High Availability Networking in AWS VPC

By: Warren Turkal

2015.03.16

signal fx

Who am I?

- Infrastructure Hacker at SignalFx
- 12+ years experience in tech
- 3+ years hacking on AWS

What I'll Be Talking About

- Traditional AWS solution
- My requirements
- Details of my solution
- Additional benefits

Traditional VPC Design Overview

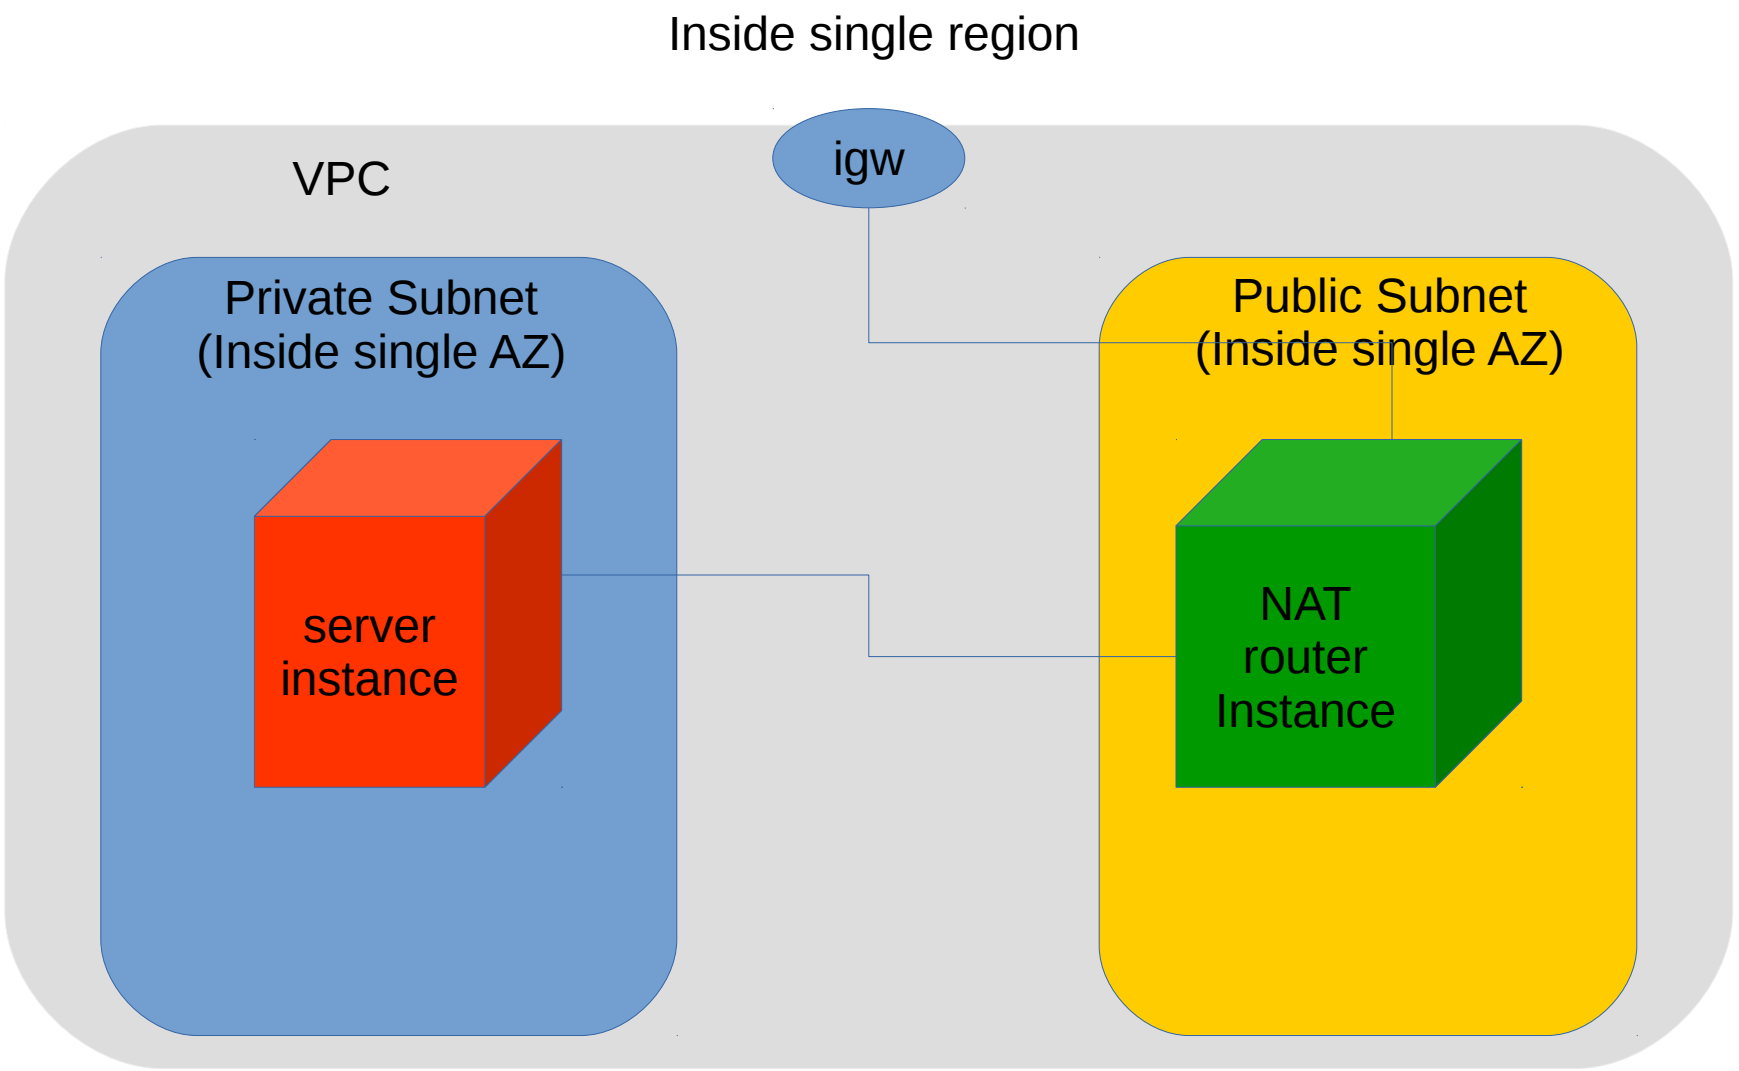

Problems with the Traditional Design

- NAT router is a SPOF
- No quick failover for NAT router failure
- No layer-3 networking between hosts in different regions

My Design Criteria

- Provider agnostic abstraction
- Abstraction boundaries reflect the cost/performance/reliability boundaries inside AWS and other cloud providers
- Isolatable zones of maintenance
- Enable developers to self serve infrastructure (e.g. hosts and load balancers)
- Work within our current processes for software deployment
- Allow experimentation

The Onion of Abstractions

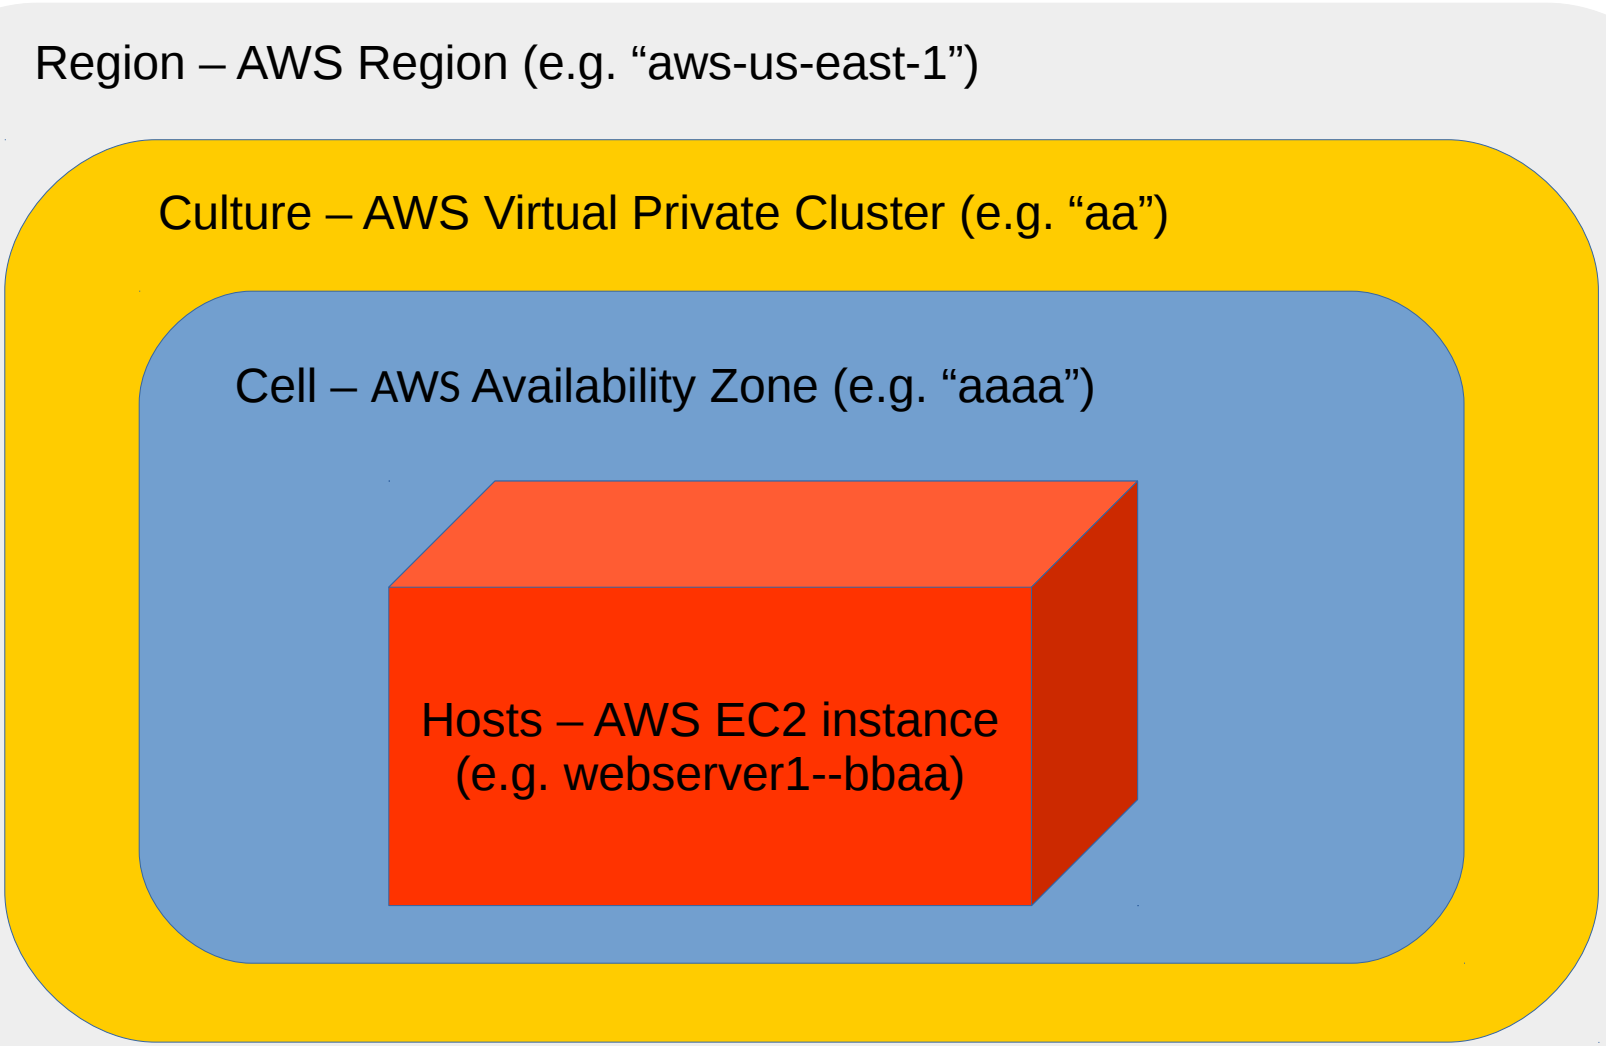

The Onion of Abstractions

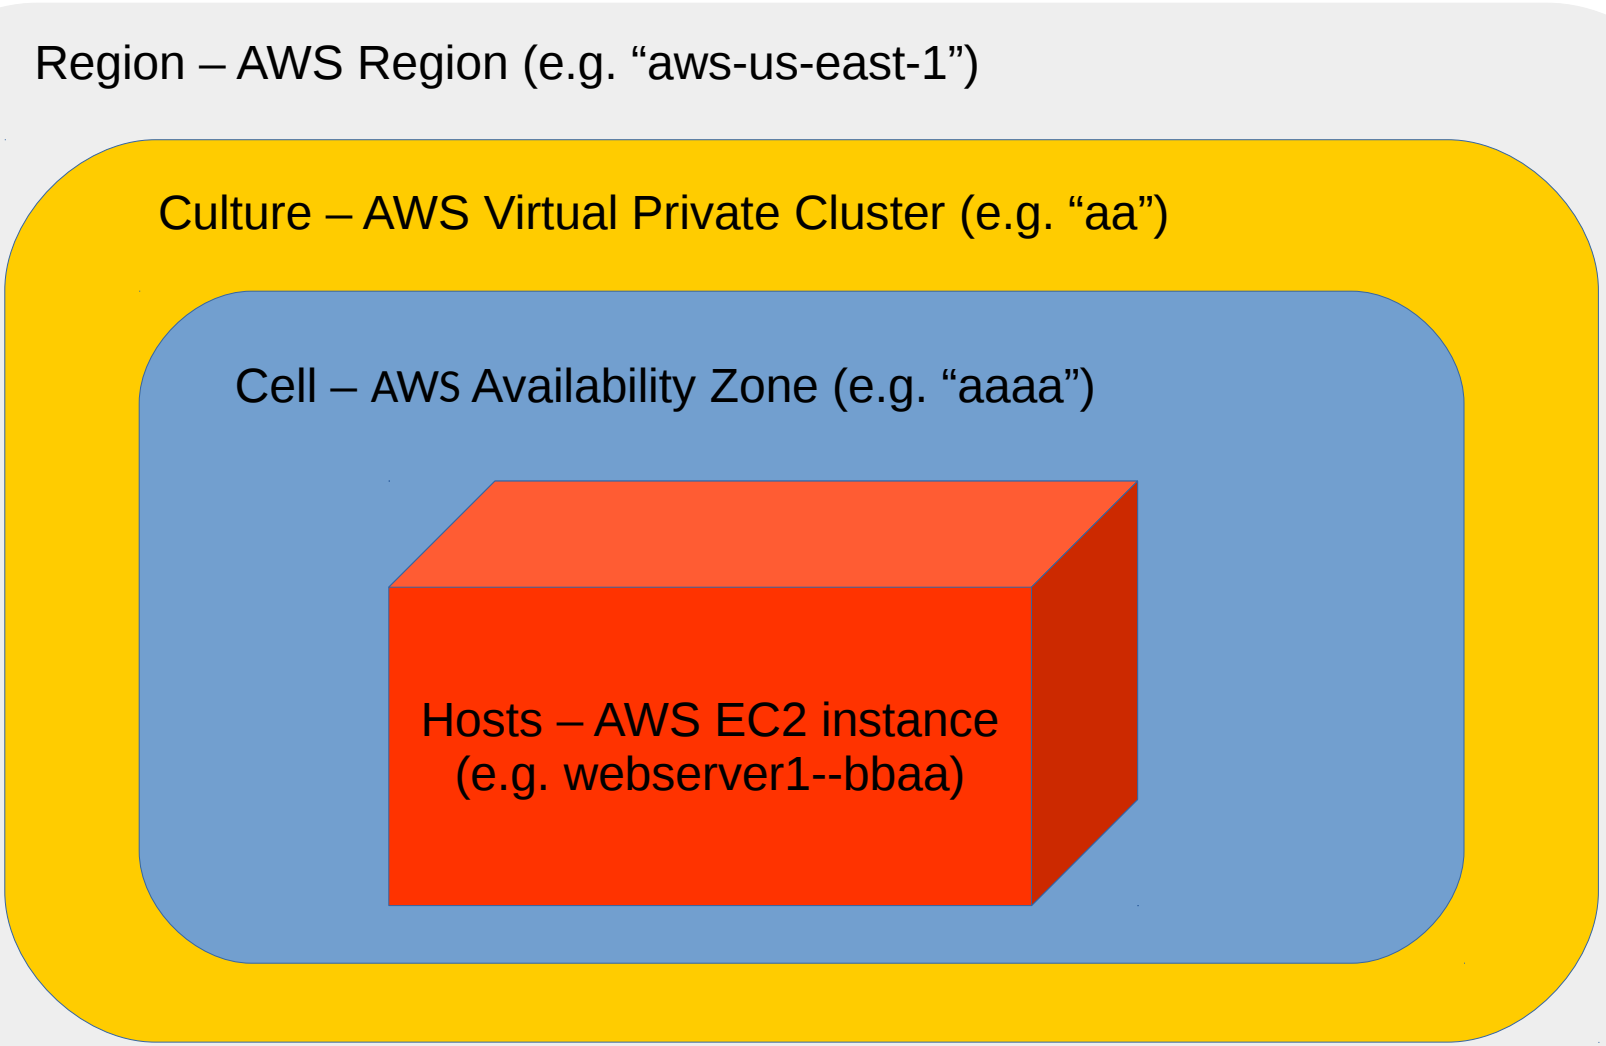

Cell Design

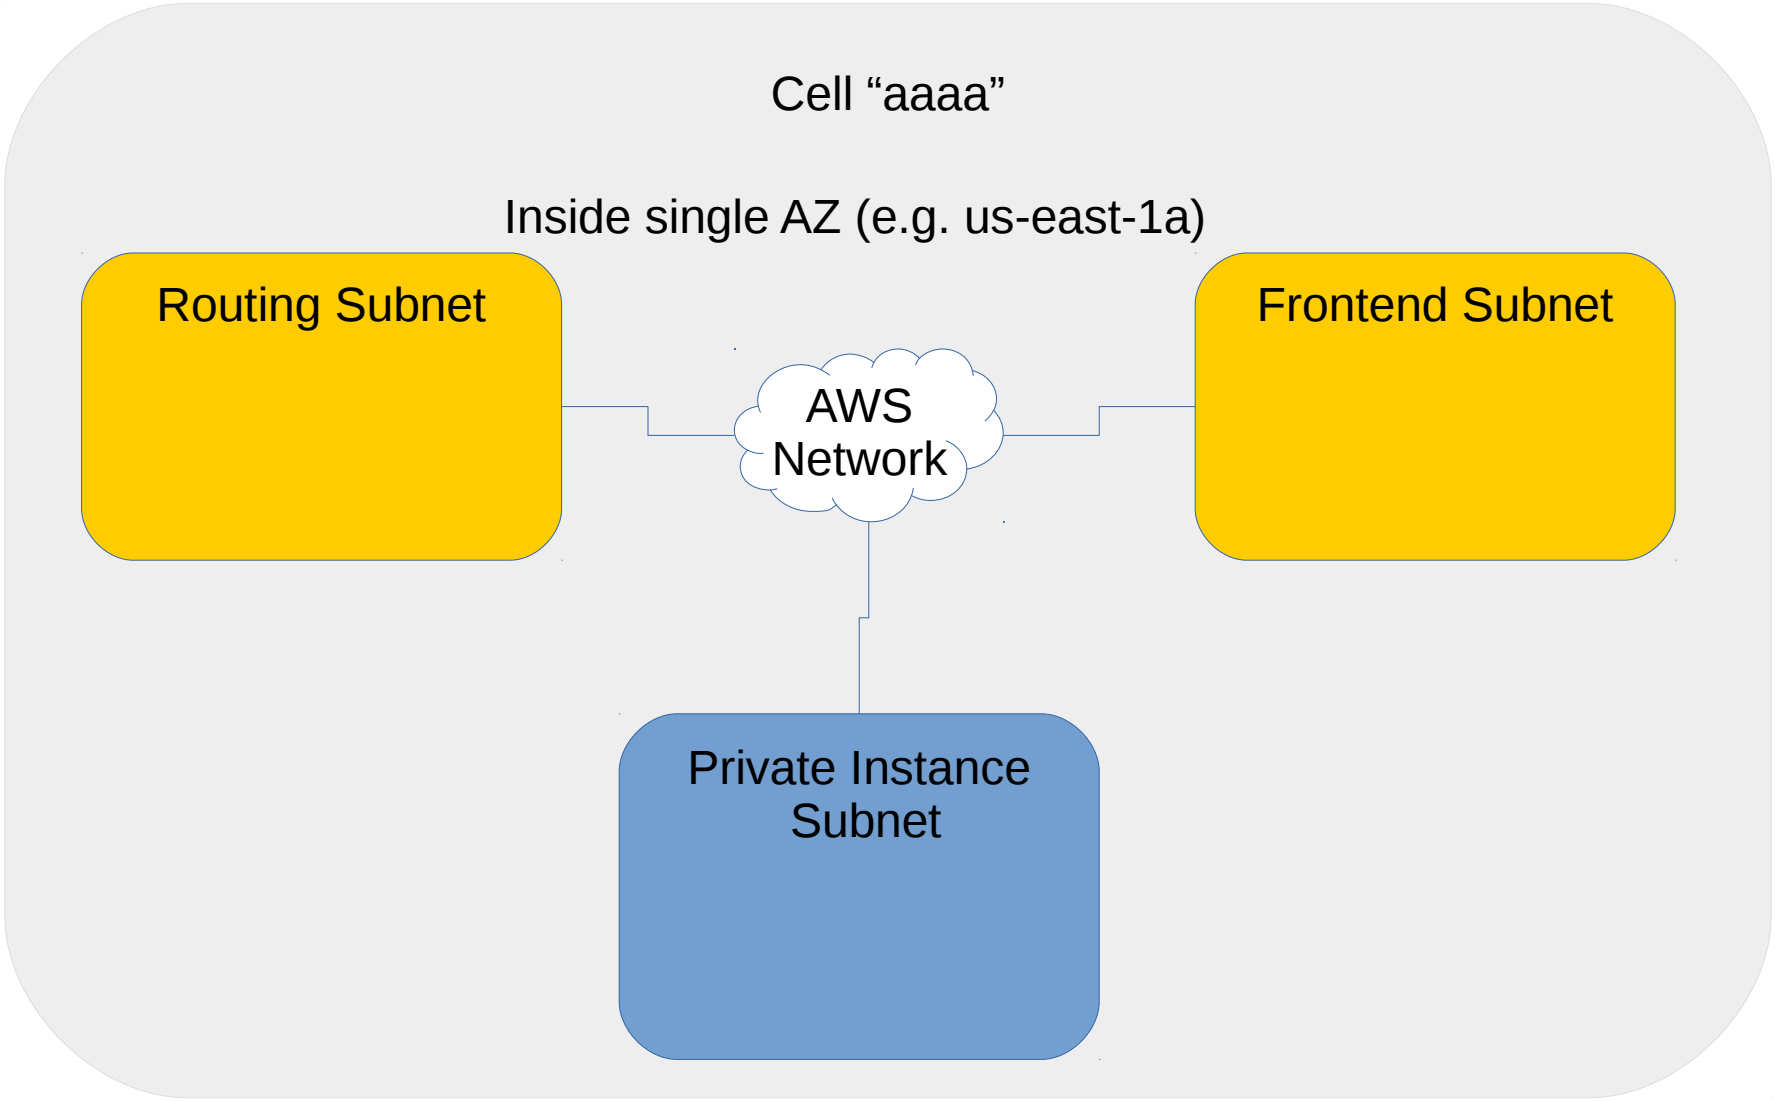
Cell “aaaa”
Inside single AZ (e.g. us-east-1a)
Routing Subnet
AWS Network
Frontend Subnet
Private Instance Subnet

Culture Design

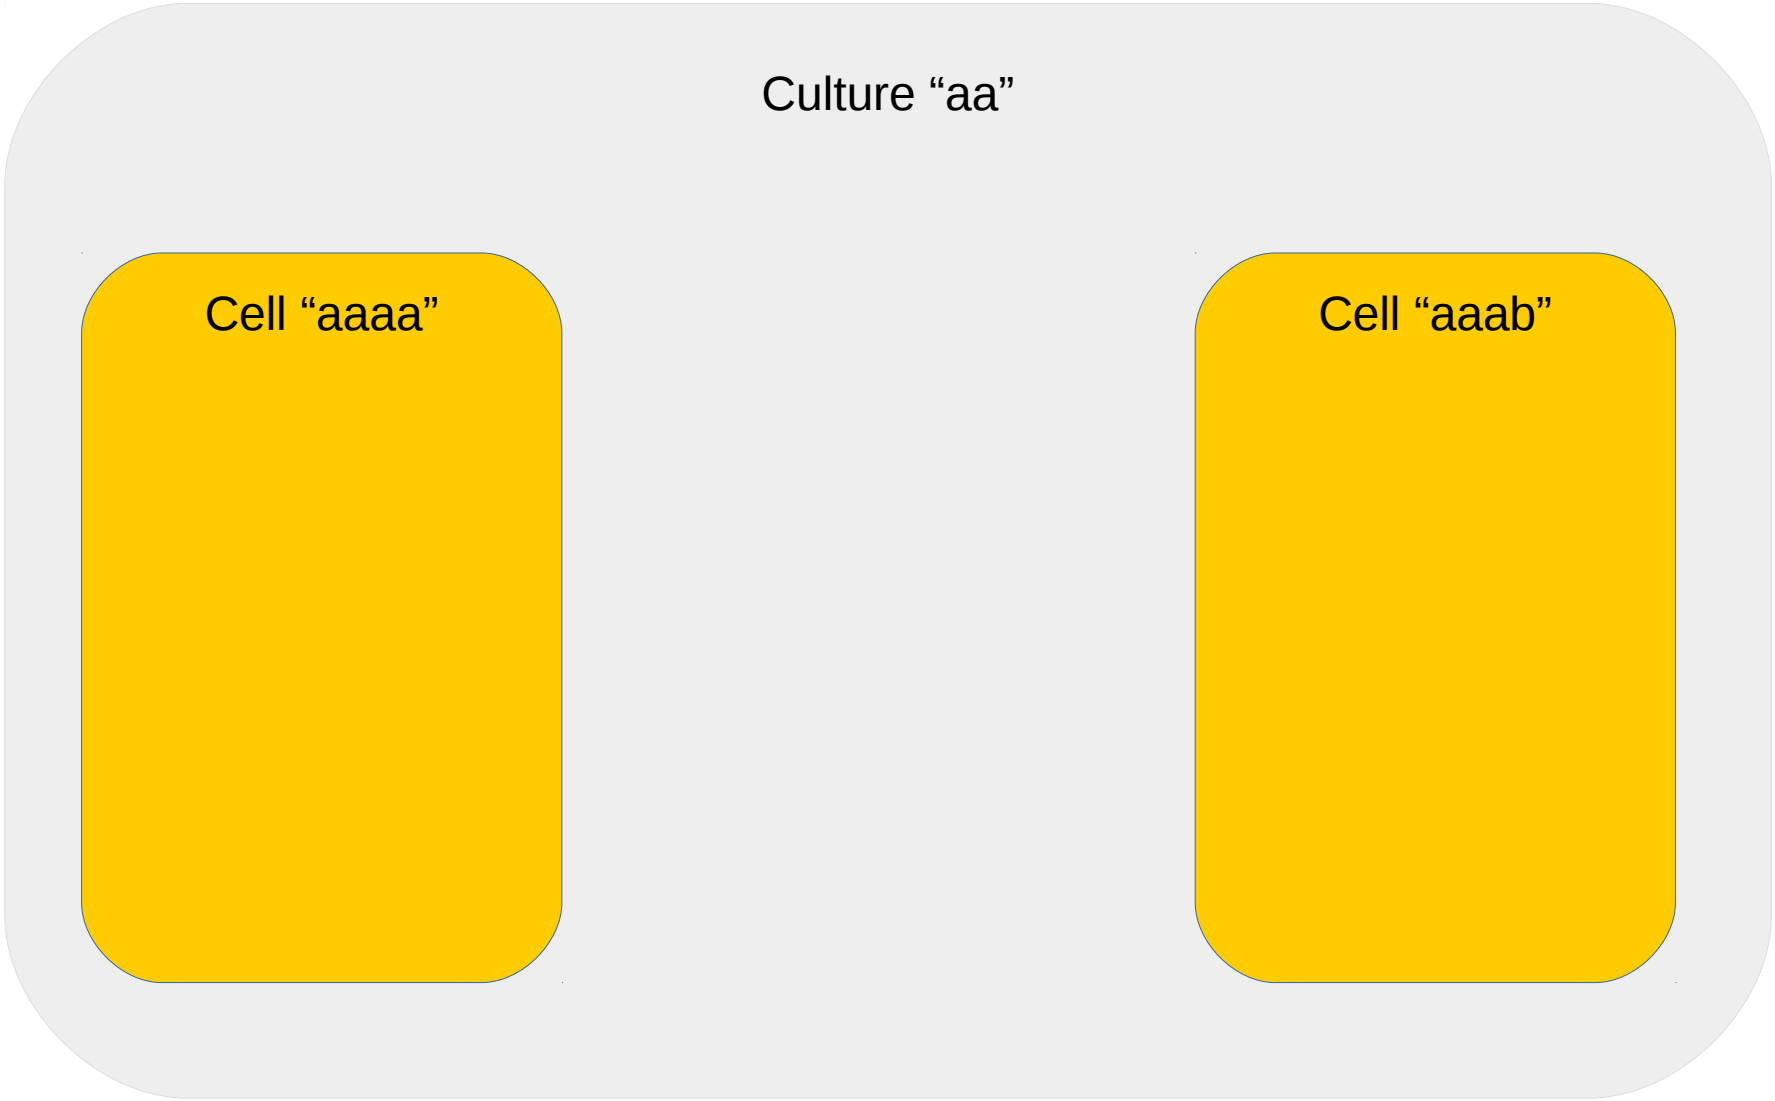

Cell Routing and Culture Wiring

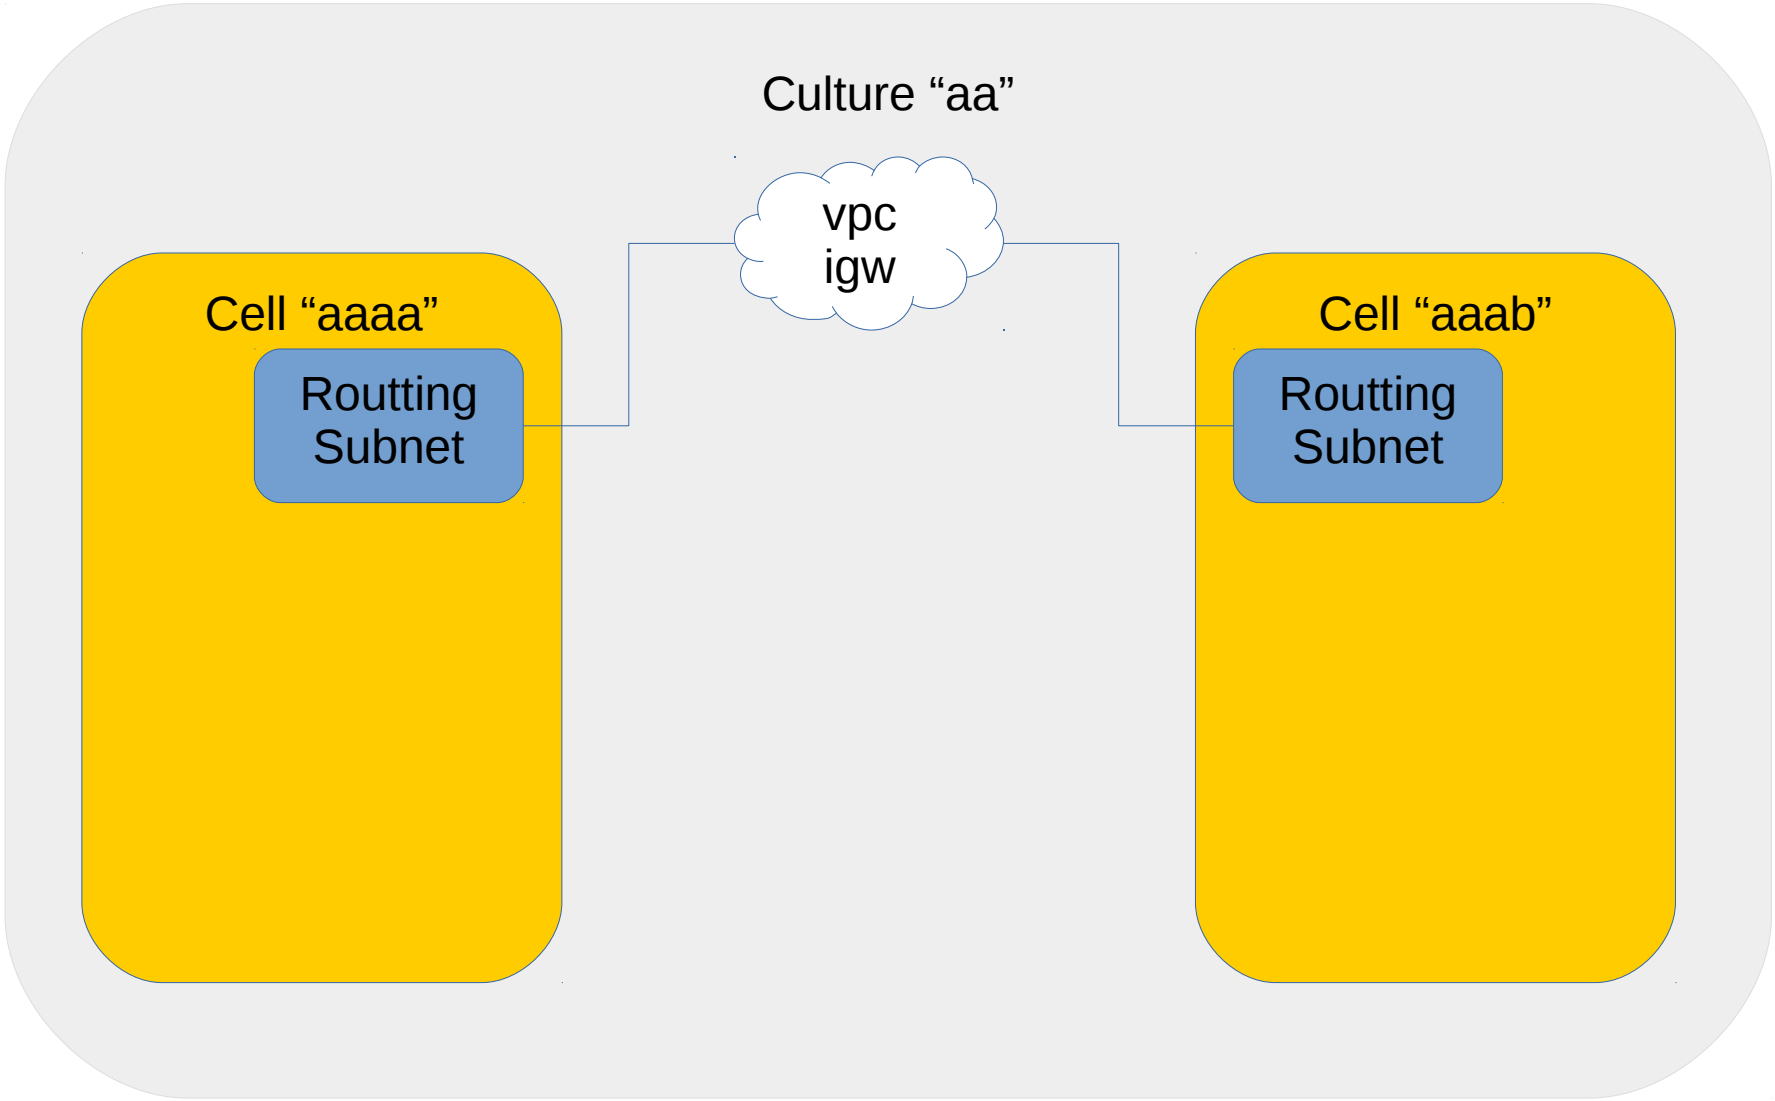

Regional/Global Design

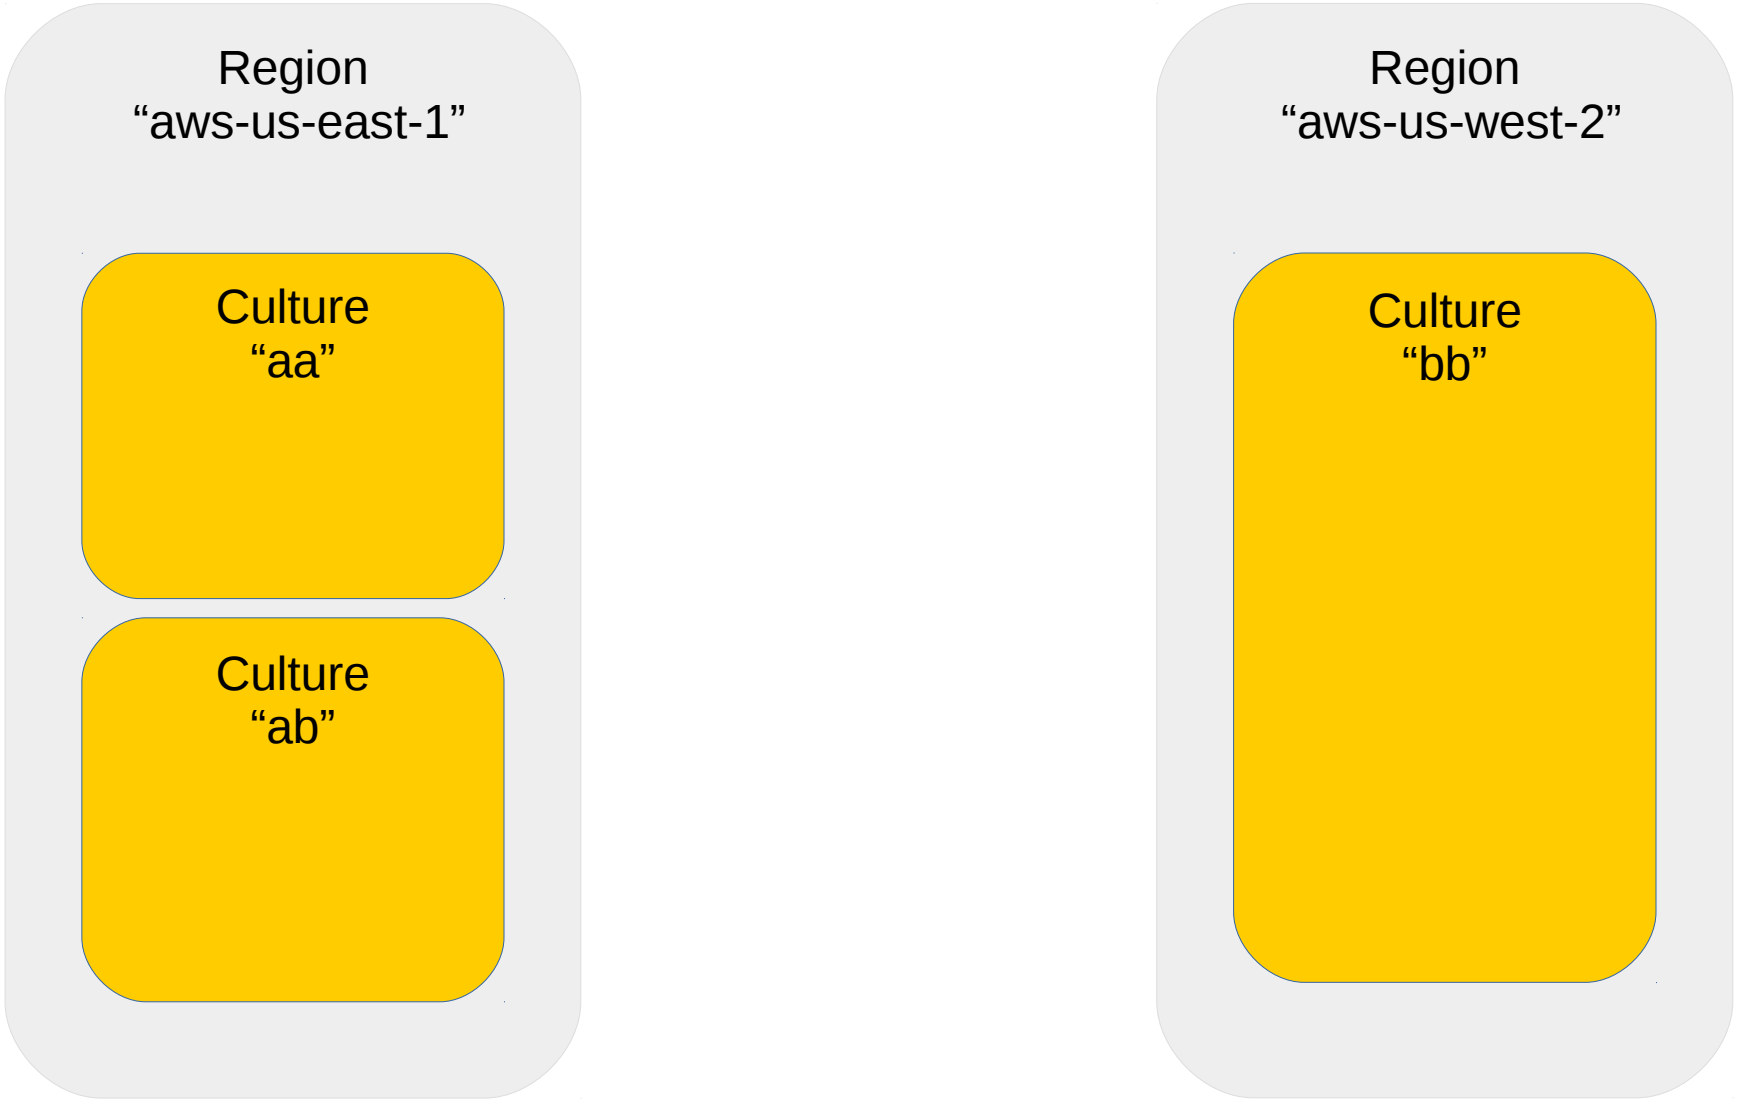

Interregion Highlevel Networking

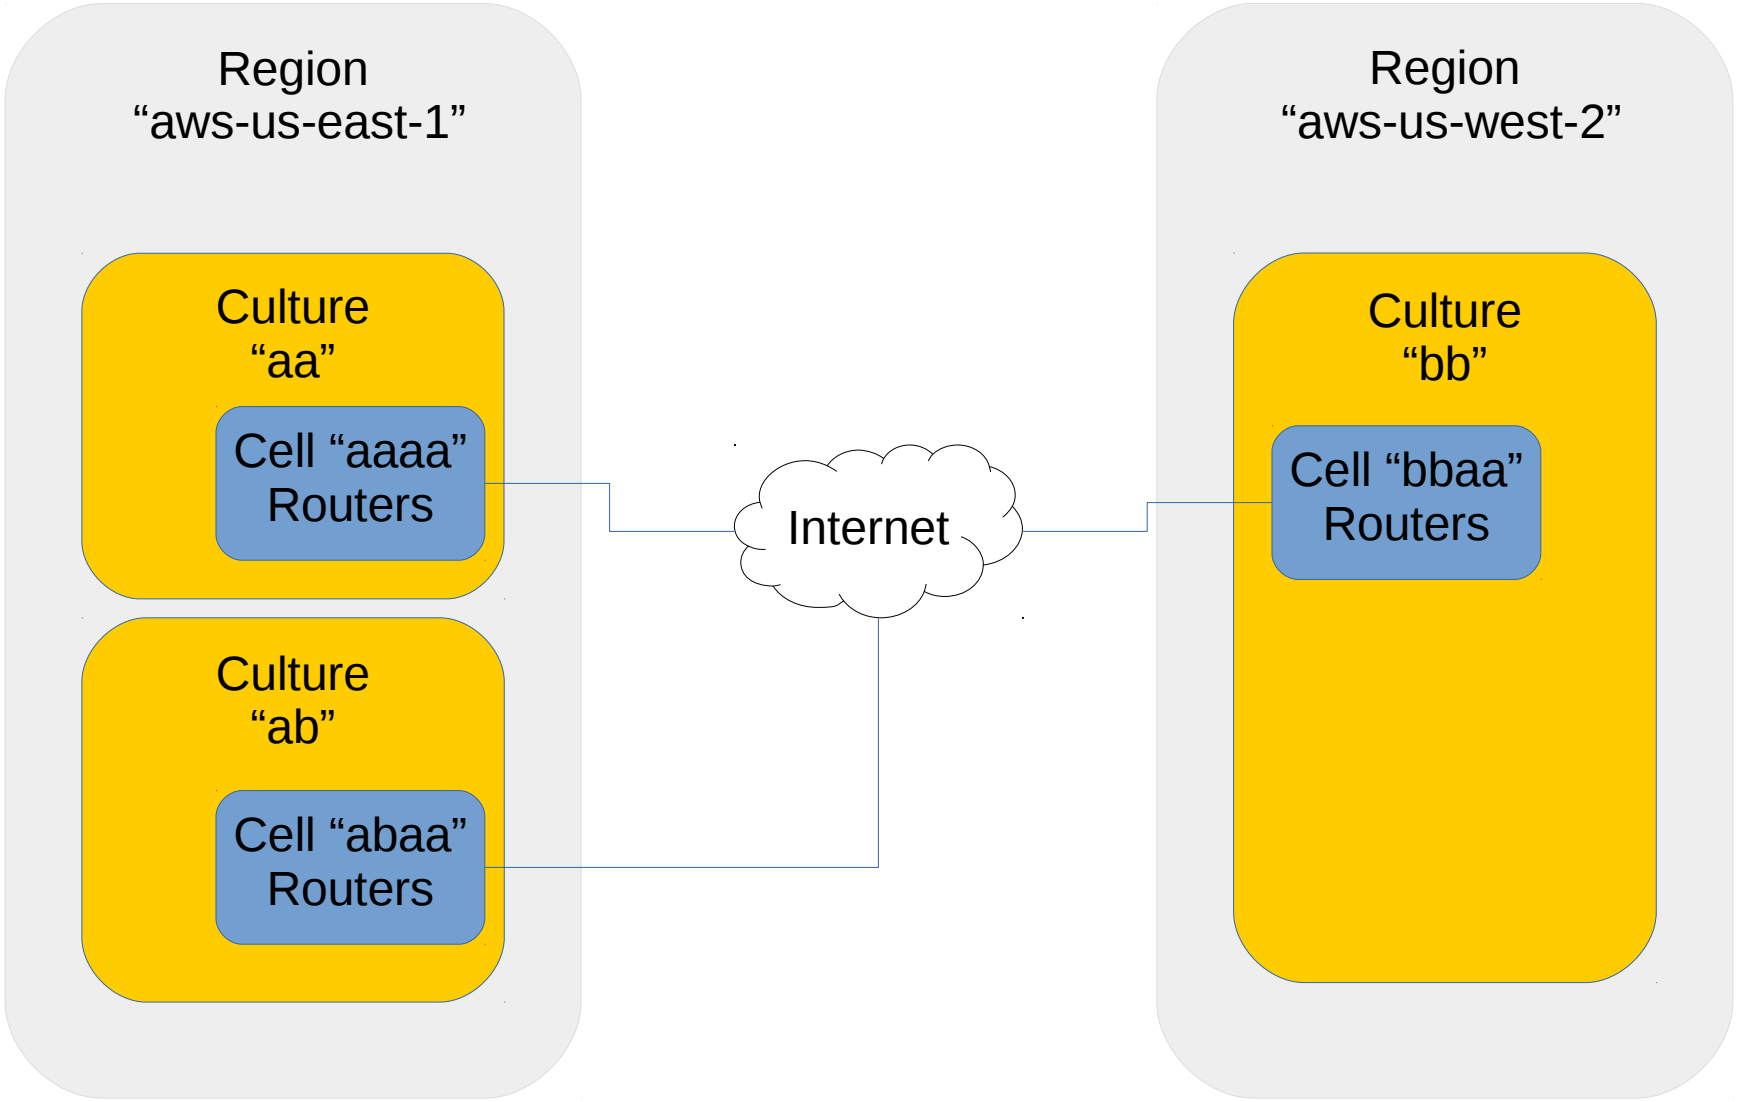

Add a Corp VPN

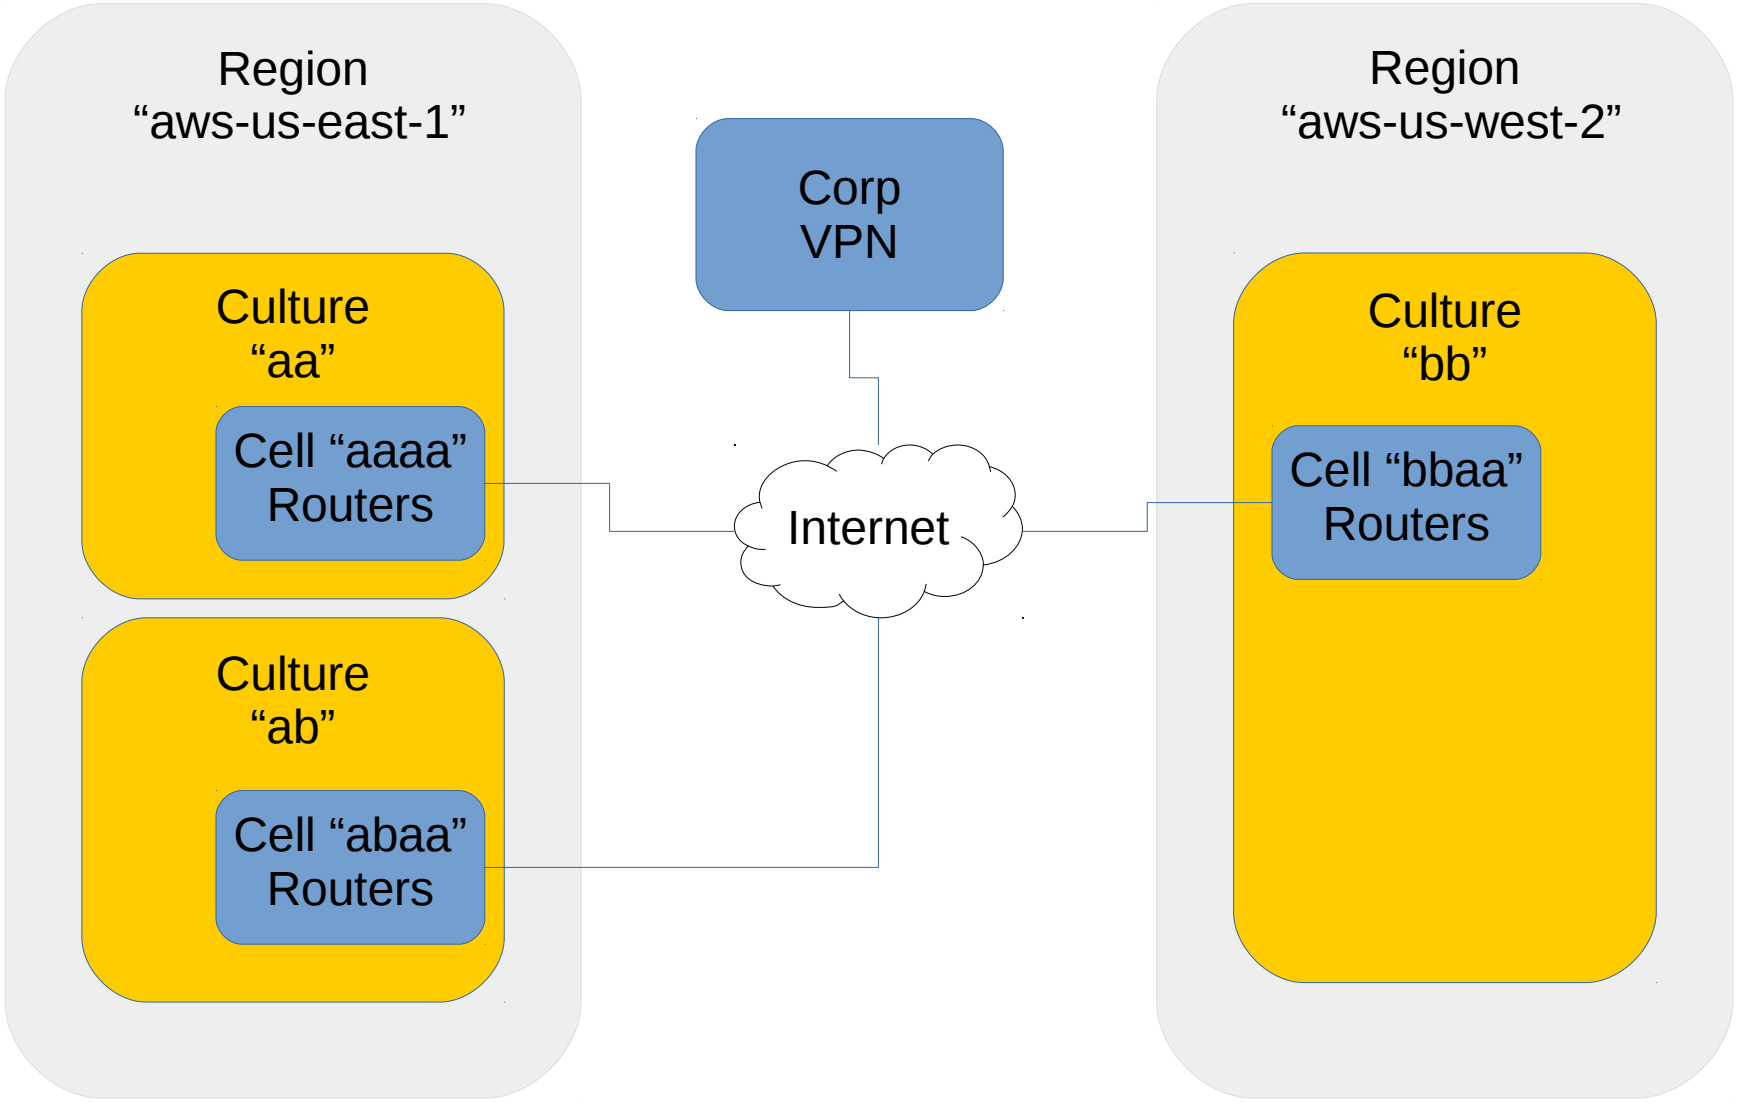

Cell Internet Connectivity

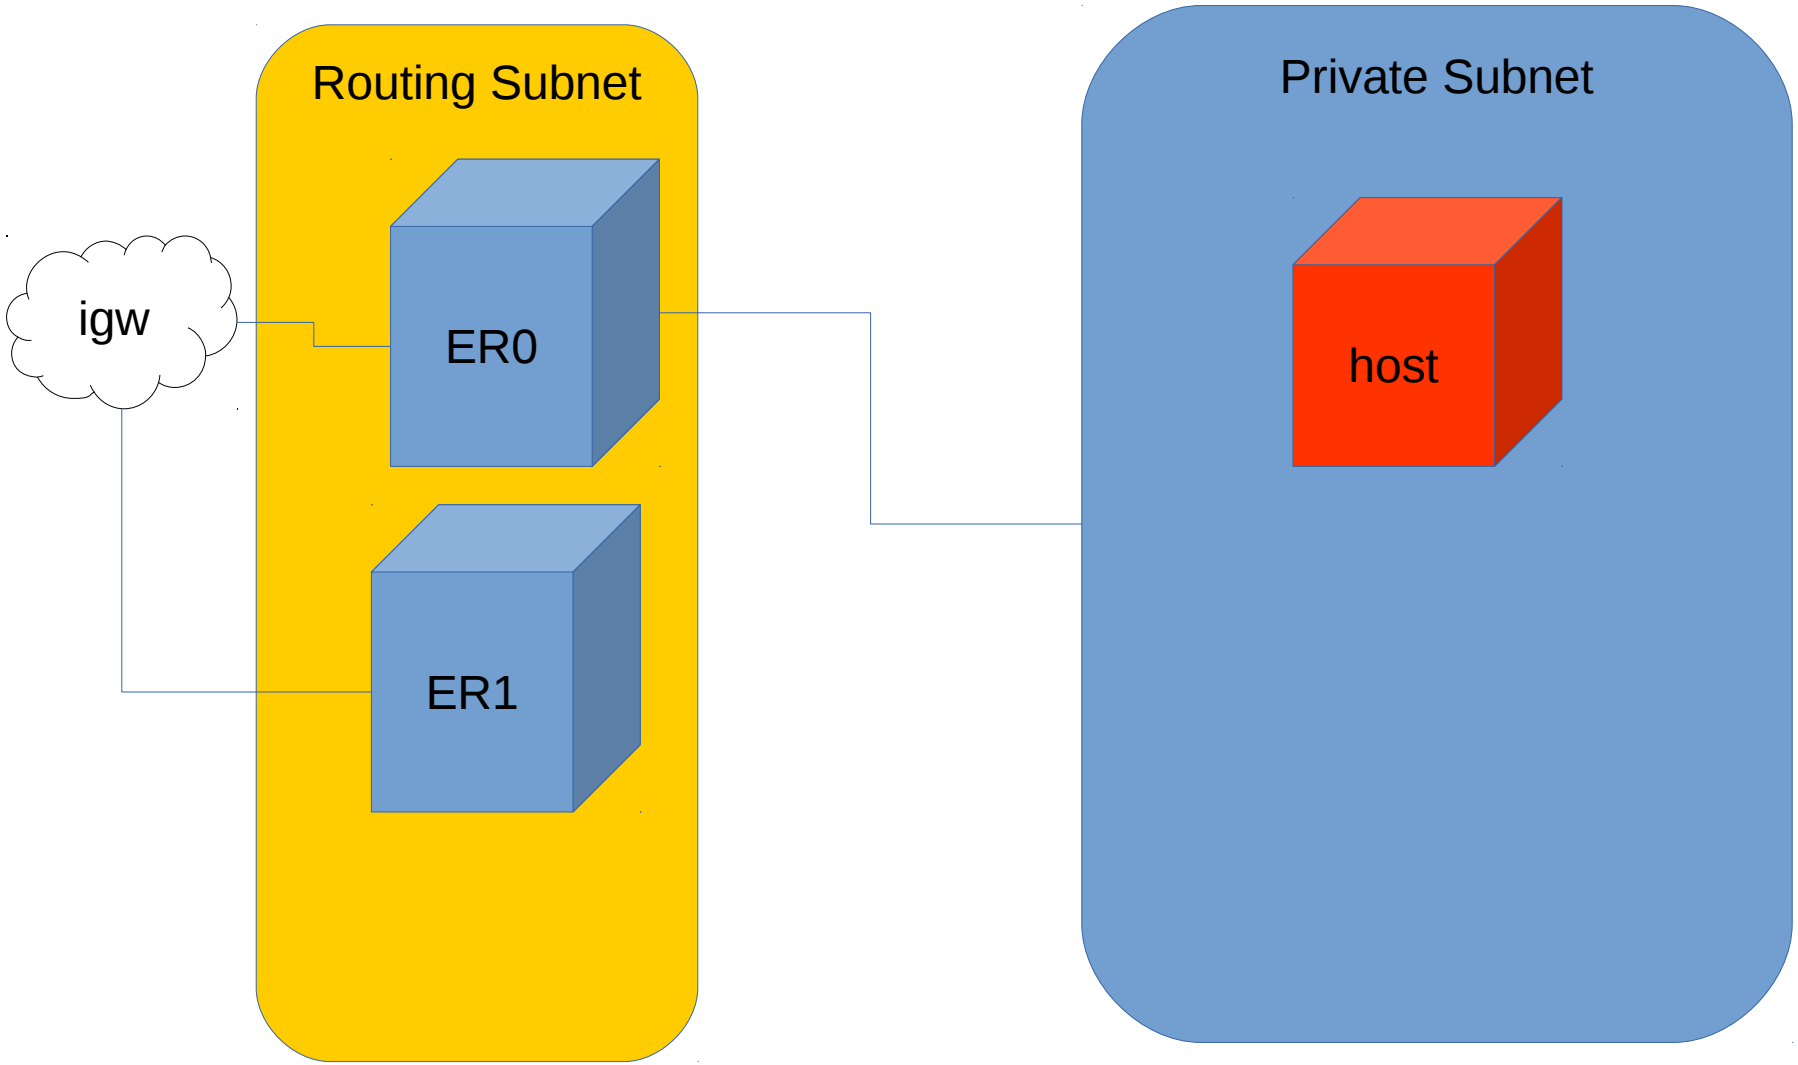

Intracell Routing

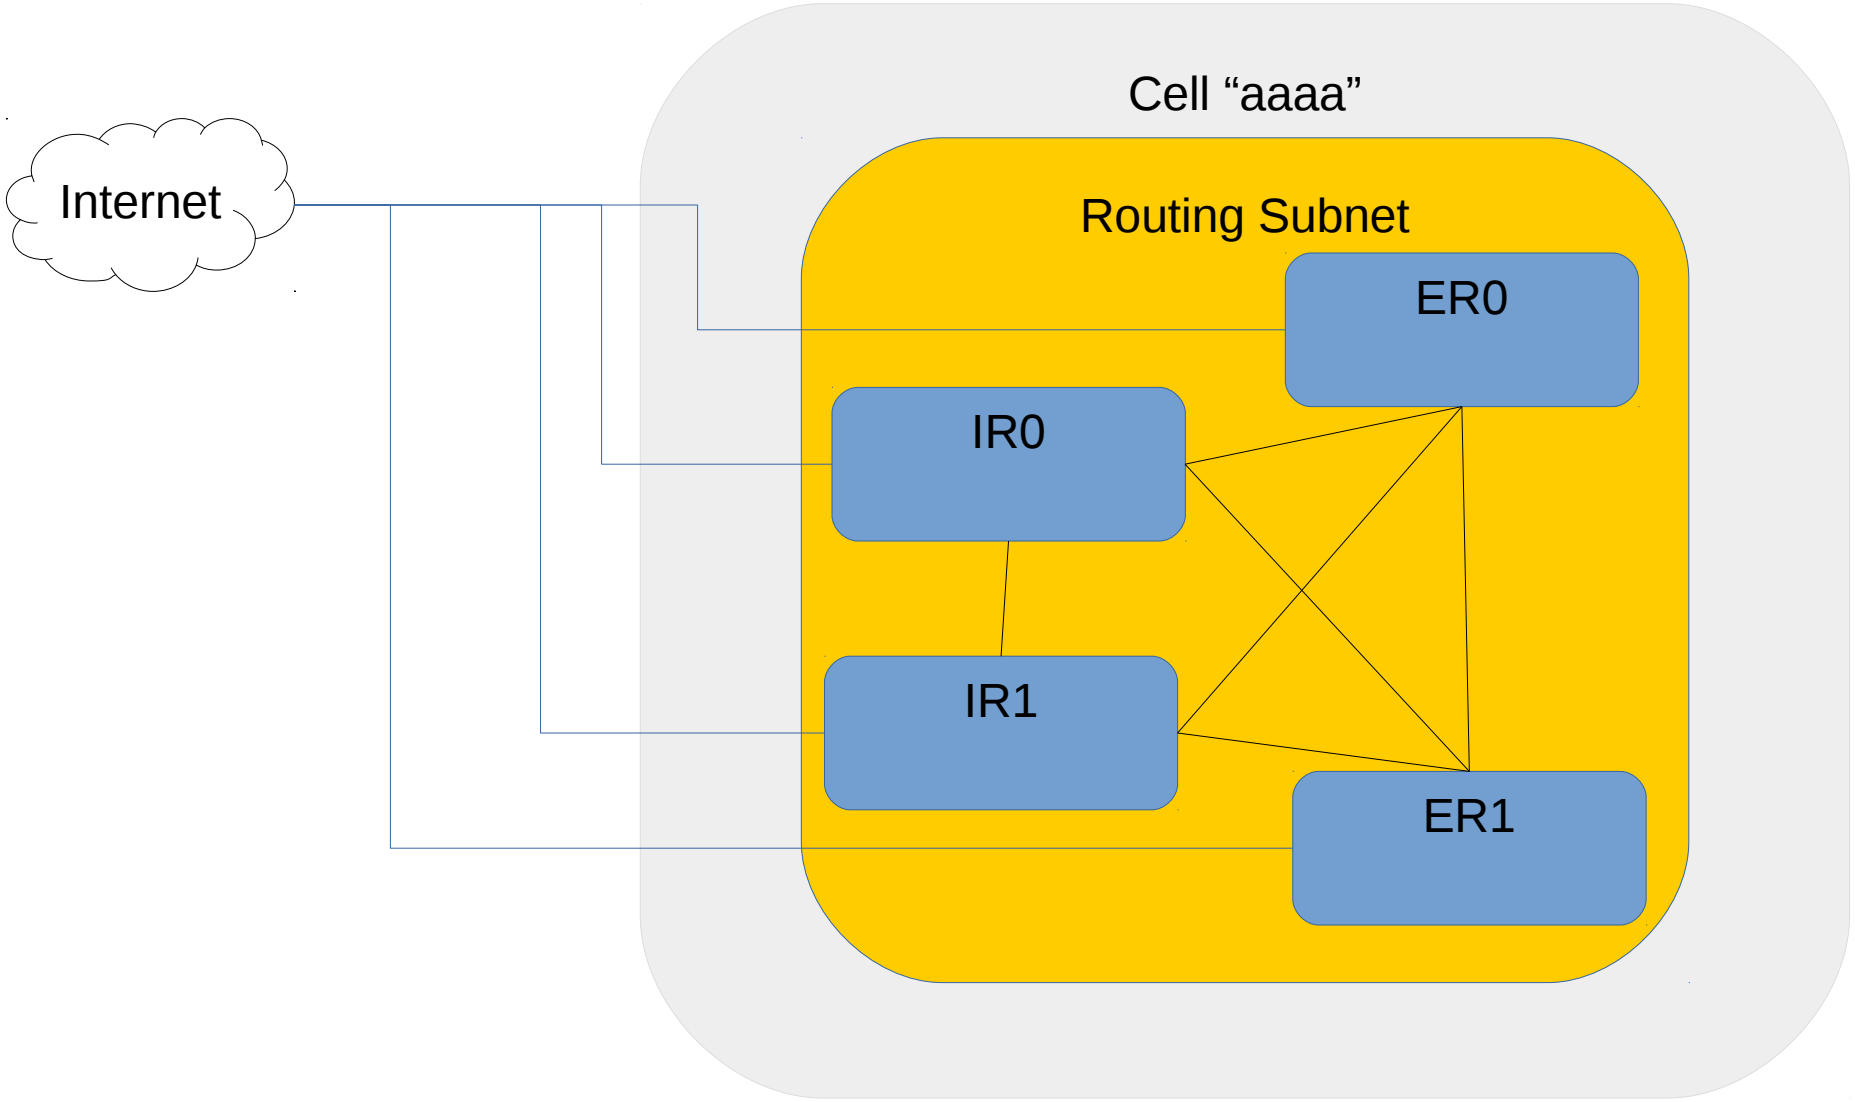

Intracell Routing

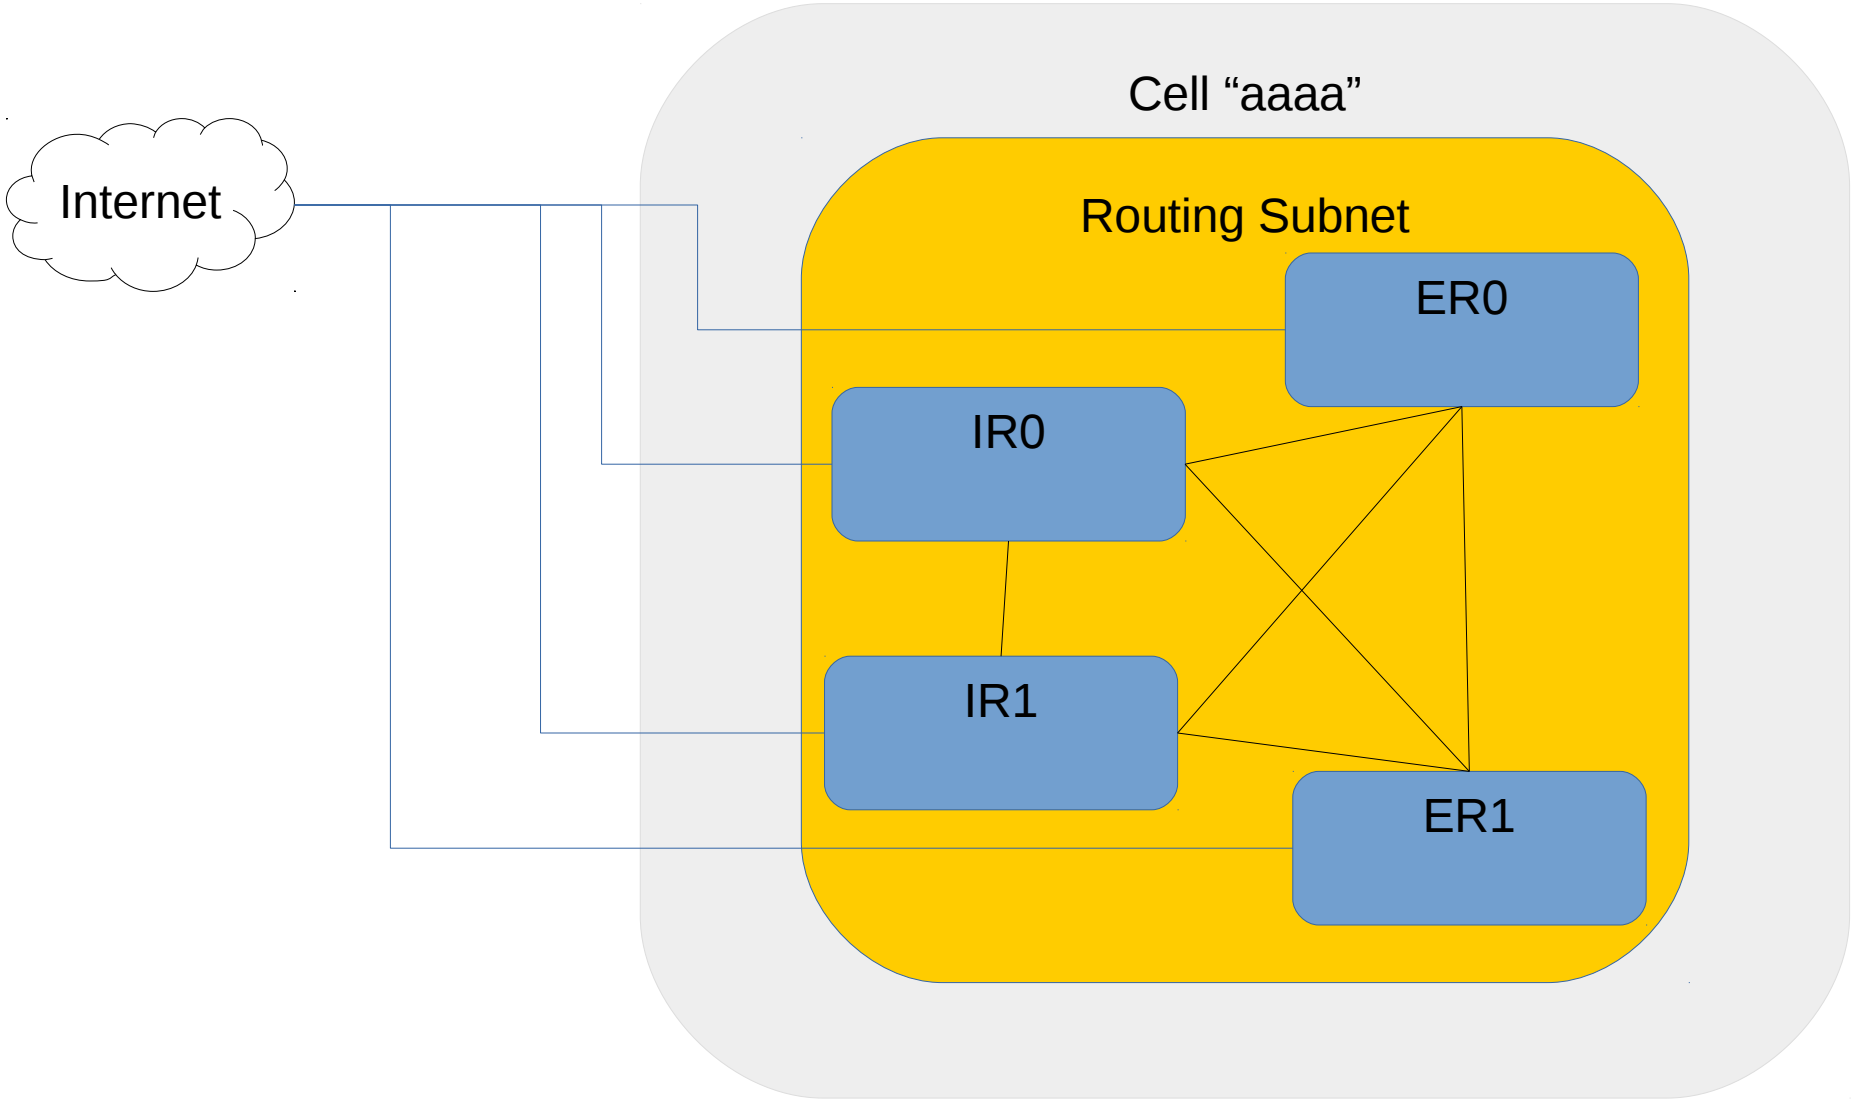

Intercell Routing

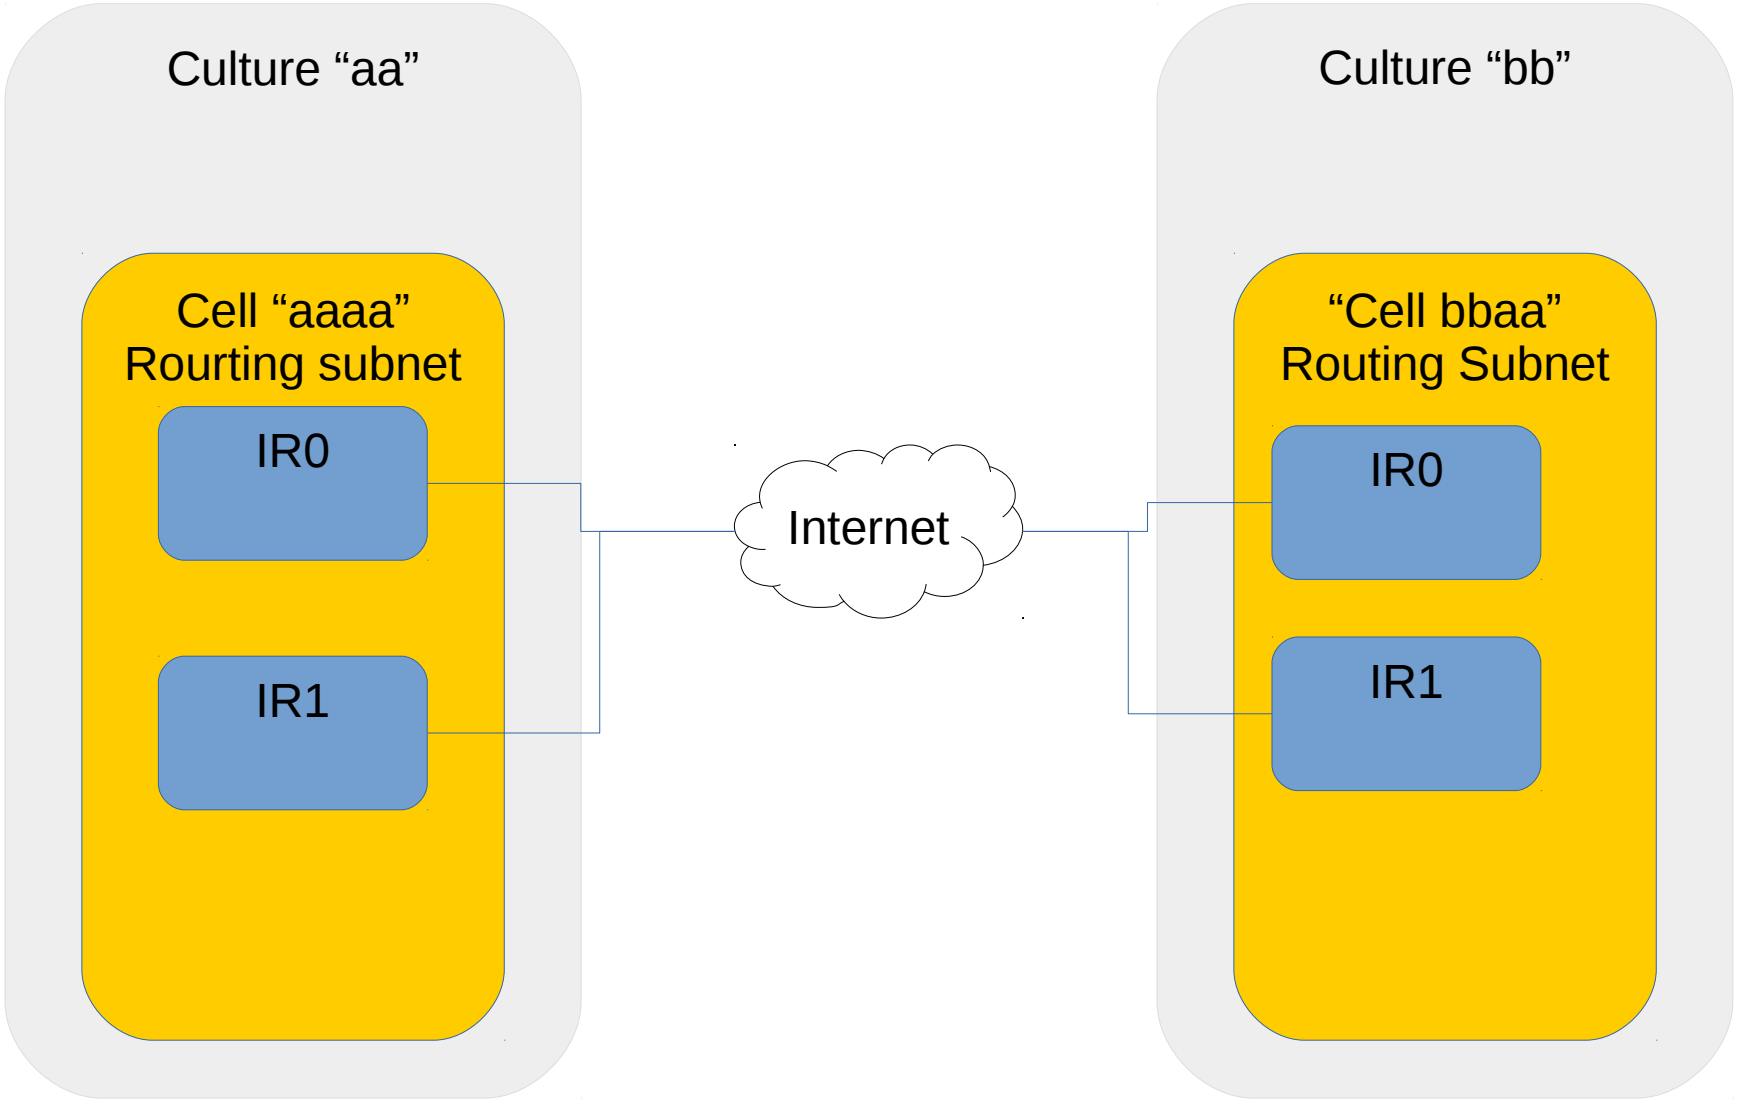

Intercell Routing (VPN tunnels)

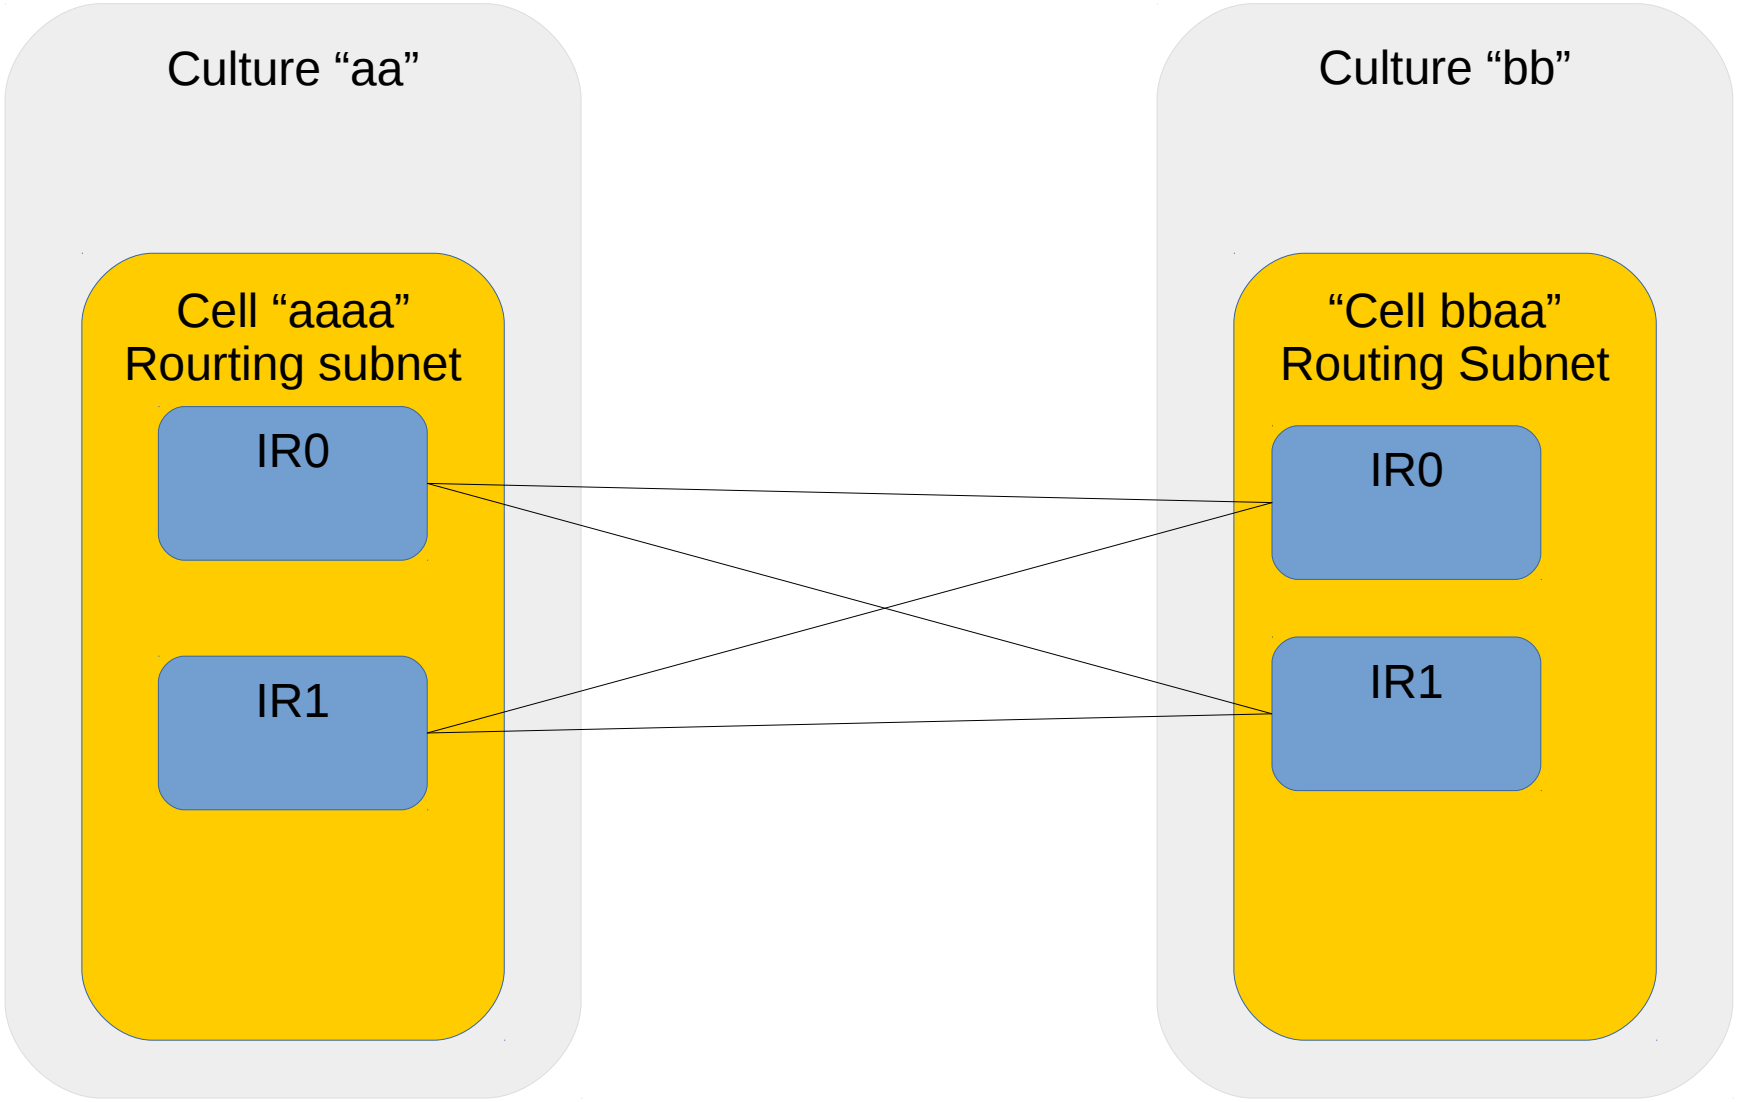

Eliminate ER Hop in Intercell Connectivity

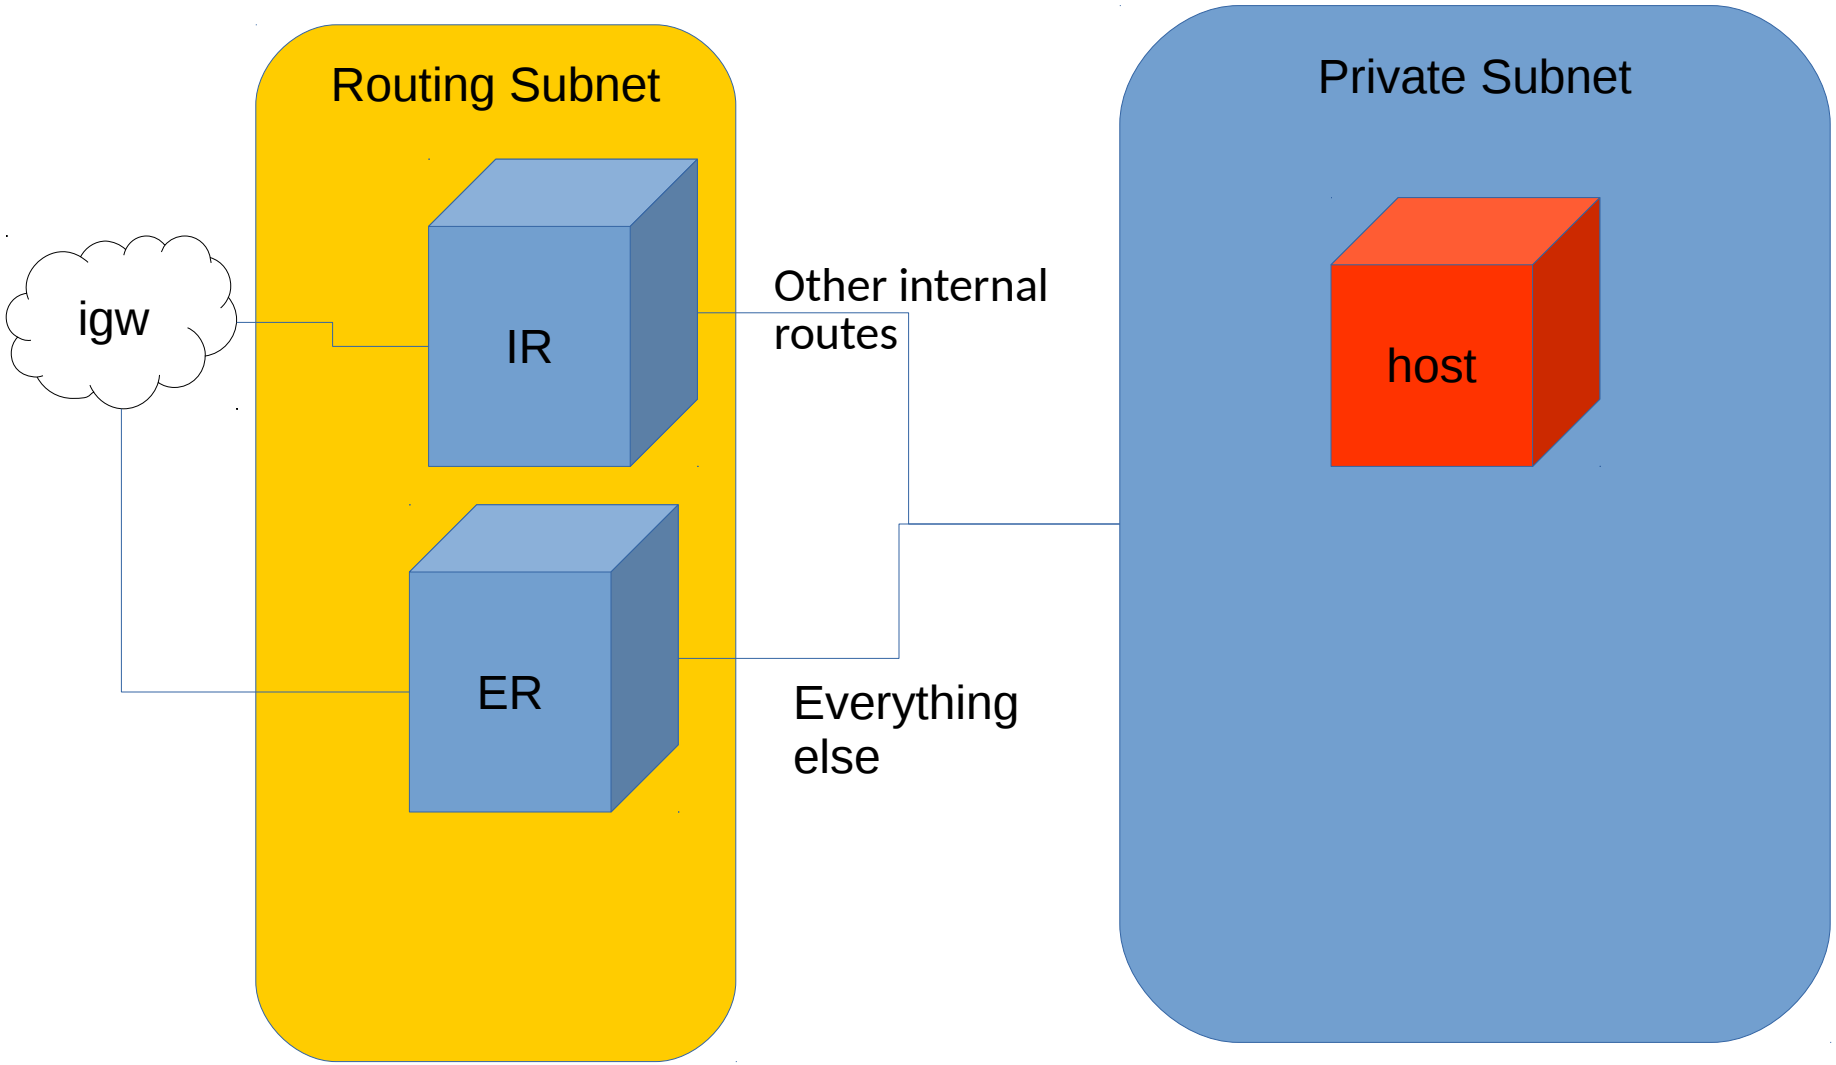

Implementation on AWS

- Cloudformation w/ custom scripting (python+troposphere lib+boto lib)
- Cultures and cells are managed with separate tools the rely on common lib
- Must keep track of dummy and intermediate addresses among other things for the config.
- Routing protocols are all implemented with Quagga.

Interesting Numbers

- A cell is composed of 86 AWS objects when running with 2 IR and 2 ER machines.
- A culture is composed of 22 AWS objects.

Going Beyond Just the Network in AWS VPC

- We have cell-aware and culture-aware tools to spin up/terminate instances and load balancers (ELBs on AWS).
- These tools integrate with out Salt deployment to make it easier for developer to self-service when spinning up new types of instances.

Example of spinning up instance:

```
sfhost add webserver bbaa c3.xlarge
```

Example of spinning up an ELB:

```
Sflb --culture bb --name api create
```

SignalFx – Launched last week

Streaming analytics on **multi-dimensional** metrics for monitoring modern apps

Free 30 day trial

Say hi at our table for demo+discussion

Q&A

Thank You!

signal fx