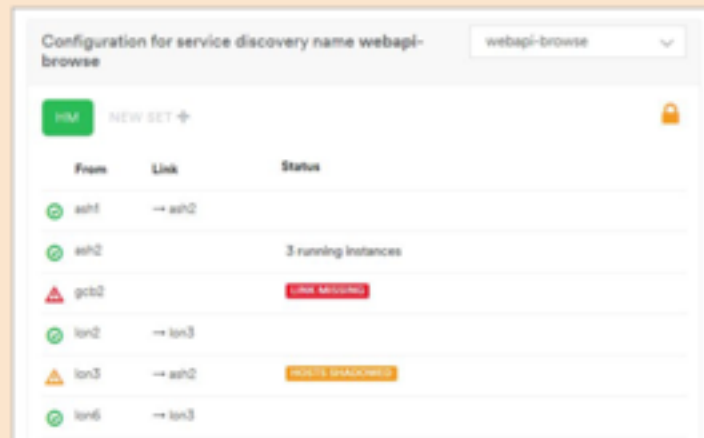
5. For any service that has registered its PagerDuty key, you can see all current incidents and escalation policies!



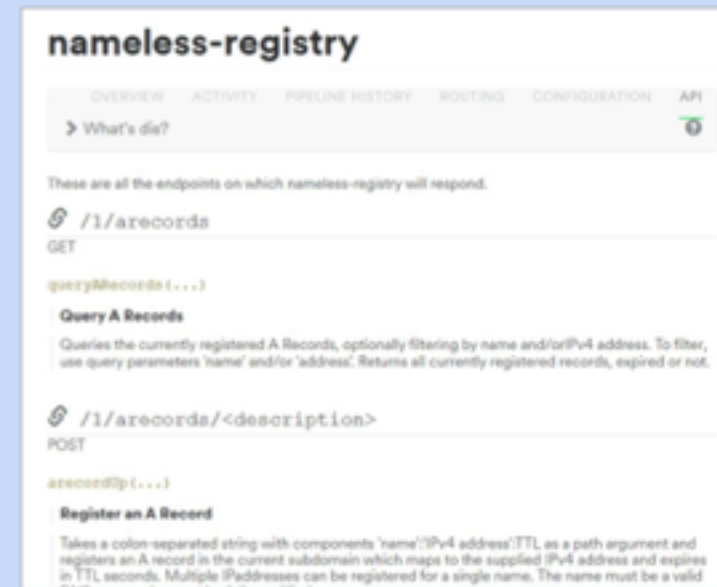
6. System-Z's homepage shows you all the system components owned by you and your teams.



7. You can see and manage routing between pods in our infrastructure!



8. You can see the published API for a service and all the endpoints on which it will respond! All services are clearly marked "public" or "private" to indicate if you can use them.



9. You can compare current configurations of running instances of Hermes-based services, making it easy to spot production anomalies!



10. You can give the Tools team **feedback** and **suggest new features** that will help you in your work! ➡ tools@ or [#tools](https://twitter.com/tools)

Impact

About 200 WAU/400 MAU of ~800 in TPD

Teams integrating features, making them easier to find

Teams talking about features => more consistent

System-Z mentioned as 'great' in 2016 'What sucks'

Swiss Army Knife (good?)

Conclusions

Microservices => many small things, big picture is hard

Metadata about microservices helps understand the system

Our metadata is dirty; this is probably unavoidable

Combining many tools => better collaboration and consistency

Questions?

