slack
from Salesforce

Network Flow Data in the Cloud

Steve Dodd, Staff Engineer
Demand Engineering

Agenda

slack
from Salesforce

Motivation

Review of Graph Theory

Implementation

Comparison to Native Solution

Data Use Cases

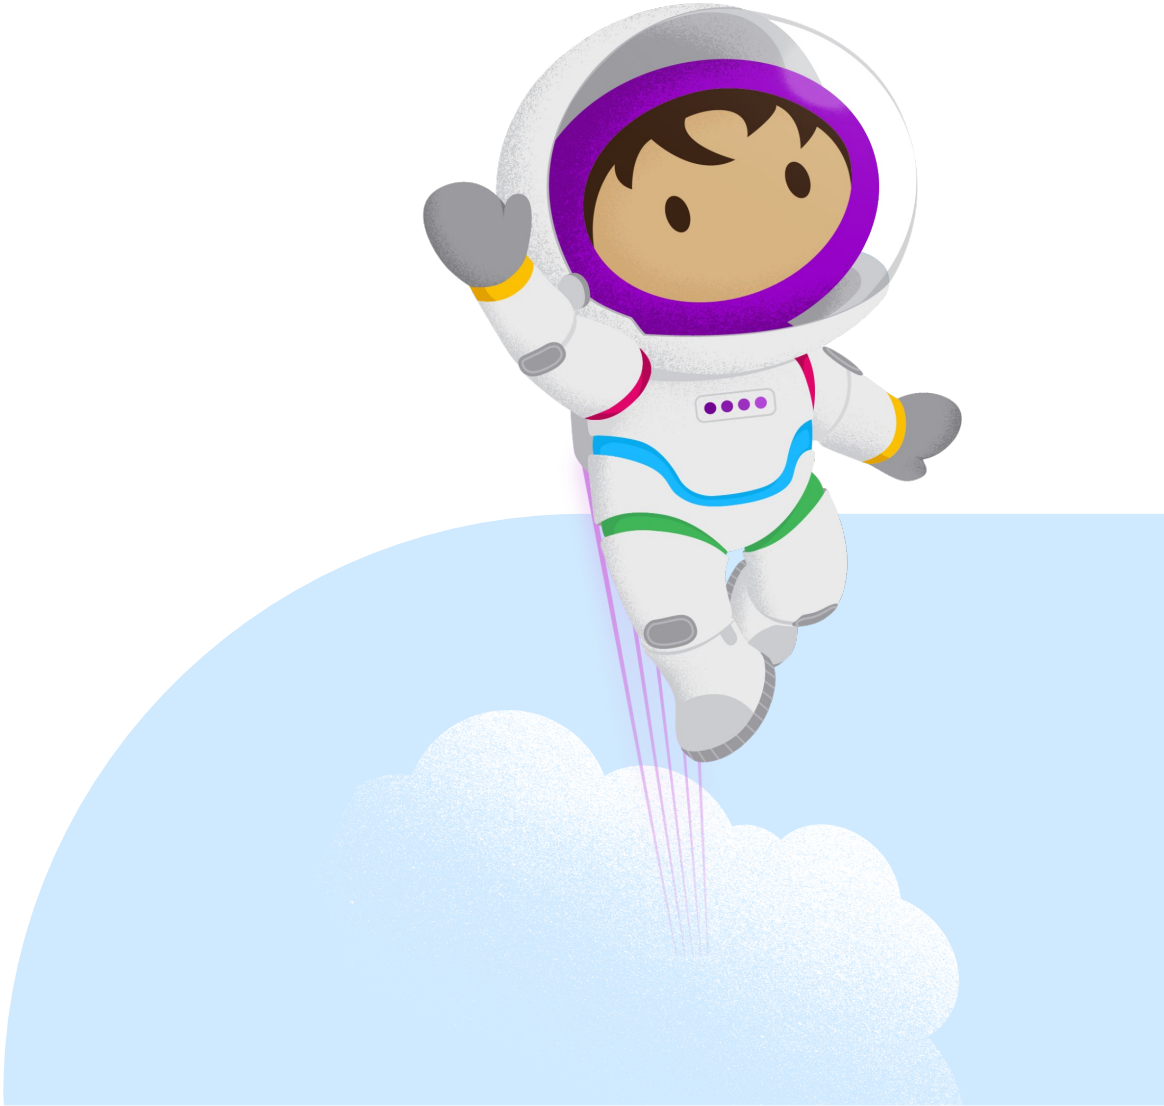

Motivation

slack from Salesforce

Mark McBride 3:53 PM

👋 As part of considering a move of wwws/caches/vitess to whitecastle, the question "how much traffic does a webapp do to other memcache/vitess/other chef roles" came up, and the answer was... it's hard to tell.

How much traffic do we have?

Where is our traffic going?

Is our ability to deliver traffic at risk?

Motivation

slack from Salesforce

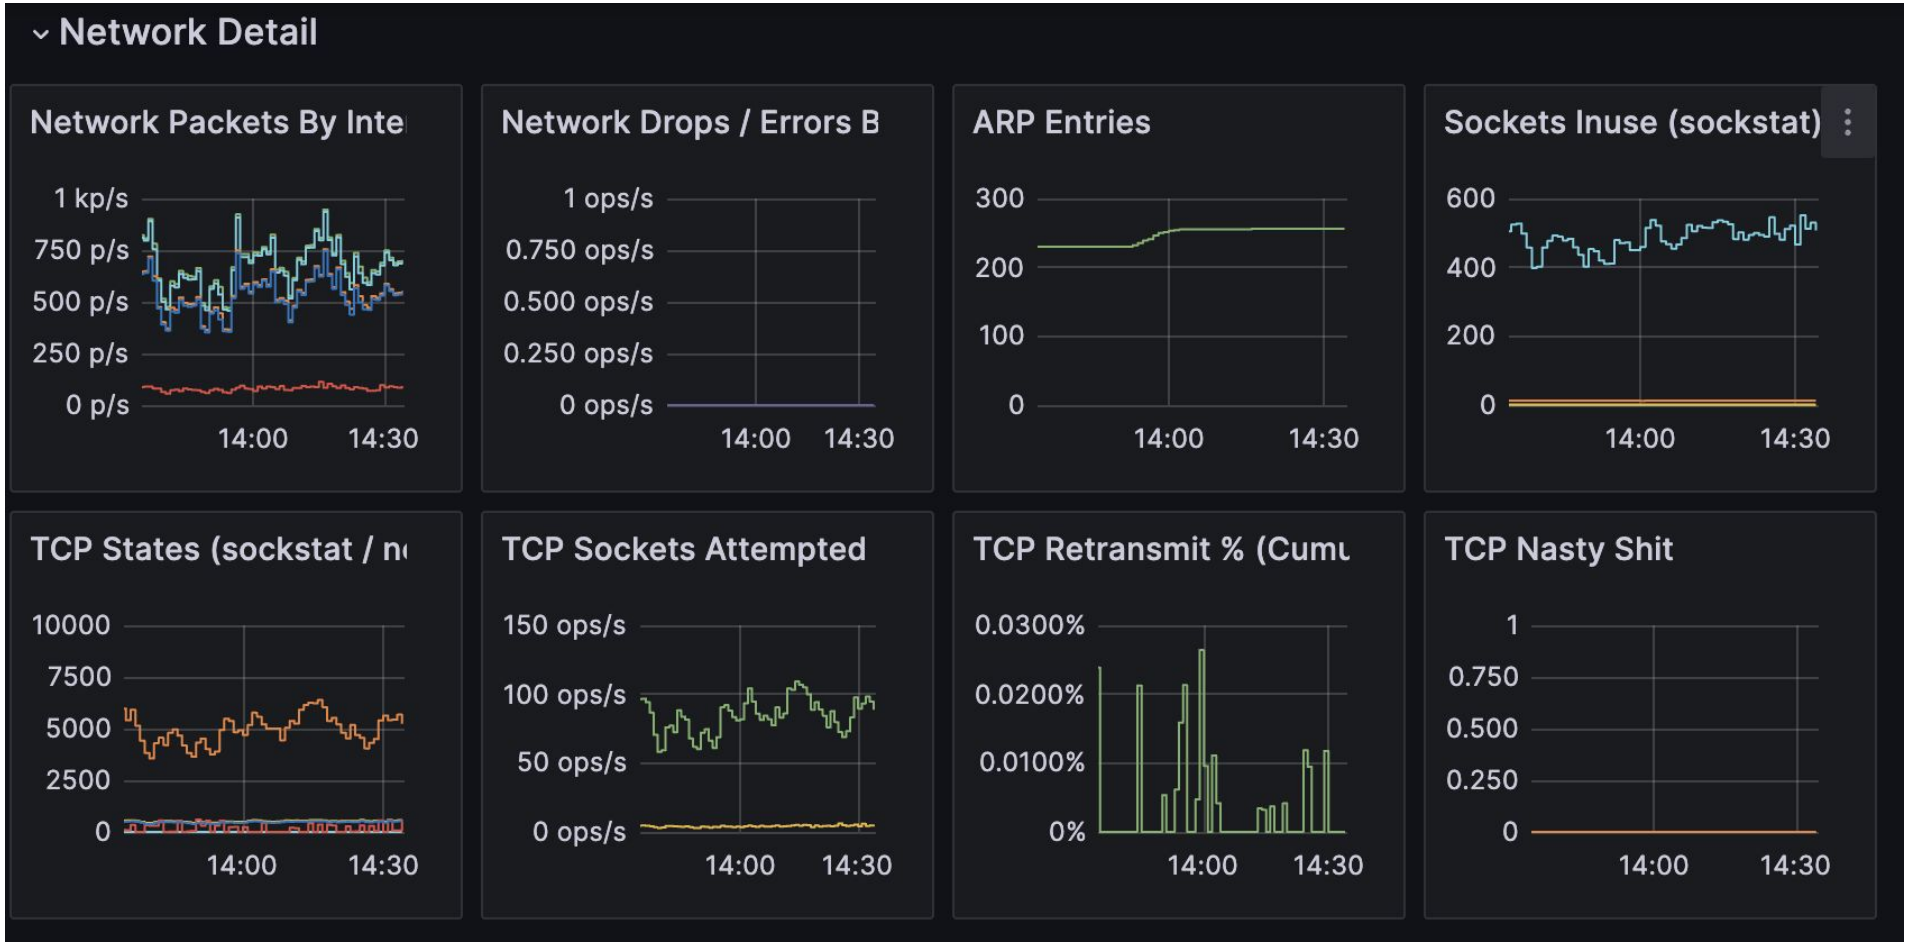

Graph Theory

Flow Network – A directed graph *G = (V, E)* with a capacity function *c* for each edge. A flow is a mapping that satisfies the constraints:

- *∀ e ∈ E*
  - *0 <= f(e) <= c(e)*
- *∀ v ∈ V*
  - *∑f(e_in) must equal ∑f(e_out) , except for the source and sink of flows*

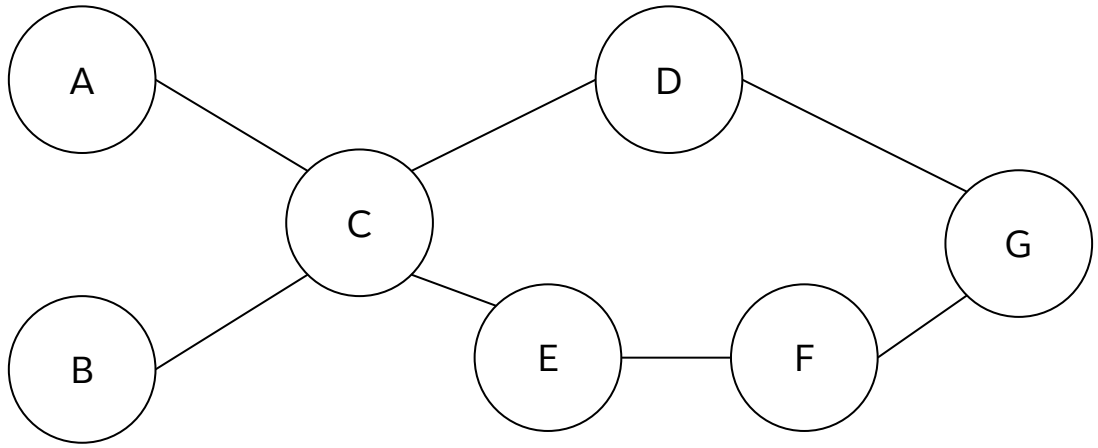

Graph Theory

slack
from Salesforce

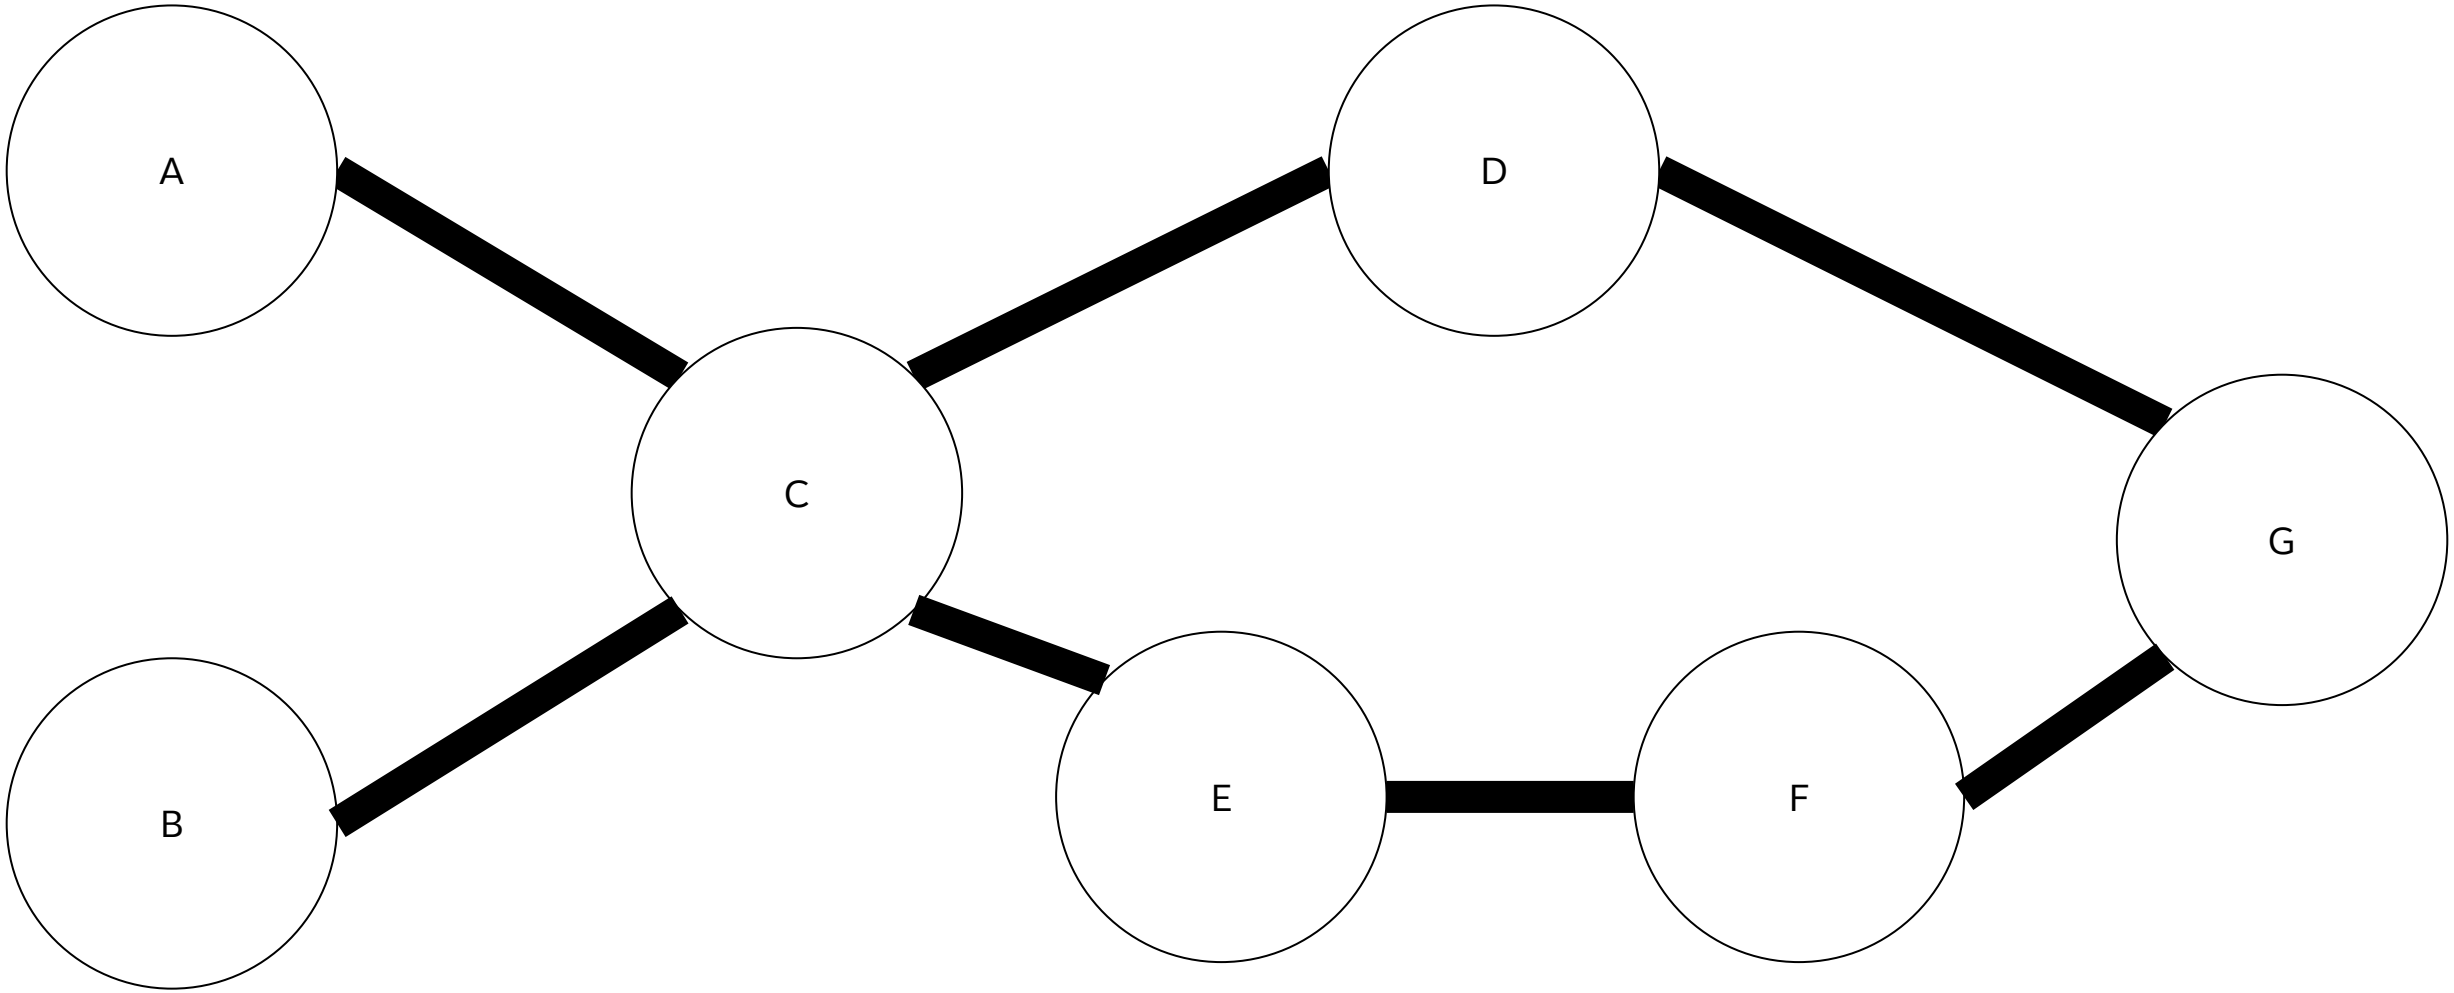

Graph Theory

slack
from Salesforce

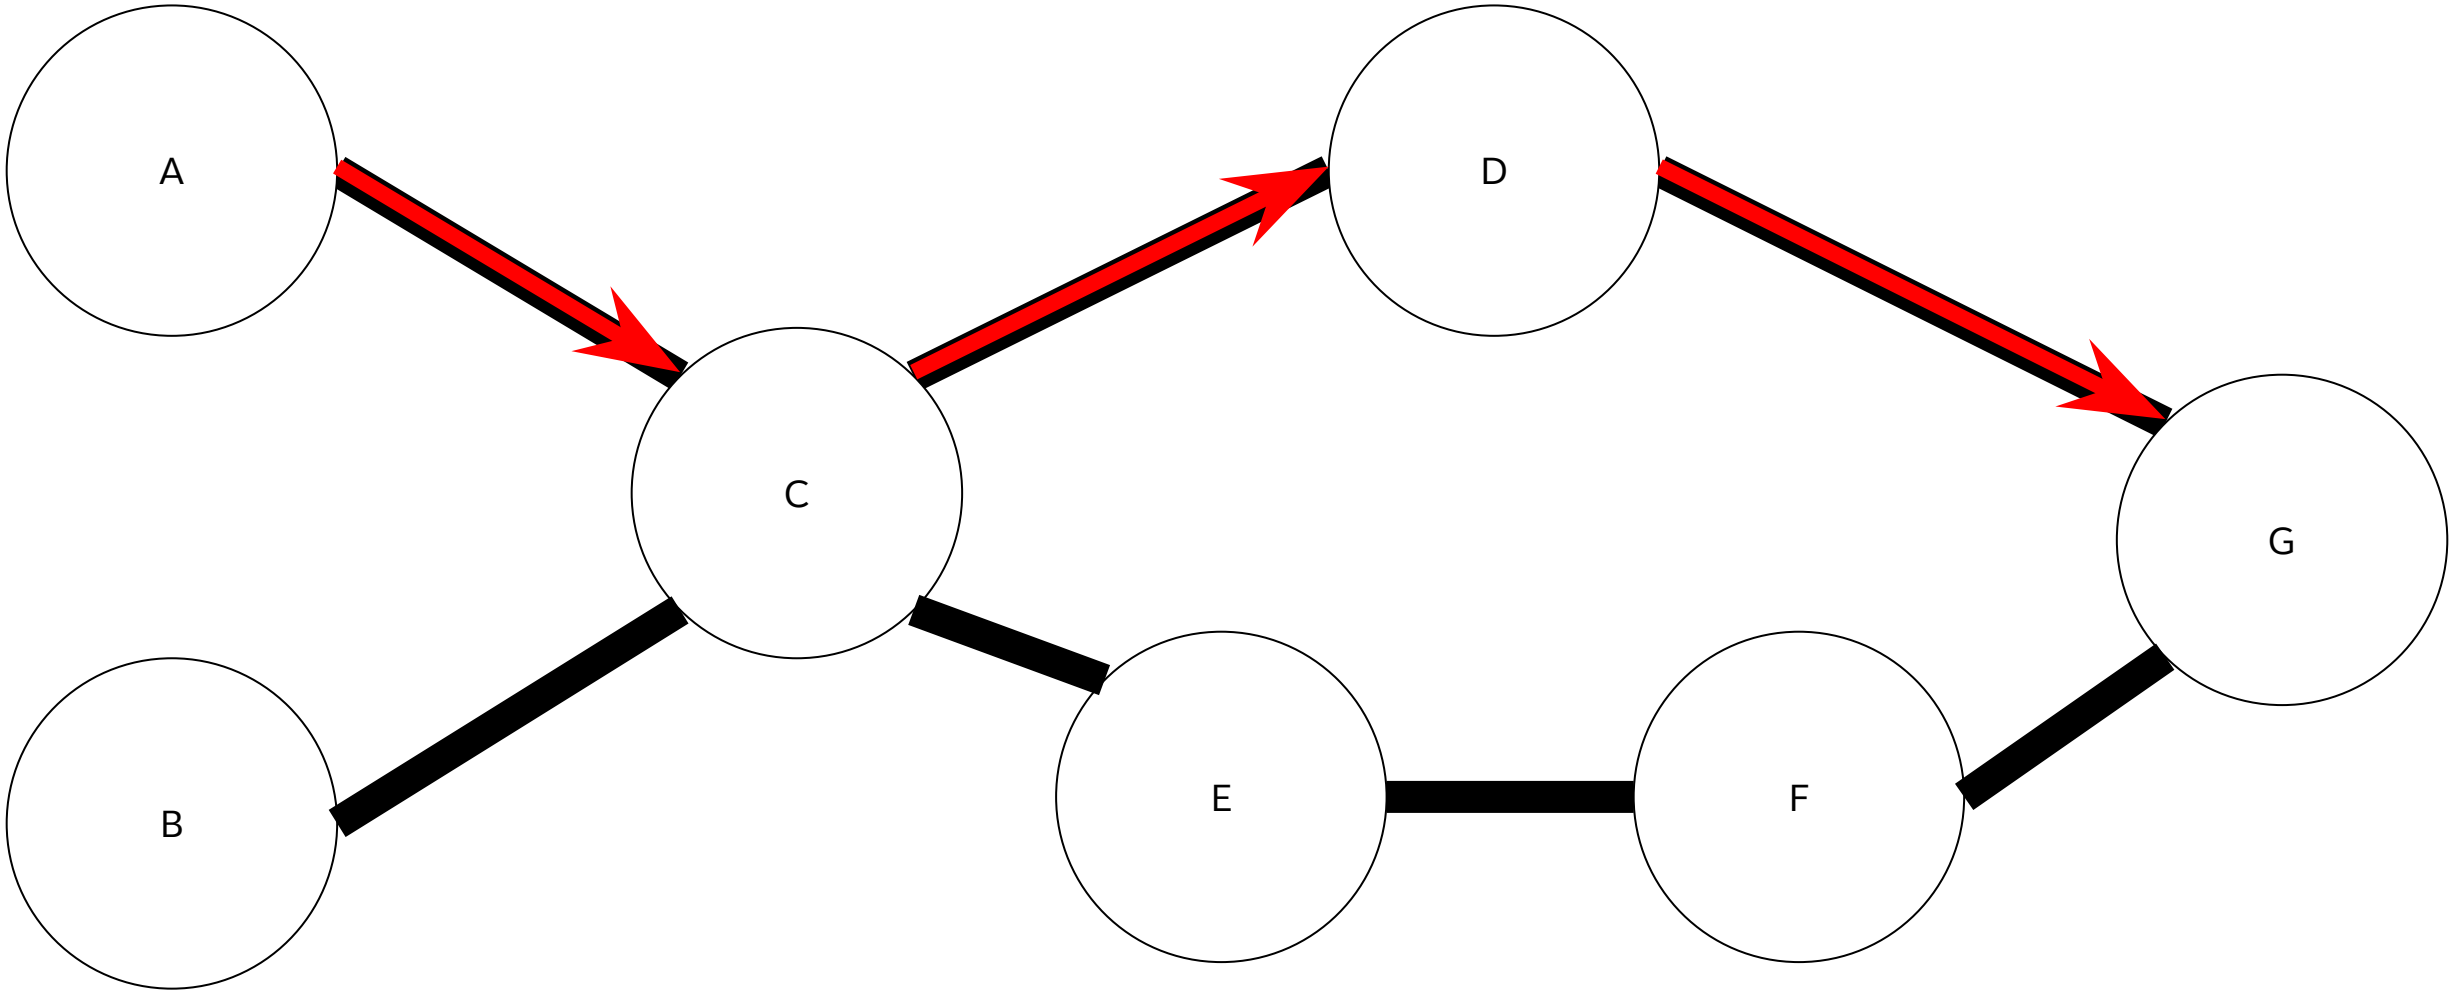

Graph Theory

slack
from Salesforce

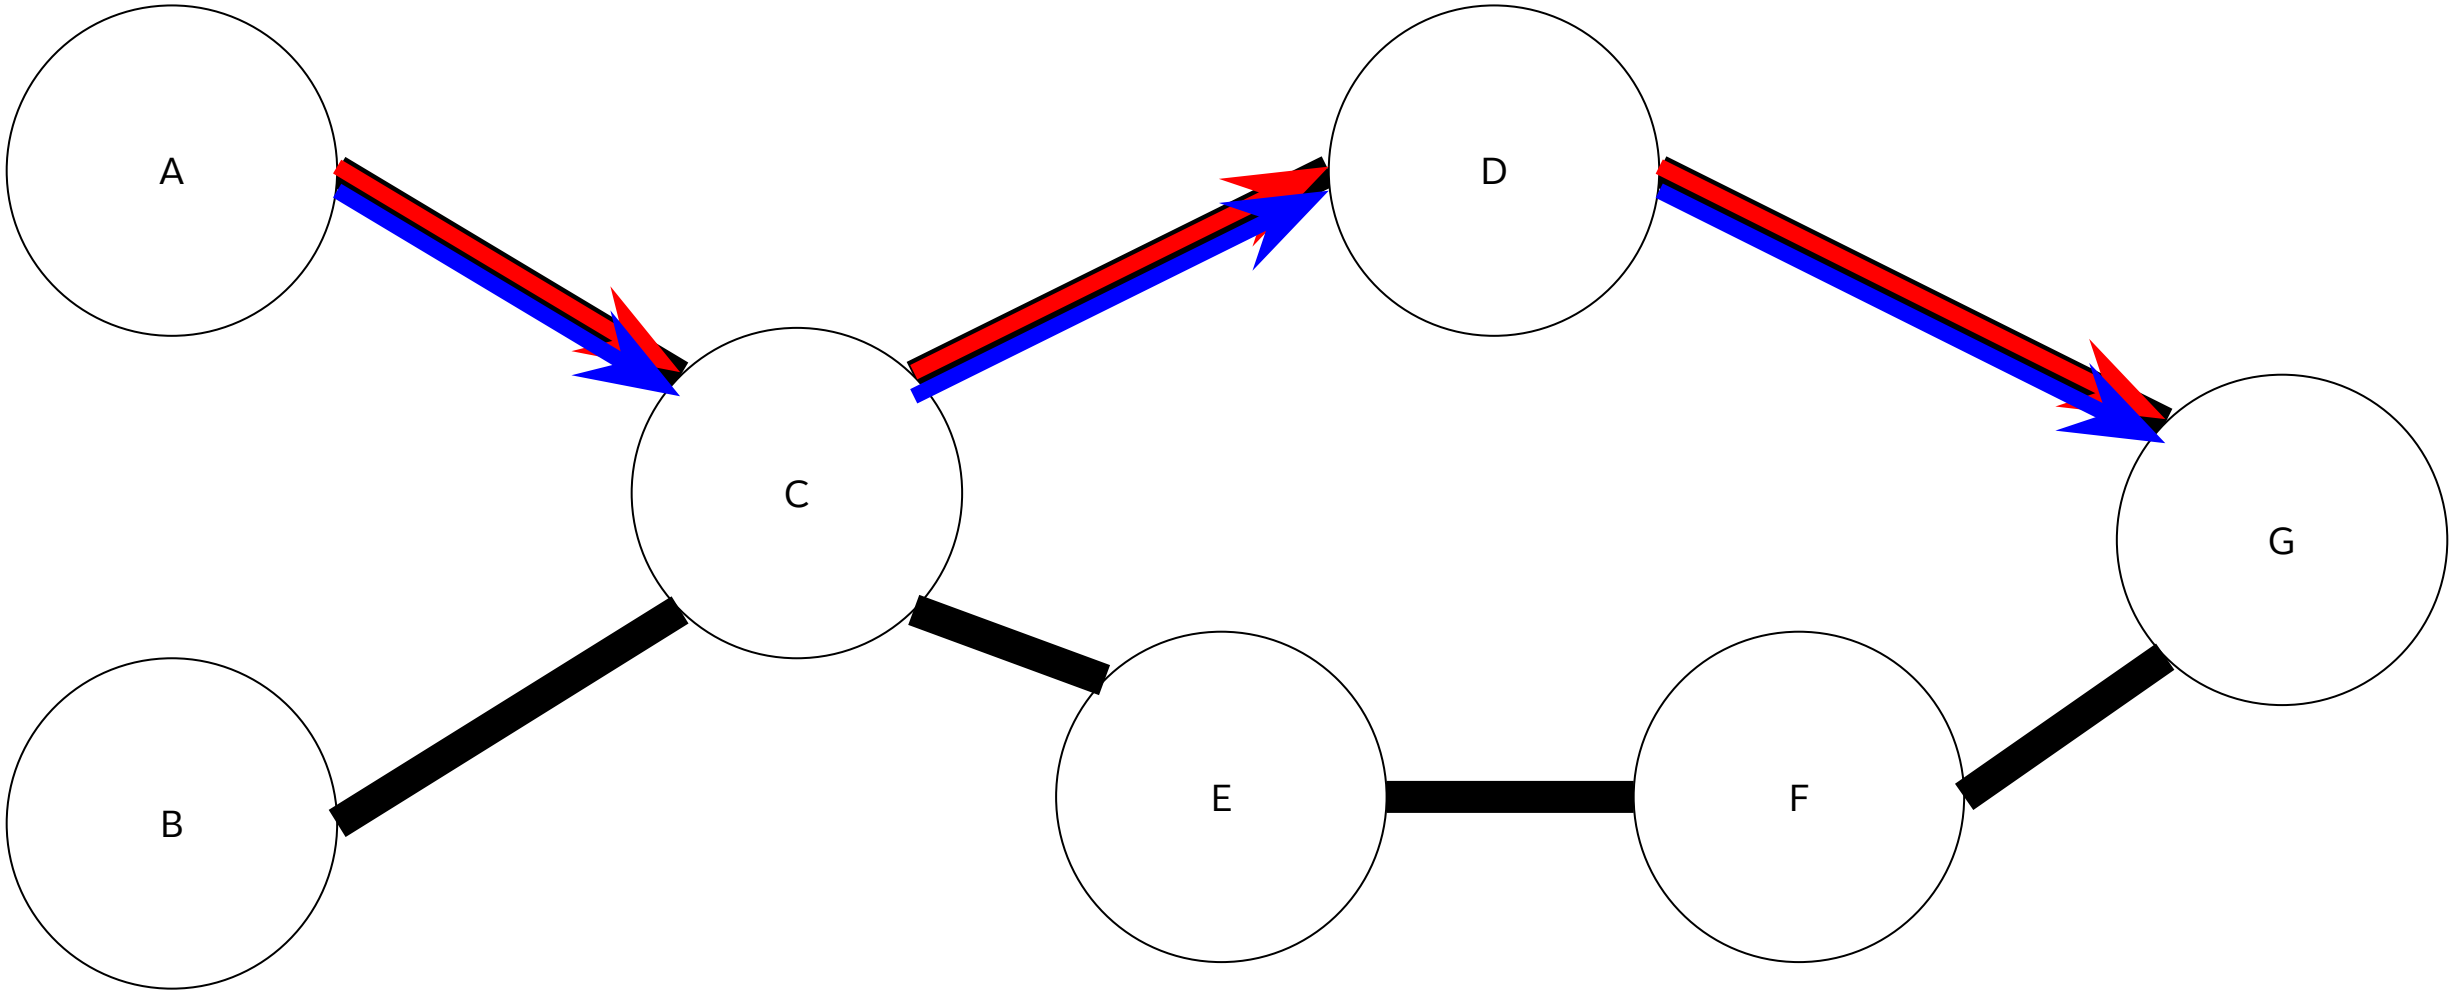

Graph Theory

slack
from Salesforce

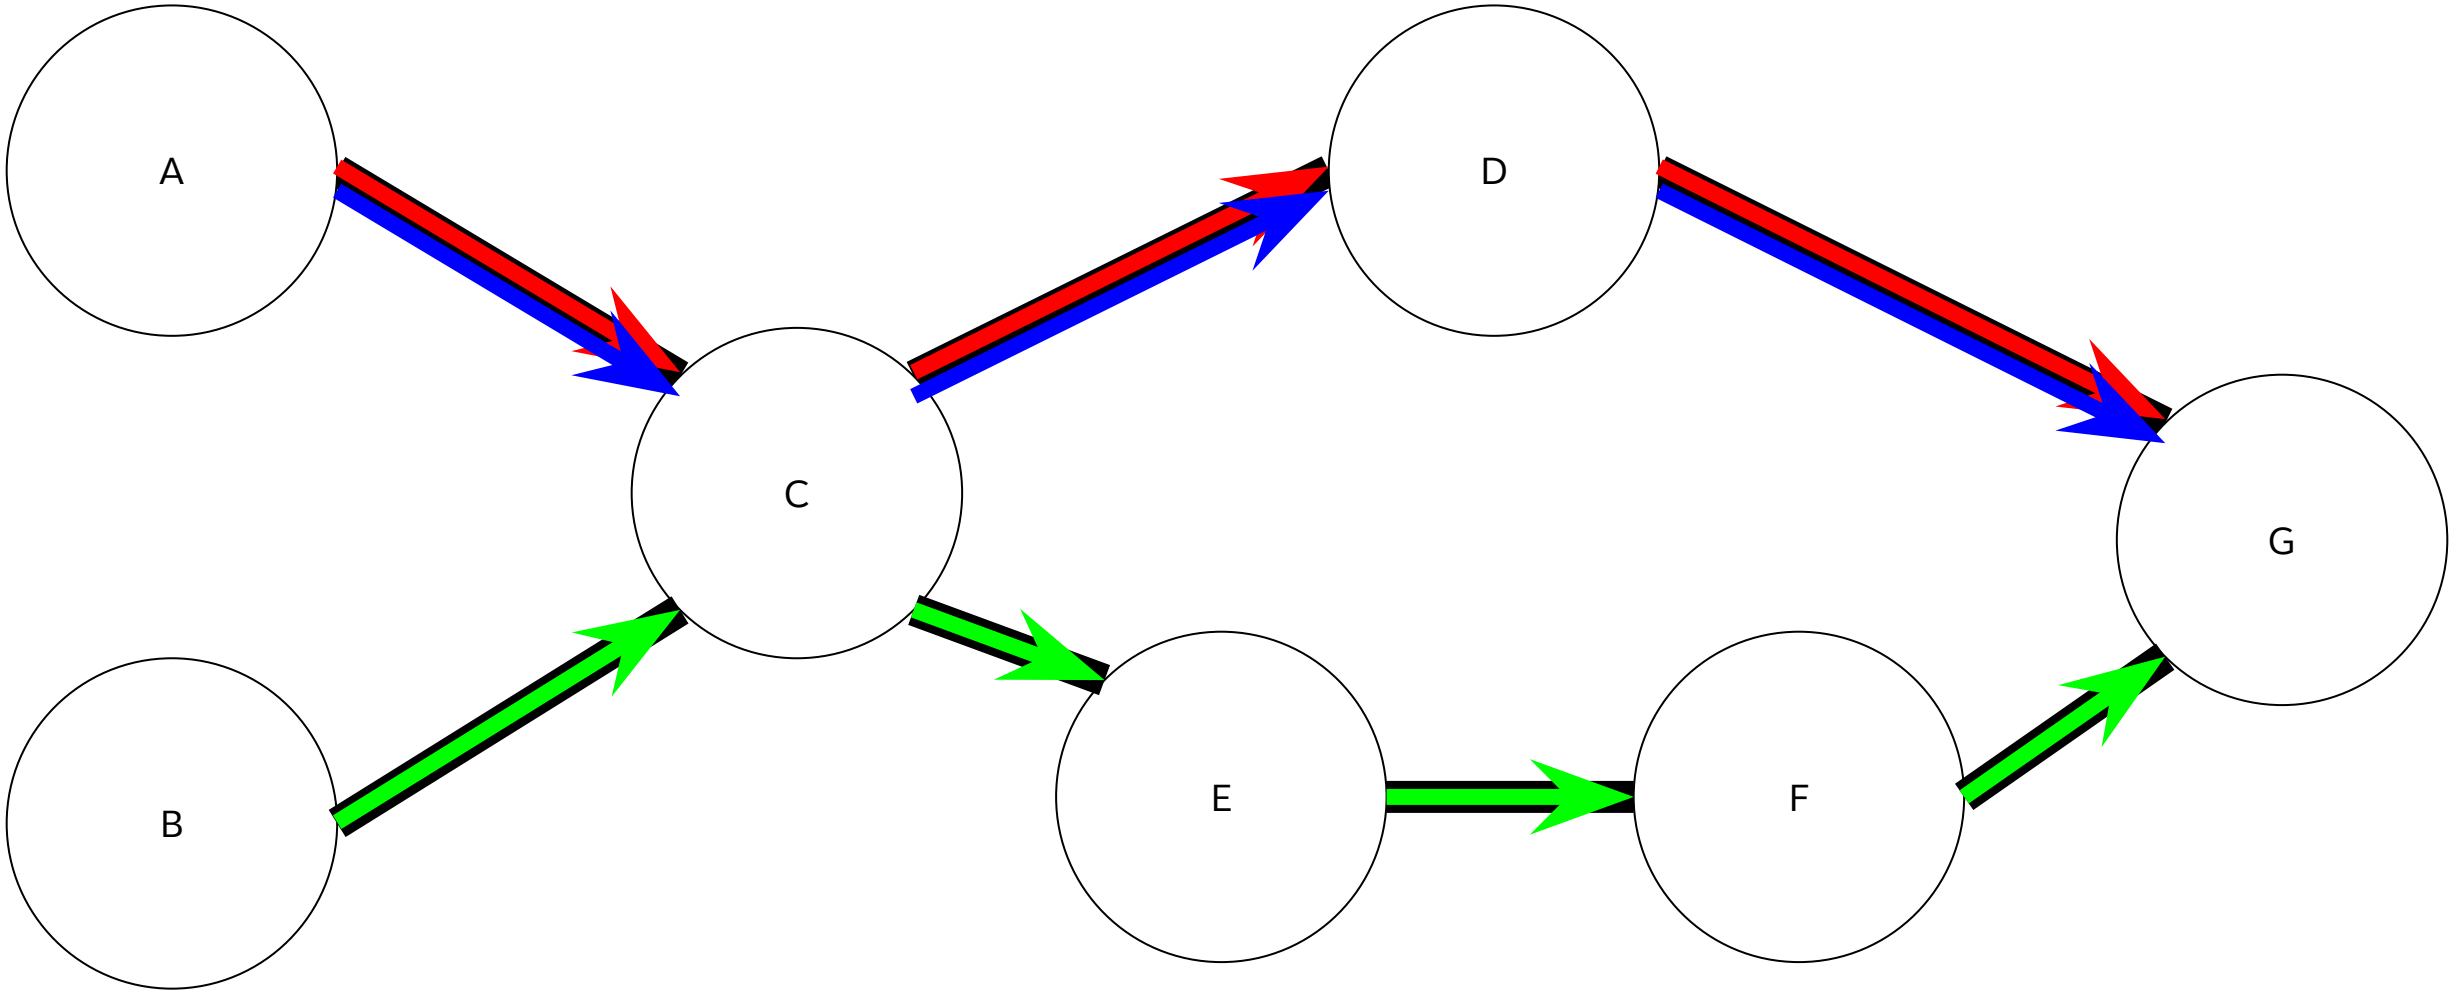

A more familiar graph

slack
from Salesforce

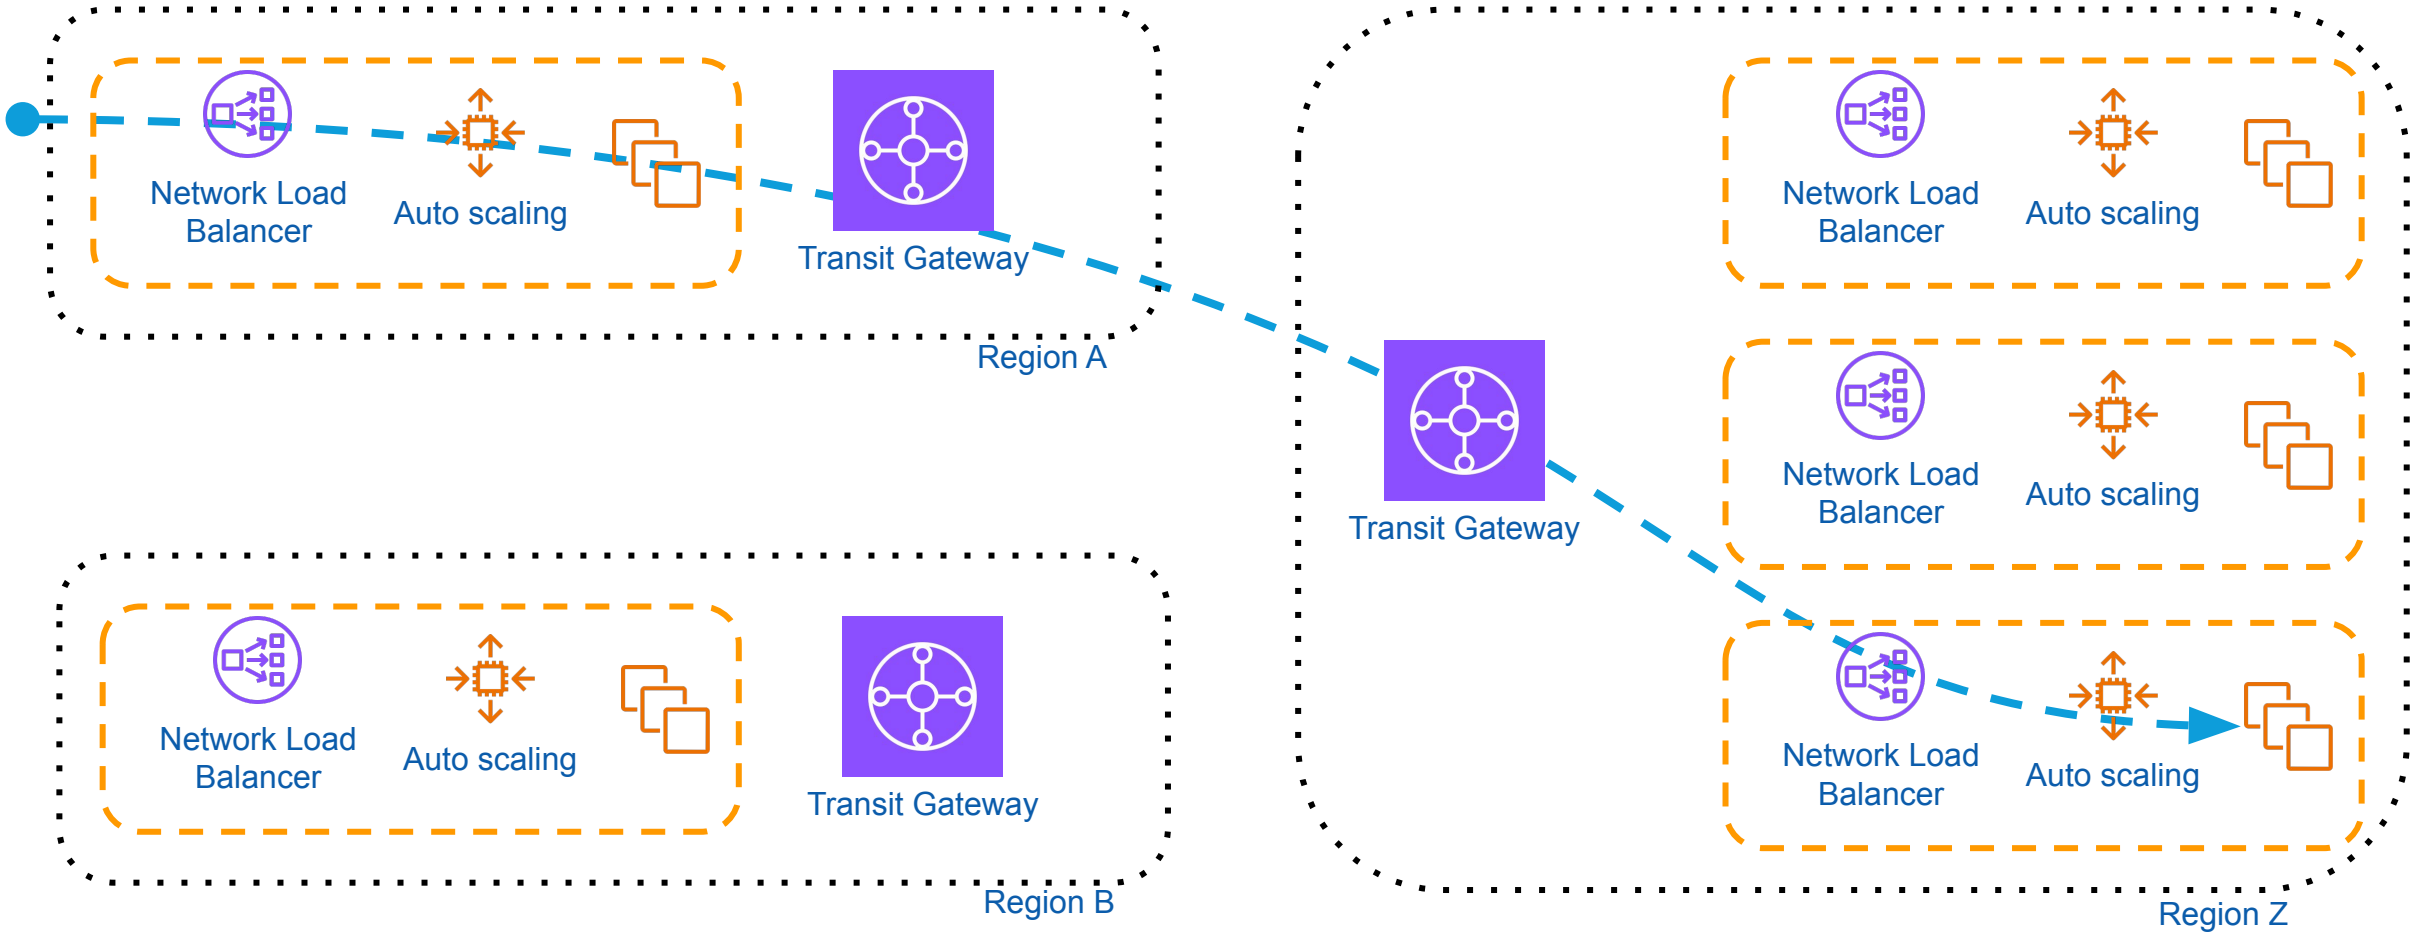

Converted

slack
from Salesforce

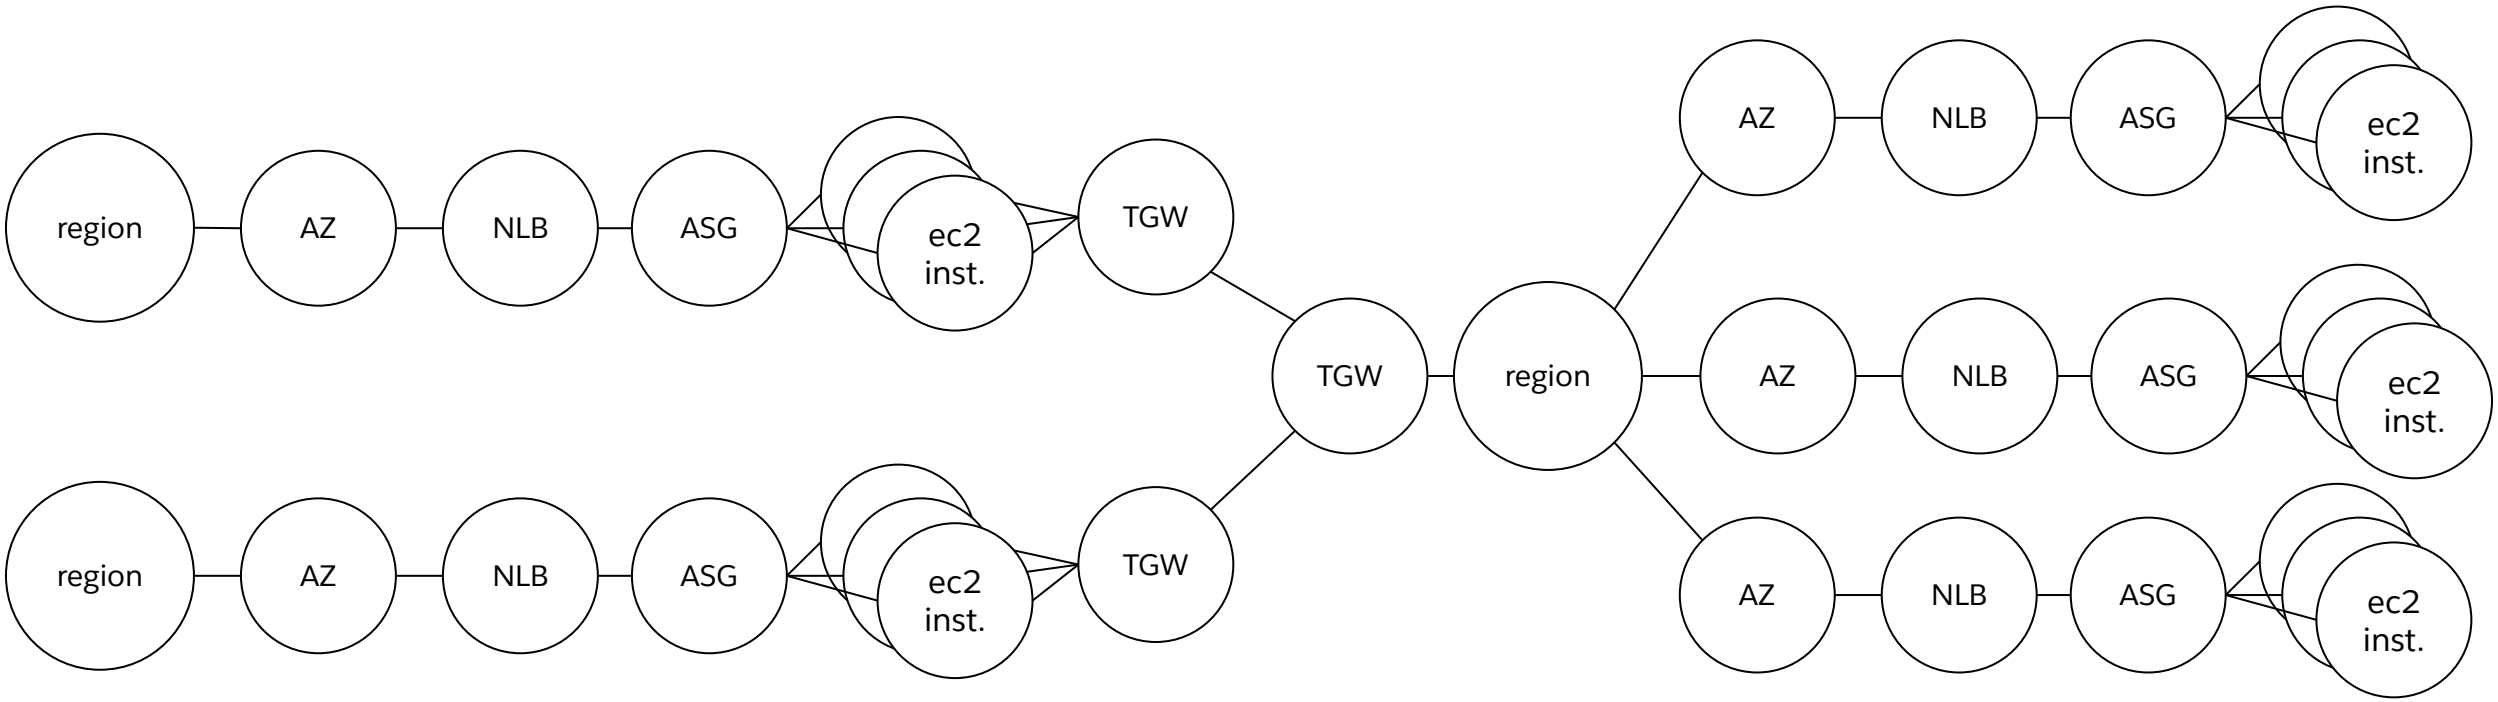

Implementation

slack
from Salesforce

Questions to answer

- What are the traffic flows?
- What is the topology of our infrastructure?
- Can the *capacity* of the topology satisfy the flows?

Implementation

sFlow

[RFC 3176]

- Use traditional network 5-tuple to identify flows. Consists of:
  - src_addr
  - dst_addr
  - src_port
  - dst_port
  - protocol
- Also tracks things like TCP flags, QOS
- Add counters to aggregate number of packet count and total payload size
- Solely based off header data! *(no payload, no PII, security teams like this)*

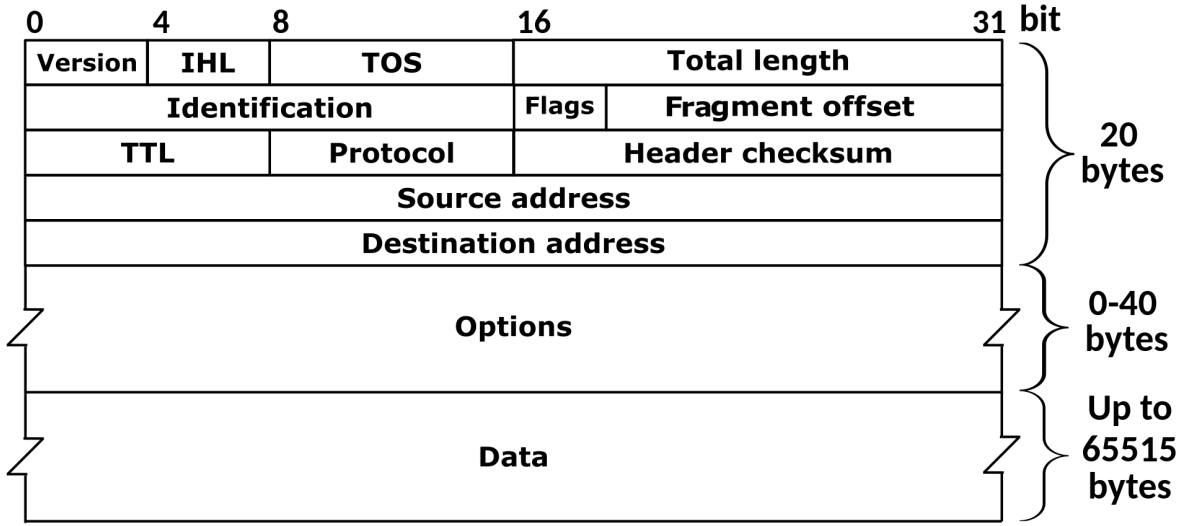


```
/* Packet IP version 4 data */

struct sampled_ipv4 {
   unsigned int length;     /* The length of the IP packet excluding
                               lower layer encapsulations */
   unsigned int protocol;   /* IP Protocol type
                               (for example, TCP = 6, UDP = 17) */
   ip_v4 src_ip;            /* Source IP Address */
   ip_v4 dst_ip;            /* Destination IP Address */
   unsigned int src_port;   /* TCP/UDP source port number or
                               equivalent */
   unsigned int dst_port;   /* TCP/UDP destination port number or
                               equivalent */
   unsigned int tcp_flags;  /* TCP flags */
   unsigned int tos;        /* IP type of service */
}
```

Implementation

Flow Collection

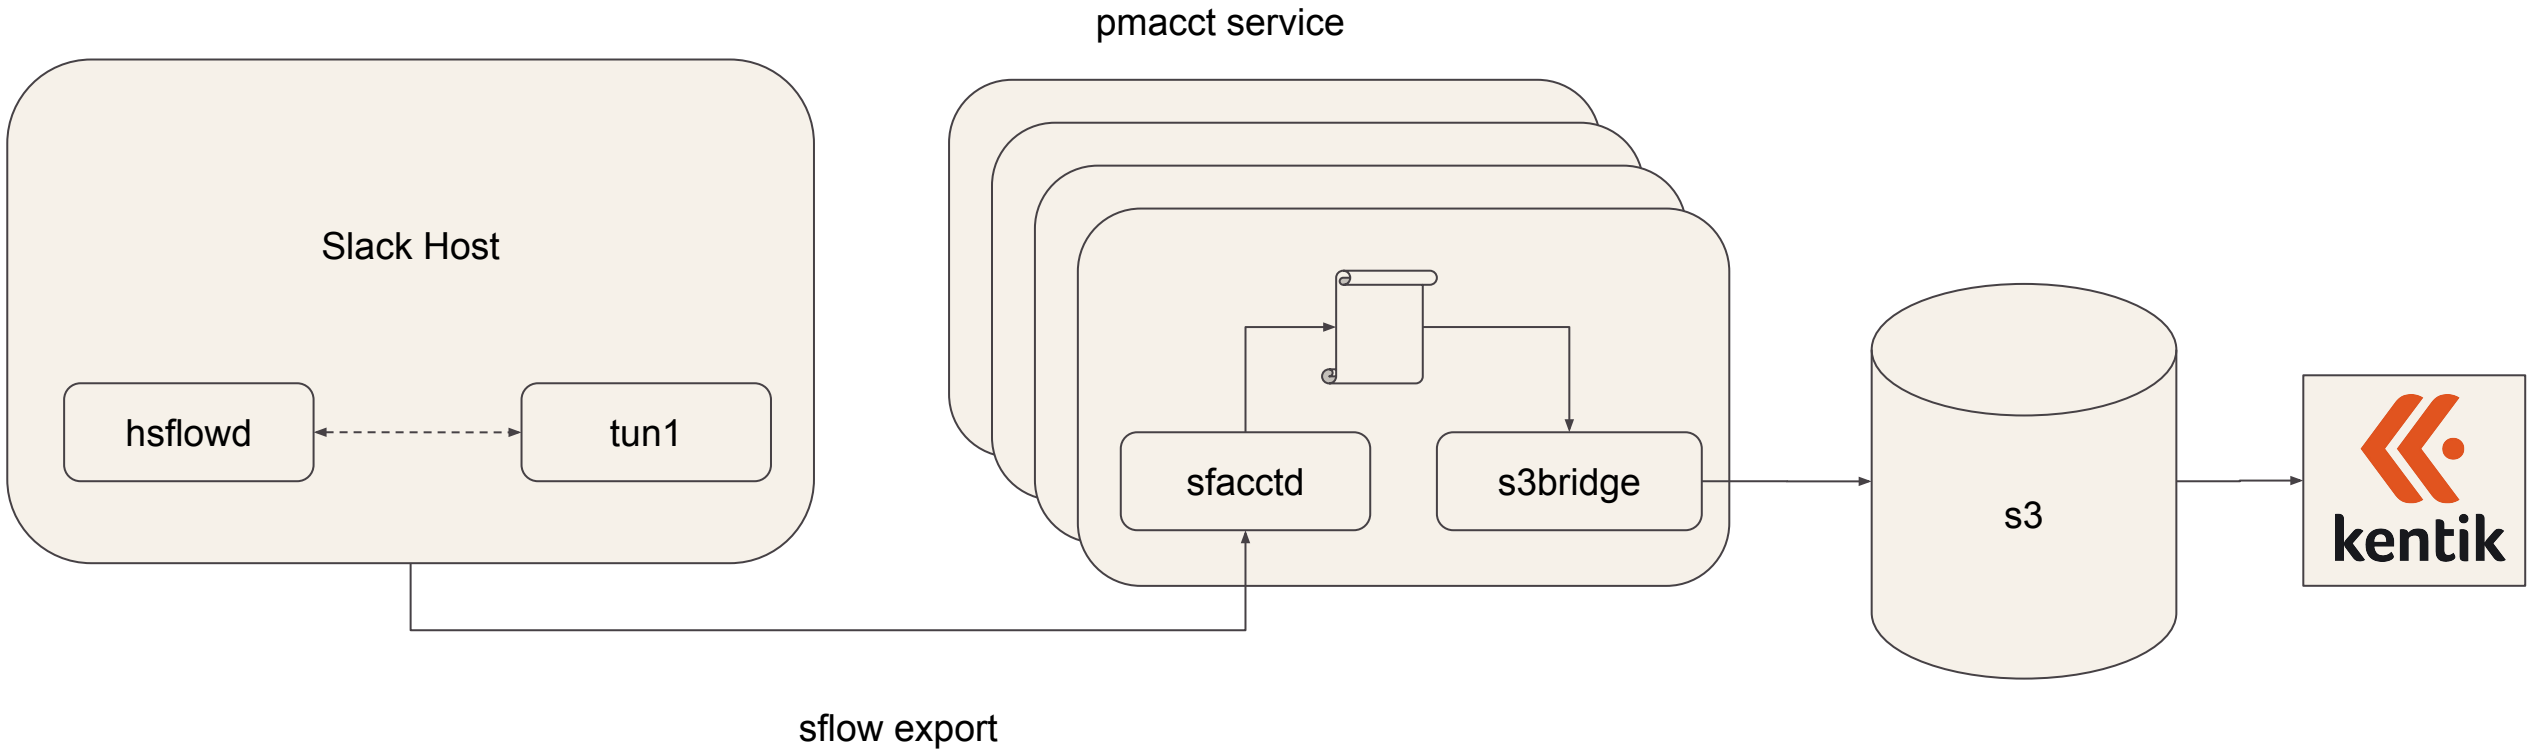

Implementation

Topology

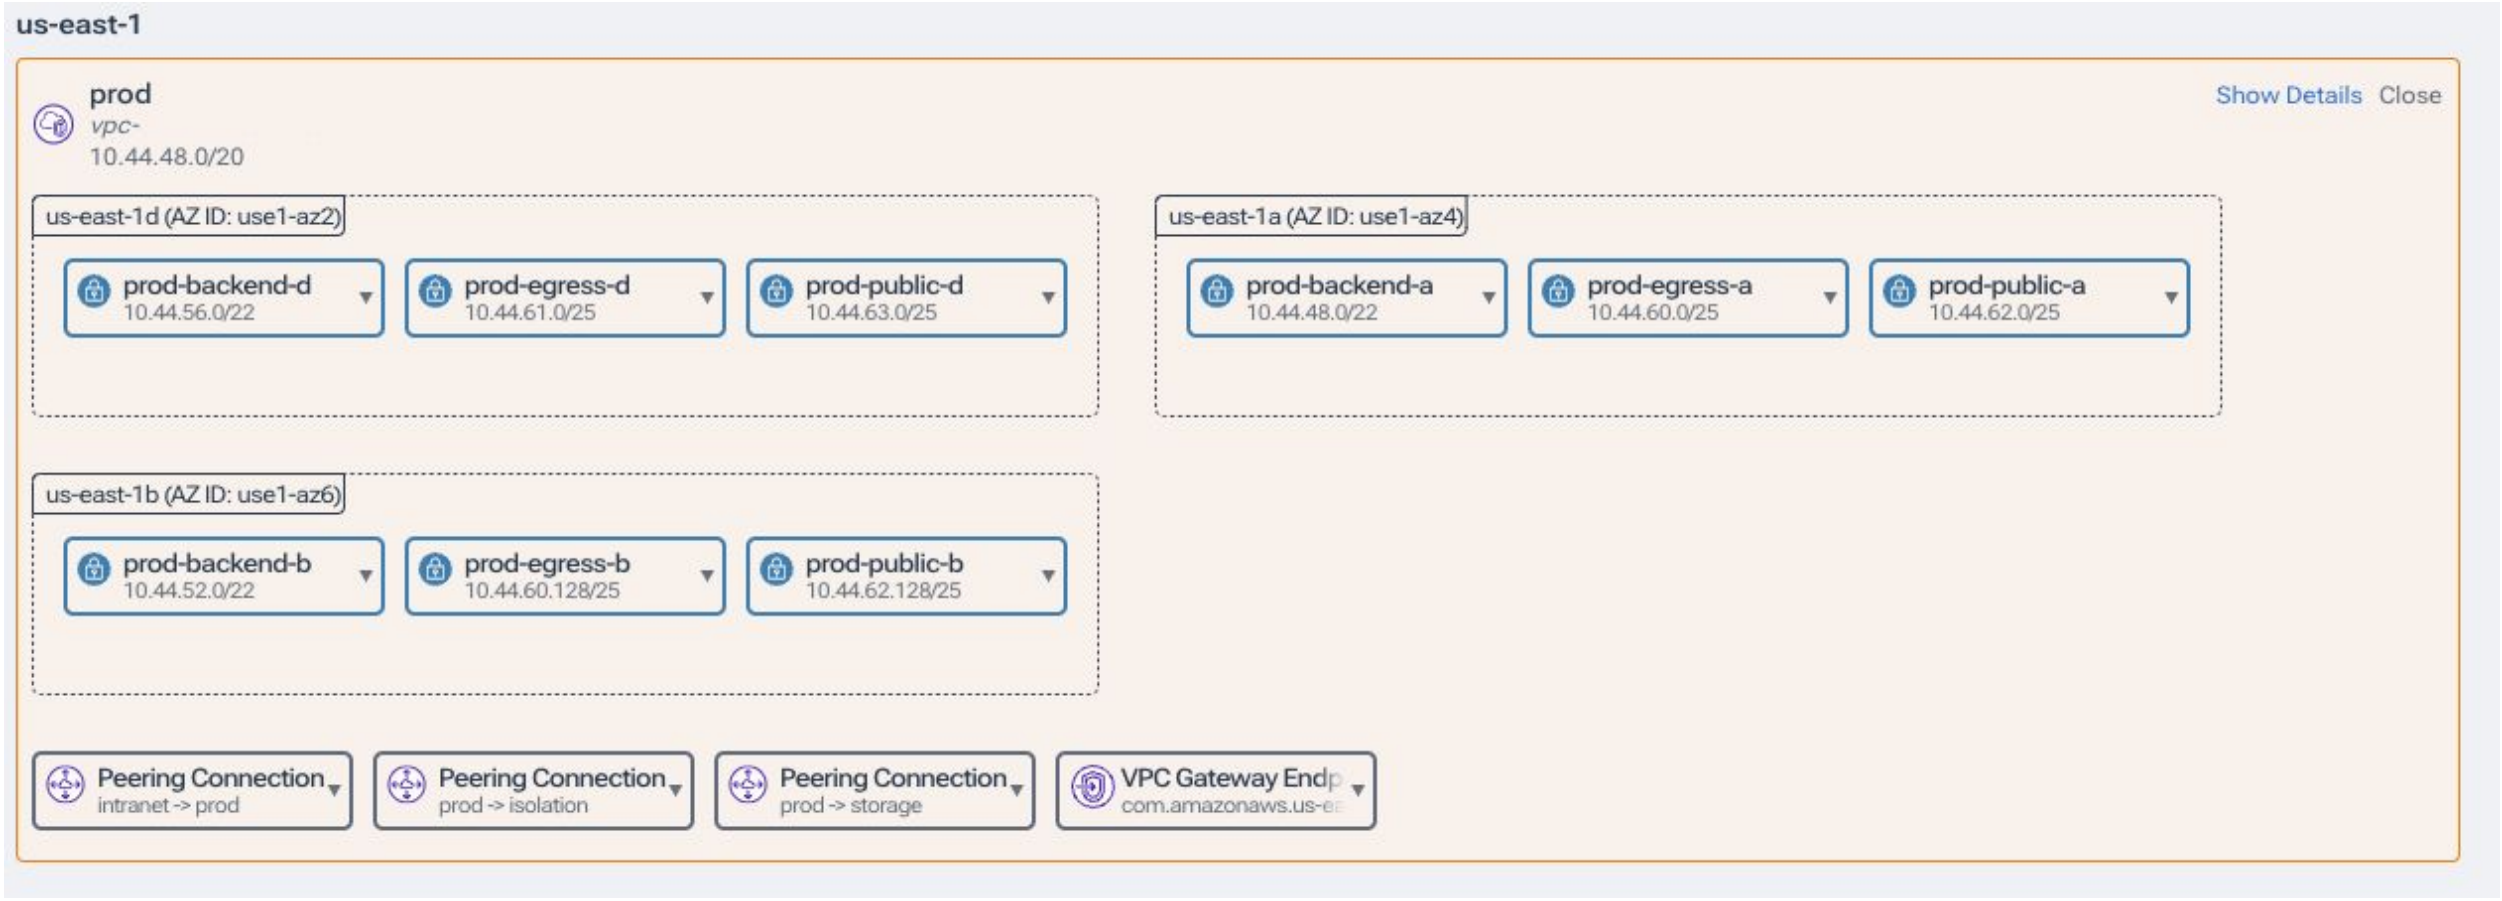

Why not vended logs?

Compare/Contrast

Vendor Solution

- Coverage for network primitives (TGWs, NAT GWs, etc.)
- Only has insight into underlay traffic headers
- Cost (💰)

Home grown

- Only coverage for compute hosts
  - Slack runs dedicated ingress/egress proxy stacks, so this is ok
- Can view Nebula overlay header information
- Flexible parsing and annotation
- 3 orders of magnitude cheaper!

Use Cases

Cost Showbacks

Shared Infra Costs – Service C is fronted by an NLB, Service A sends 2x traffic as Service B

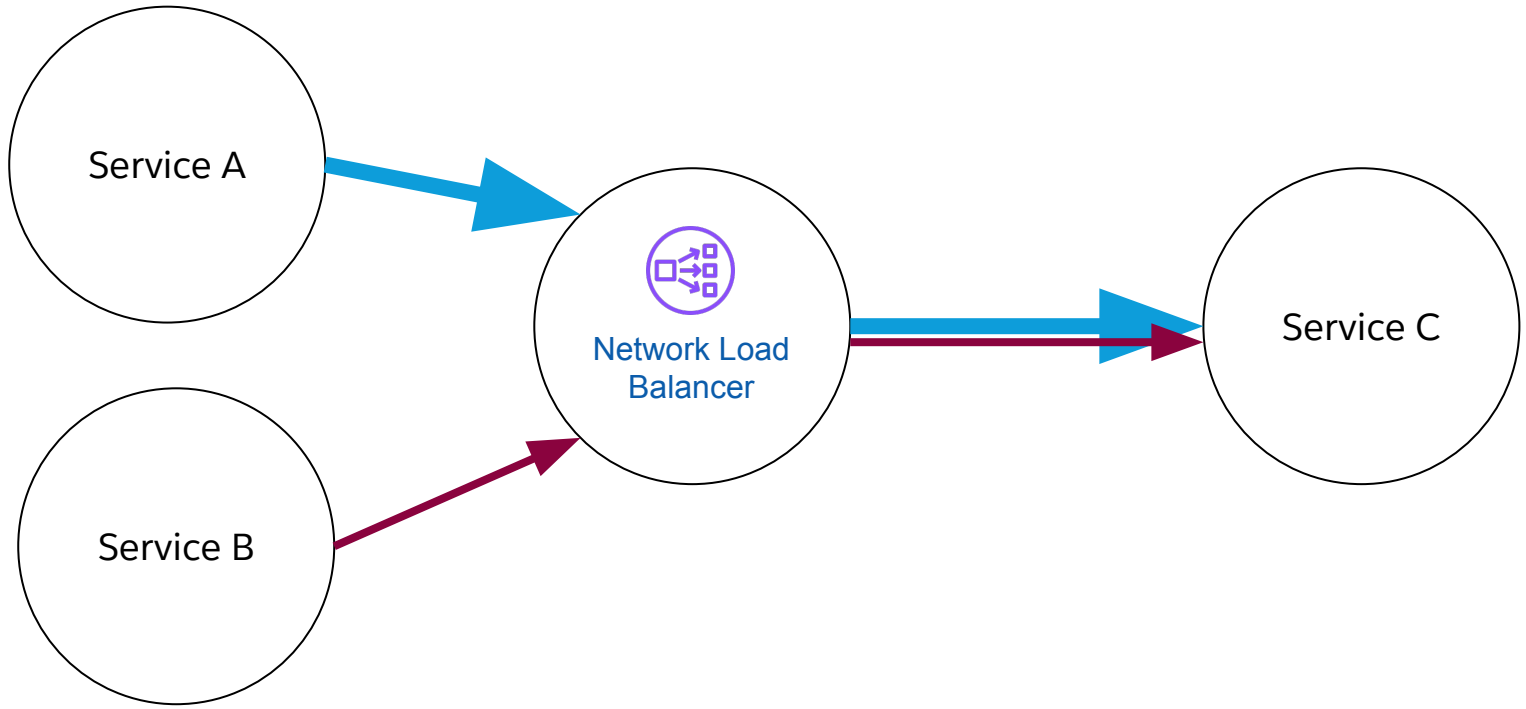

Use Cases

Cost Showbacks

Intra vs Inter AZ traffic affinity – What is the locality of service to service traffic flows?

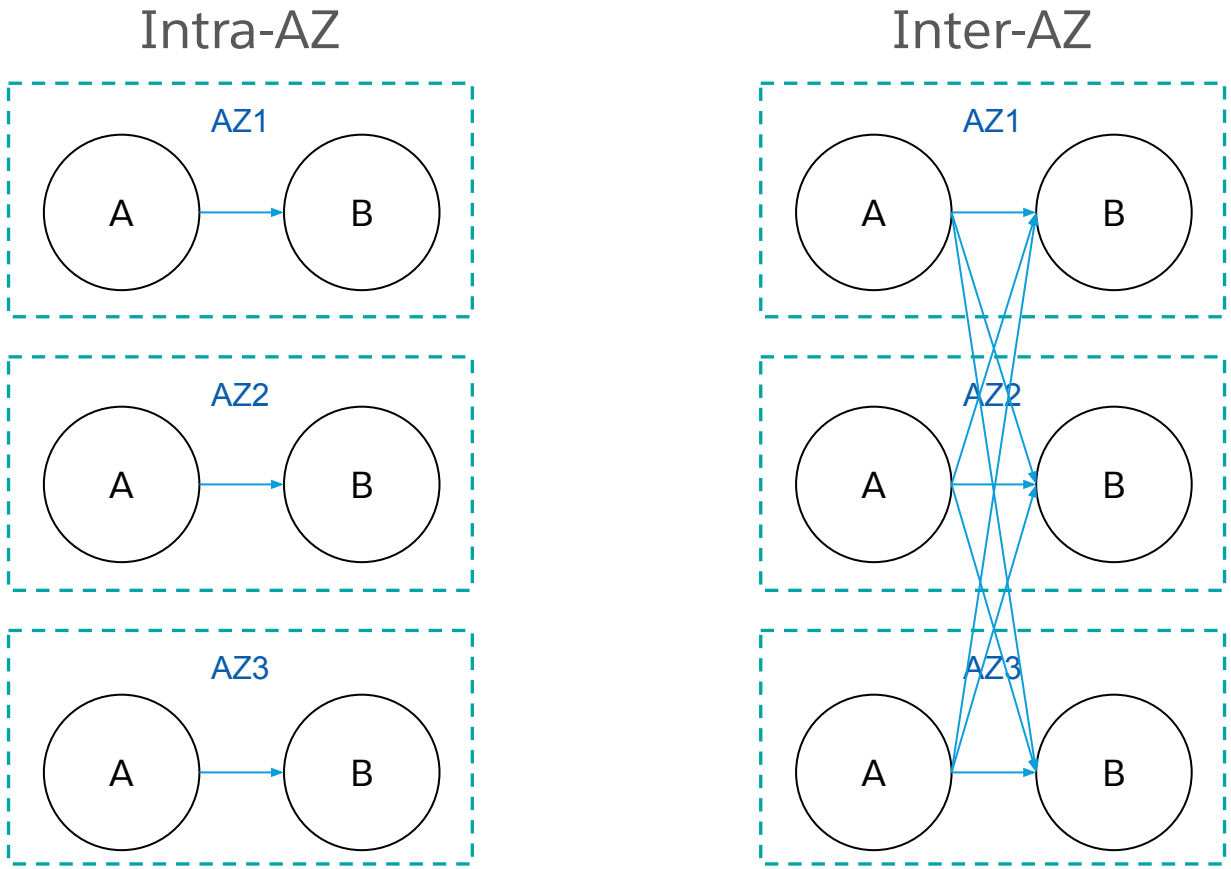

Use Cases

Capacity Planning

VPC Peering to TGW migration

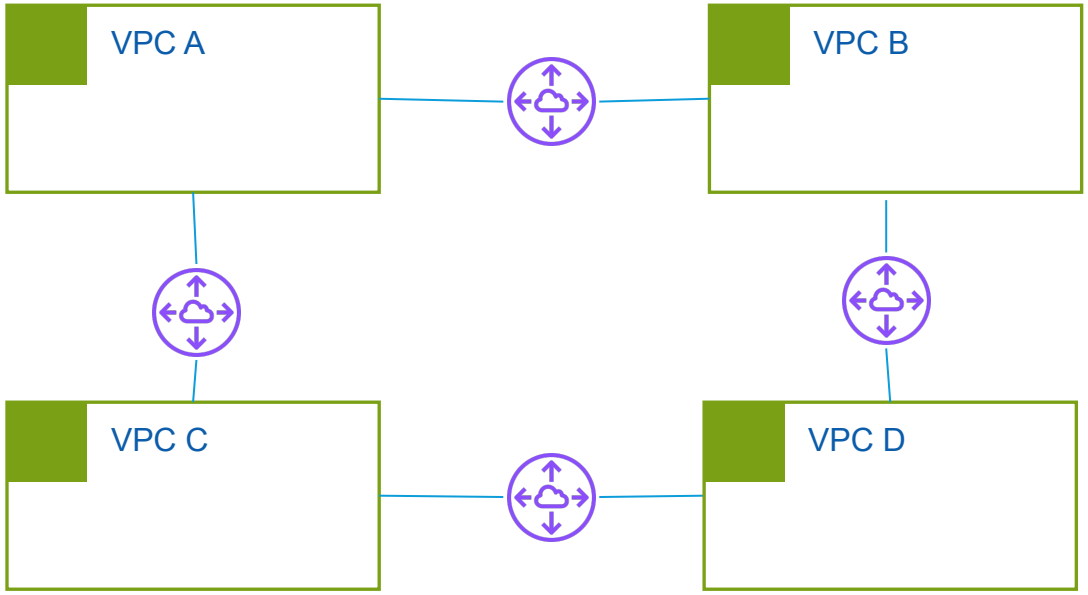

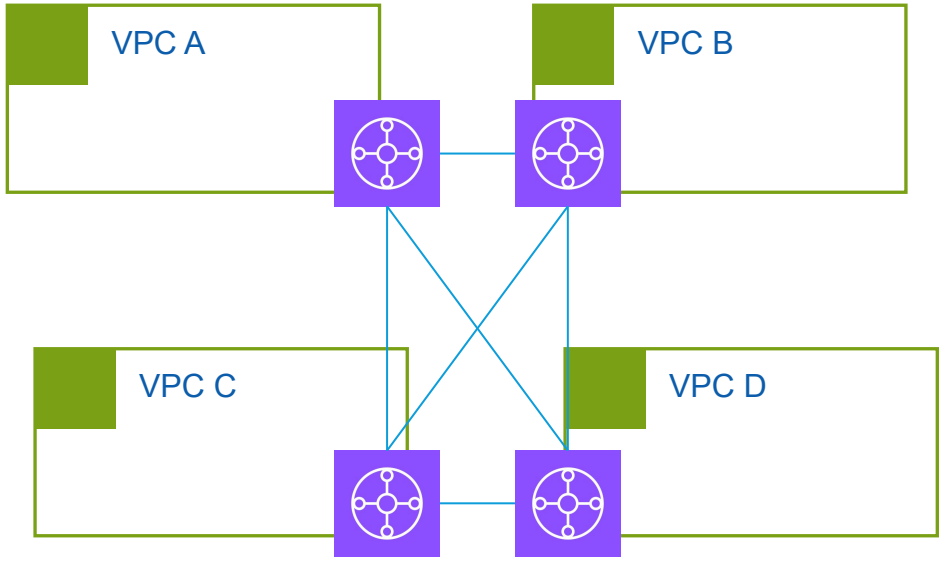

Use Cases

Capacity Planning

Node/link failure analysis

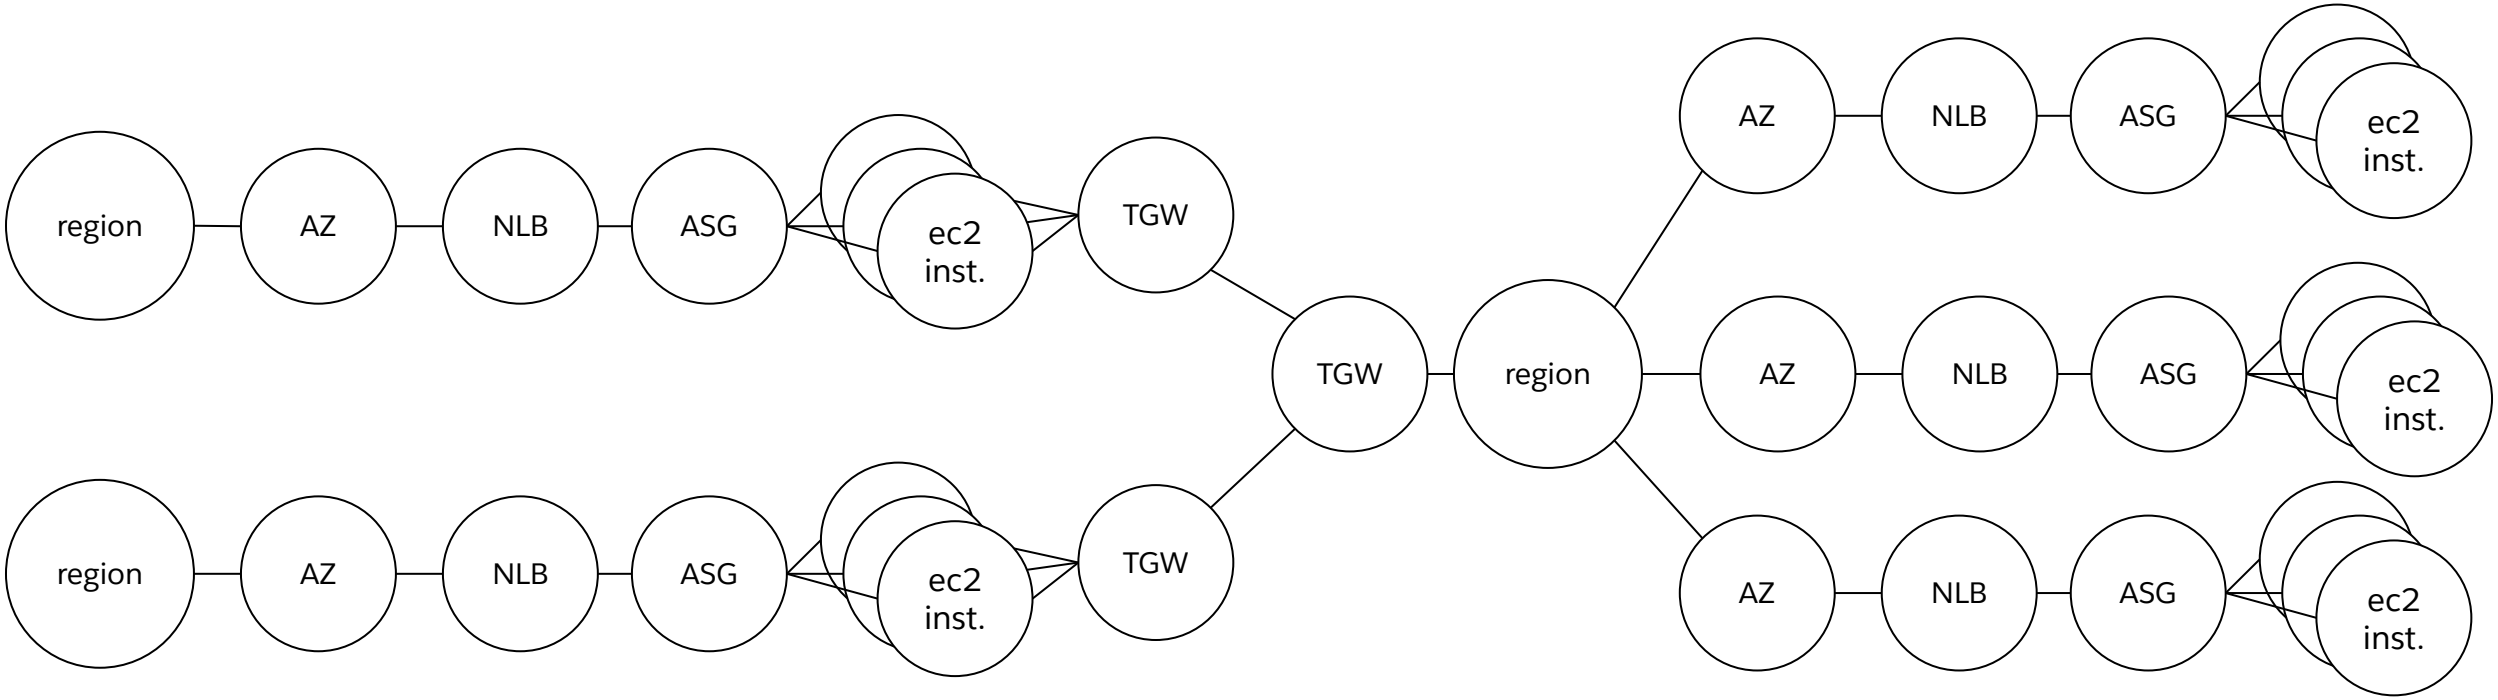

Use Cases

Capacity Planning

Node/link failure analysis

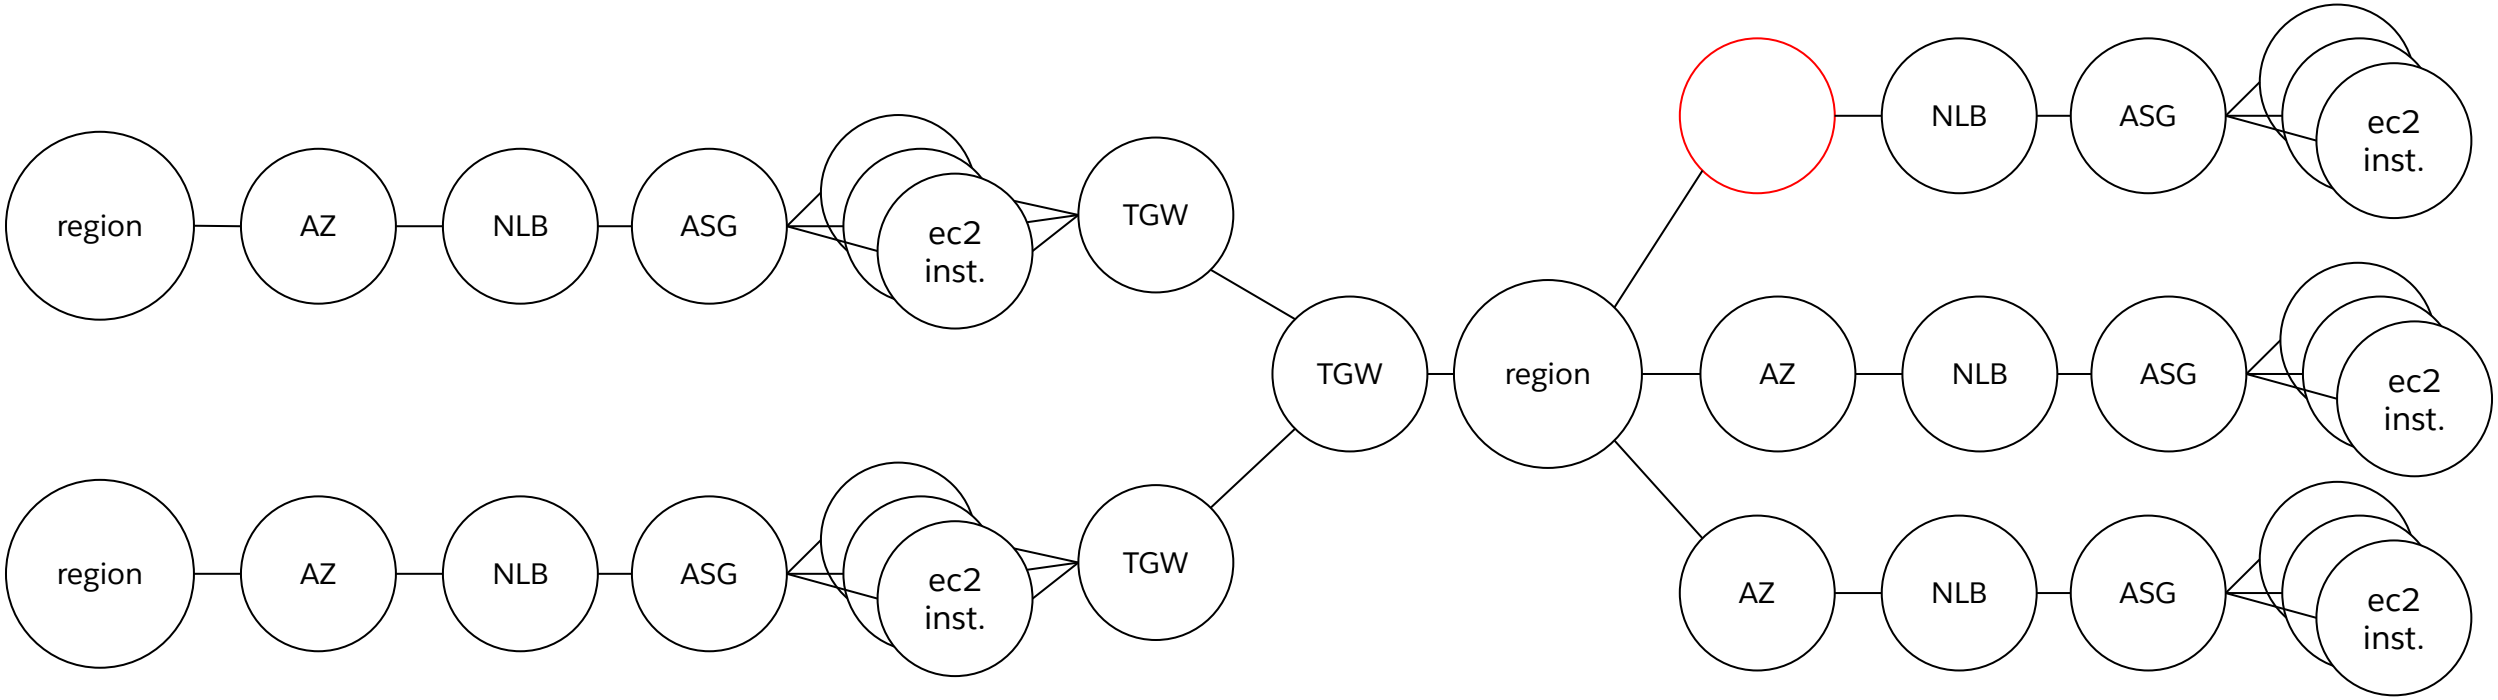

Recap

slack
from Salesforce

Enabling network flow data

- Generate flow logs
- Model network topology
- Evaluate traffic loads with forwarding rules
- Make informed decisions

Q & A

slack
from Salesforce