

#RP30

Automated Deployment and Scaling of Named Data Networks in Cloud Environments

Sean Liao
Supervisor: Zhiming Zhao

Next 20 mins...

Some academic(s) arrive to tell us
that (once again) they have Fixed the Internet,
and (once again) it runs on top of the current actually-working
internet,
and (once again) if you sign up you can communicate with as
many as twelve other computers.

n-gate.com *in reference to SCION*


Introduction

ENVRI-FAIR

ENVironmental Research
Infrastructures (ENVRI)


connection to

European Open Science
Cloud (EOSC)




Single Domain				Multi Domain			
							
AnaEE	EPOS	EUROFLEETS	ACTRIS	EuroGOOS	AQUACOSM	DANUBIUS	EMSO
DISSCO		EURO-ARGO	ARISE	ICOS		DANUBIUS	
ELIXIR		JERICO-RI	EISCAT_3D	IS-ENES		eLTER	
EMPHASIS		SEADATANET	EUFAR	SIOS		EMBRC	
INTERACT			EUROCHAMP 2020			LifeWatch	
			HEMERA				
			IAGOS				


Data distribution



Content Distribution Network



Federated Clouds



Named Data Networking (NDN)

Named Data Networking


NSF Future Internet Architecture Program

Information distribution network

Potential benefits:

- content caching
- network level security of data

1.2.3.4/a/b/c/d



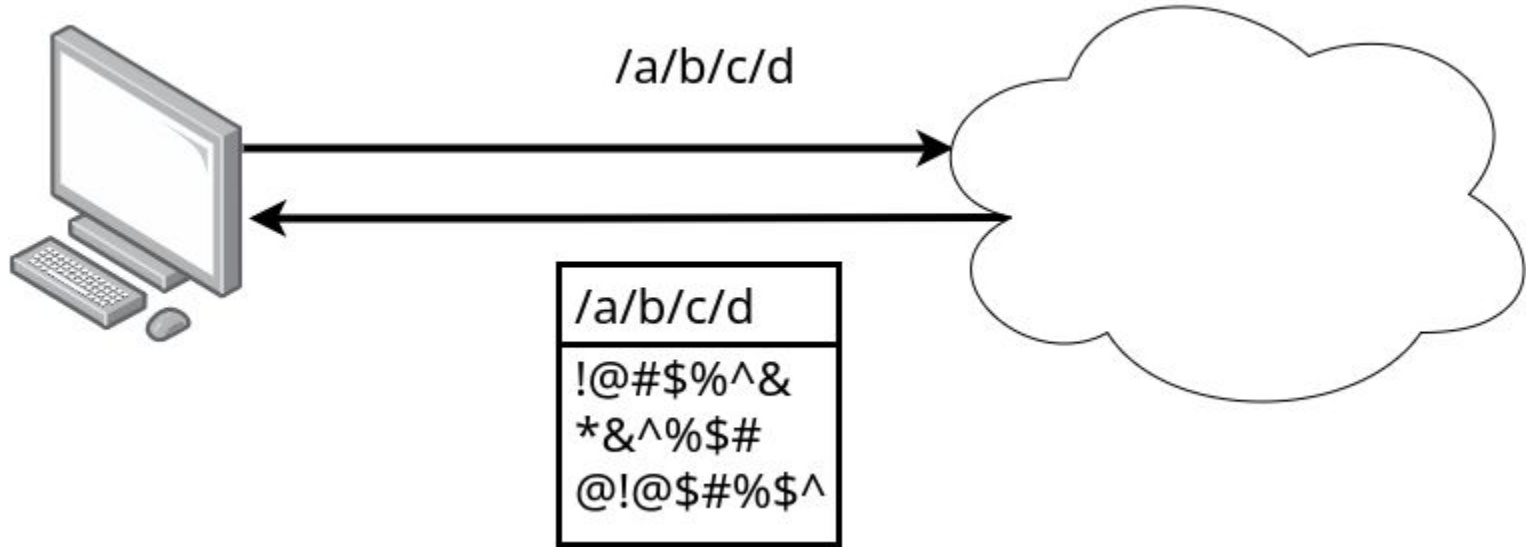
Route directly to data

IPv4: 1.2.3.4


IPv6: 1:2:3:4:5:6:7:8

NDN: /arbitrary/strings/infinite/address/space

1-1 Request-Response (Interest-Data)



In network caching



Run it

In the Cloud

Overlay over IP

Simplify deployment

Scalable

Existing Tools

Router: NFD


Link State Routing: NLSR

Dynamic route updates on static network

Connectivity: FCH

Find closest hub/gateway


Overlay over TCP / UDP



Processing Overhead

Hop Processing Time

Subtract network roundtrip time: 0.174ms





Testing notes

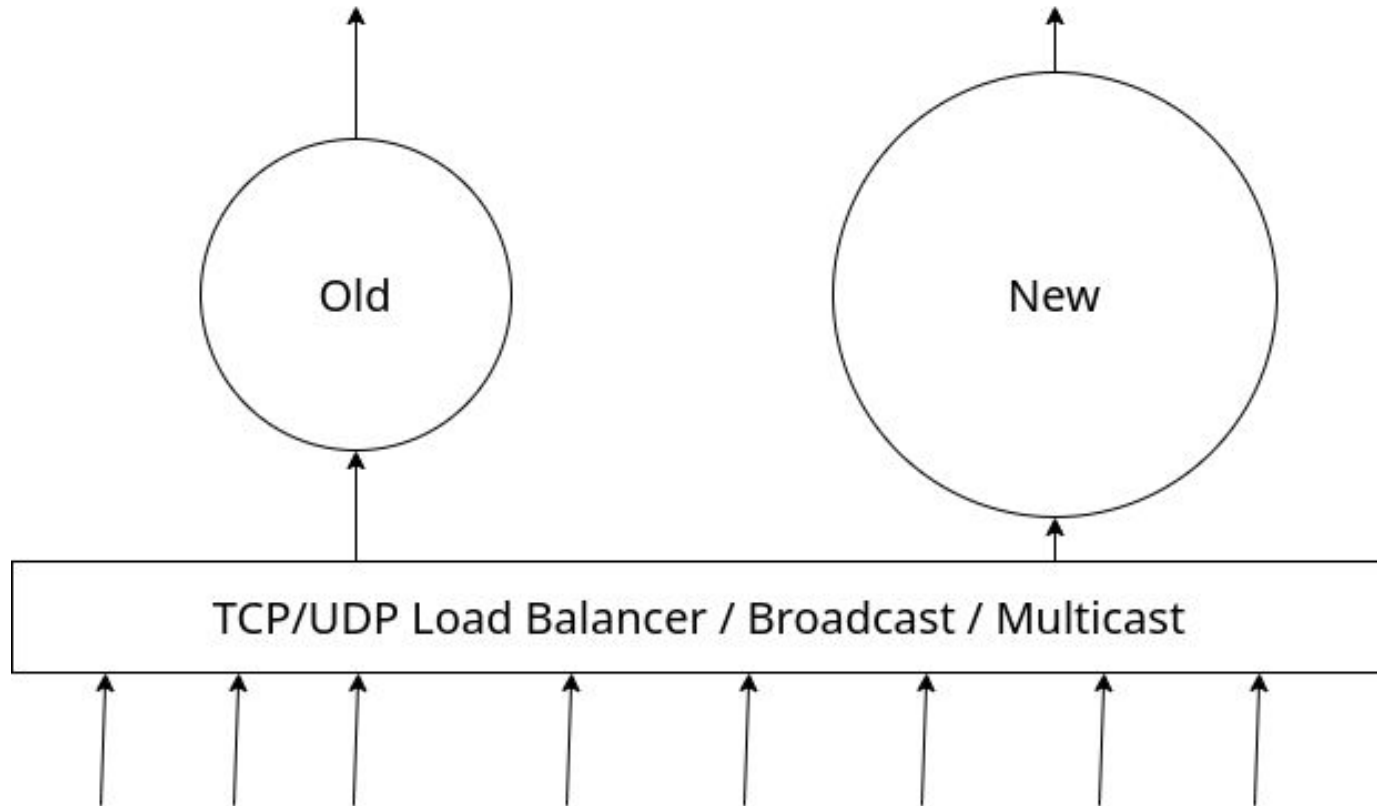
Performance severely degrades with 50000+ cached / in-flight requests

Scaling Up

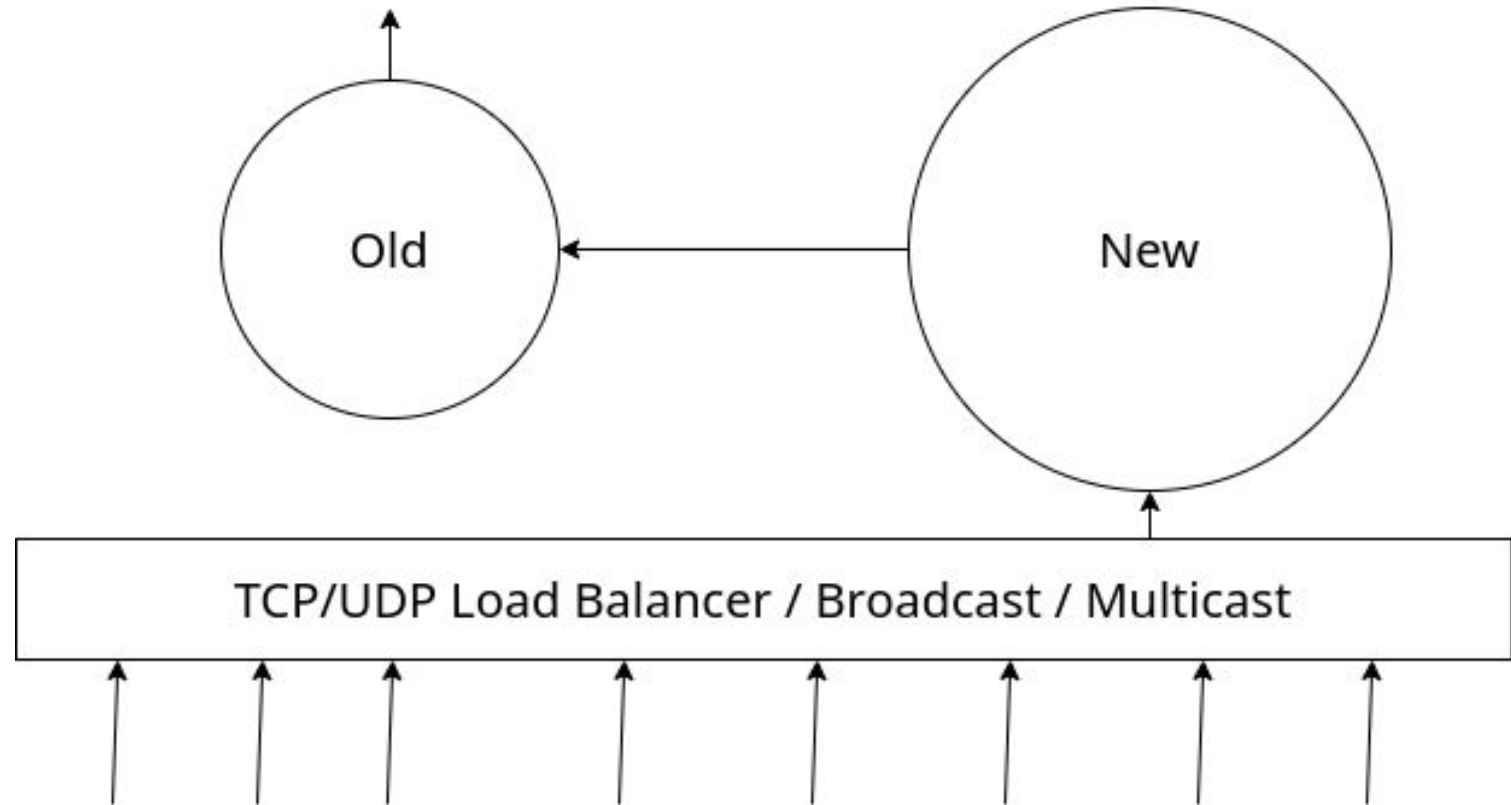
Problem: Growing a node




Replace




New Node in Front



New Node Behind (Chain)




New Node Behind (Load Balancing)




Growing the Network


Problem: New Node




Broadcast / Multicast



Bootstrap Gossip




Central Discovery Server




Network Architecture

3 Layer Cache




Mesh Network




Automation

In-band Management



Out of Band Management



Proof of Concept

Necessary Configuration

Discovery Server

- Write down its address / Give it a preconfigured address


Load Balancer

- Address of Discovery Server
- Write down its address / Give it a preconfigured address

Caching Server

- Address of Load Balancer

Roundtrip Time



Proof of Concept Performance

Scraping state of router through CLI (on a timer)

TCP connections propagate routes and updates

Coarse grained partitioning of routes

Conclusion

Does it work?

Reuse existing router in a load balancing configuration

Minimal configuration, self connecting network

Lessons Learned

Now

Rapidly evolving research testbed
for new ideas

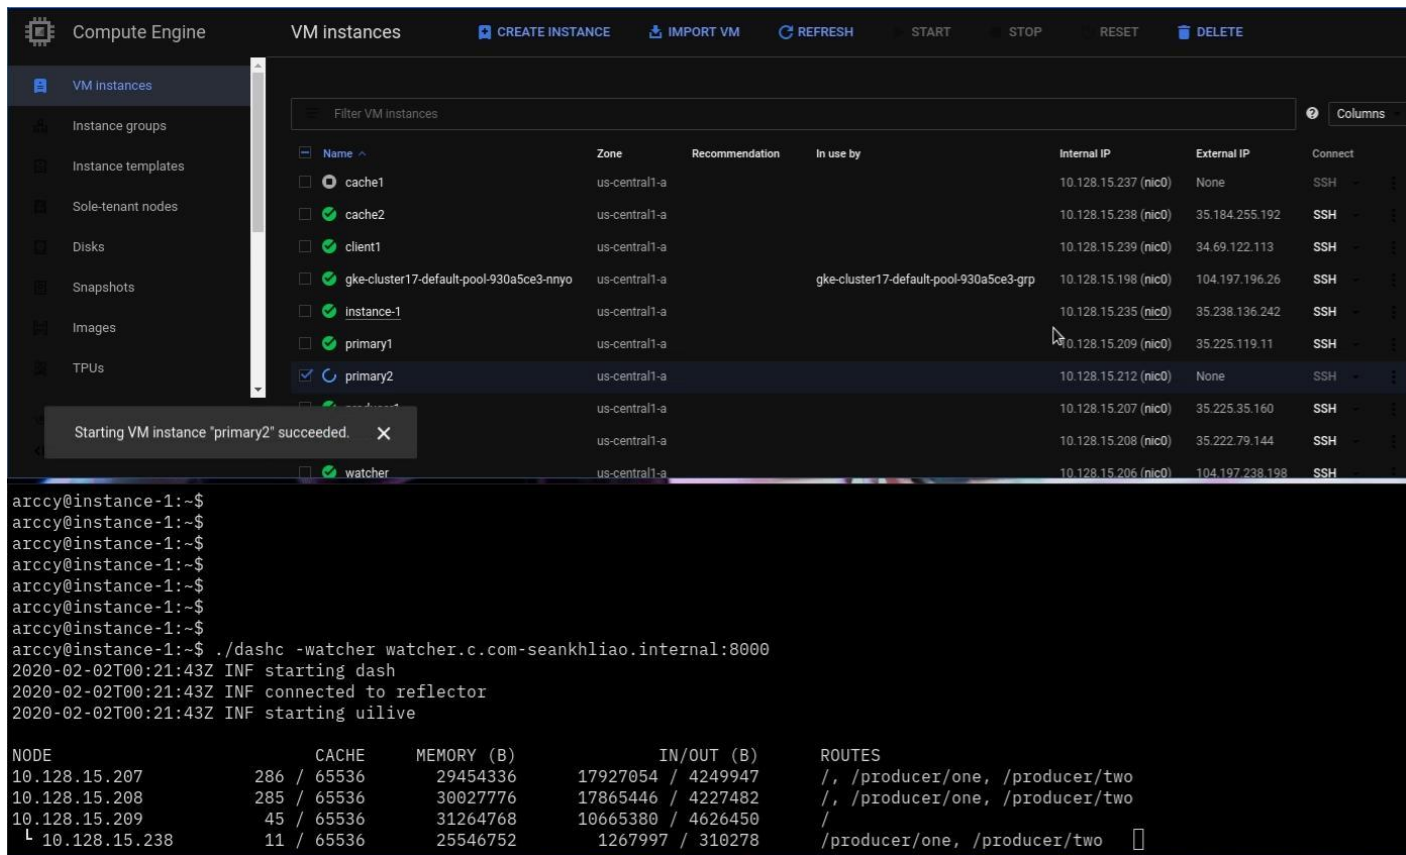
Things break or are not optimized

Future

Some new network based
on these ideas

Or maybe not (IPv6...)

Demo: add load balancer (40s)



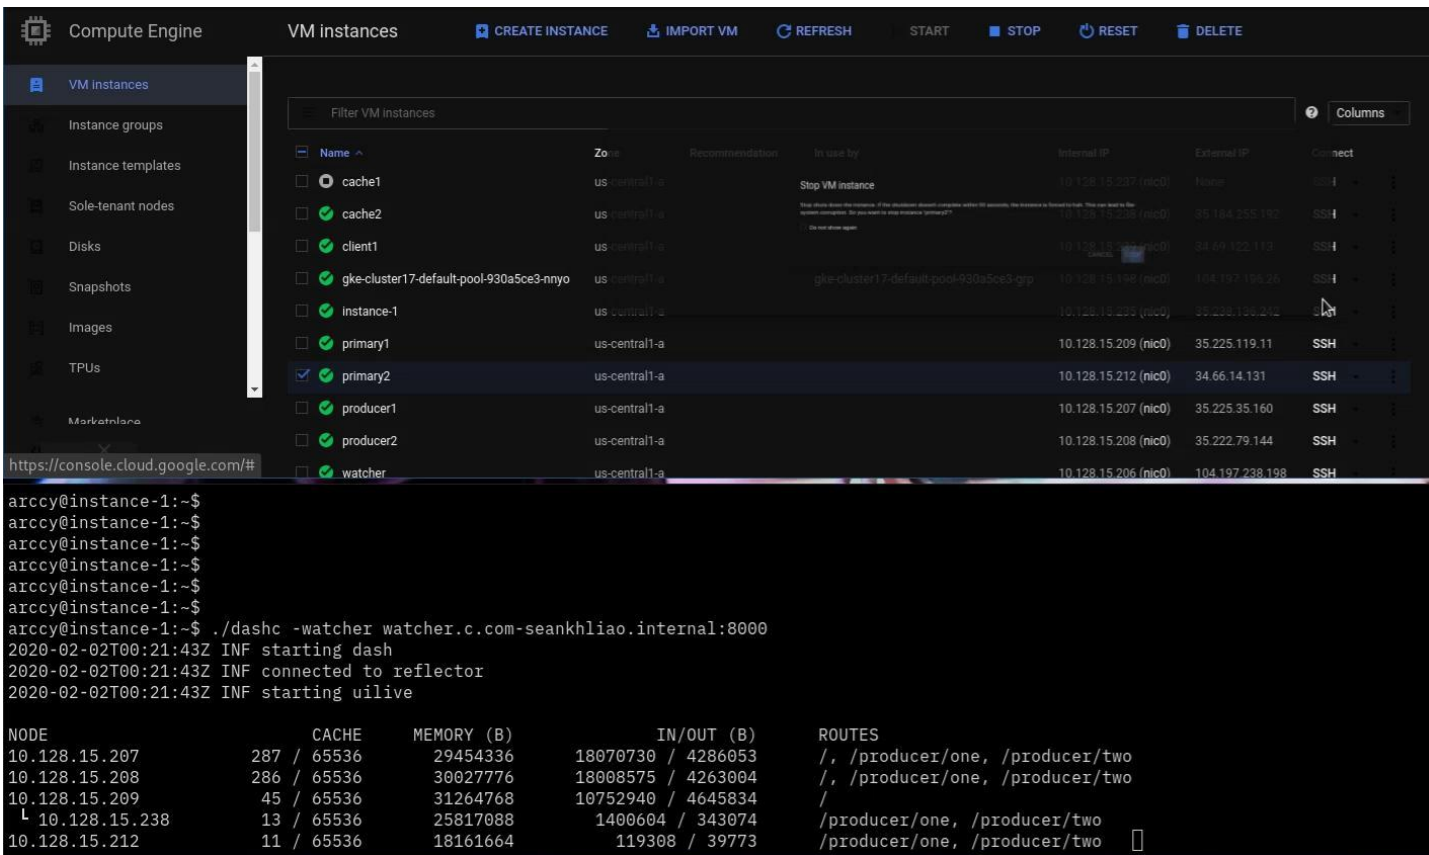
The screenshot shows the Google Cloud Platform VM instances page. The left sidebar contains navigation links for VM instances, instance groups, instance templates, sole-tenant nodes, disks, snapshots, images, and TPUs. The main area displays a table of VM instances with columns for Name, Zone, Recommendation, In use by, Internal IP, External IP, and Connect. A notification at the bottom left states "Starting VM instance 'primary2' succeeded." Below the table, a terminal window shows the output of a command to start a dash service.

Name	Zone	Recommendation	In use by	Internal IP	External IP	Connect
cache1	us-central1-a			10.128.15.237 (nic0)	None	SSH
cache2	us-central1-a			10.128.15.238 (nic0)	35.184.255.192	SSH
client1	us-central1-a			10.128.15.239 (nic0)	34.69.122.113	SSH
gke-cluster17-default-pool-930a5ce3-nnyo	us-central1-a		gke-cluster17-default-pool-930a5ce3-grp	10.128.15.198 (nic0)	104.197.196.26	SSH
instance-1	us-central1-a			10.128.15.235 (nic0)	35.238.136.242	SSH
primary1	us-central1-a			10.128.15.209 (nic0)	35.225.119.11	SSH
primary2	us-central1-a			10.128.15.212 (nic0)	None	SSH
watcher	us-central1-a			10.128.15.207 (nic0)	35.225.35.160	SSH
	us-central1-a			10.128.15.208 (nic0)	35.222.79.144	SSH
	us-central1-a			10.128.15.206 (nic0)	104.197.238.198	SSH

```
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$ ./dashc -watcher watcher.c.com-seankhliao.internal:8000
2020-02-02T00:21:43Z INF starting dash
2020-02-02T00:21:43Z INF connected to reflector
2020-02-02T00:21:43Z INF starting uilive
```

NODE	CACHE	MEMORY (B)	IN/OUT (B)	ROUTES
10.128.15.207	286 / 65536	29454336	17927054 / 4249947	/, /producer/one, /producer/two
10.128.15.208	285 / 65536	30027776	17865446 / 4227482	/, /producer/one, /producer/two
10.128.15.209	45 / 65536	31264768	10665380 / 4626450	/
10.128.15.238	11 / 65536	25546752	1267997 / 310278	/producer/one, /producer/two

Demo: remove load balancer (16s)



Compute Engine VM instances

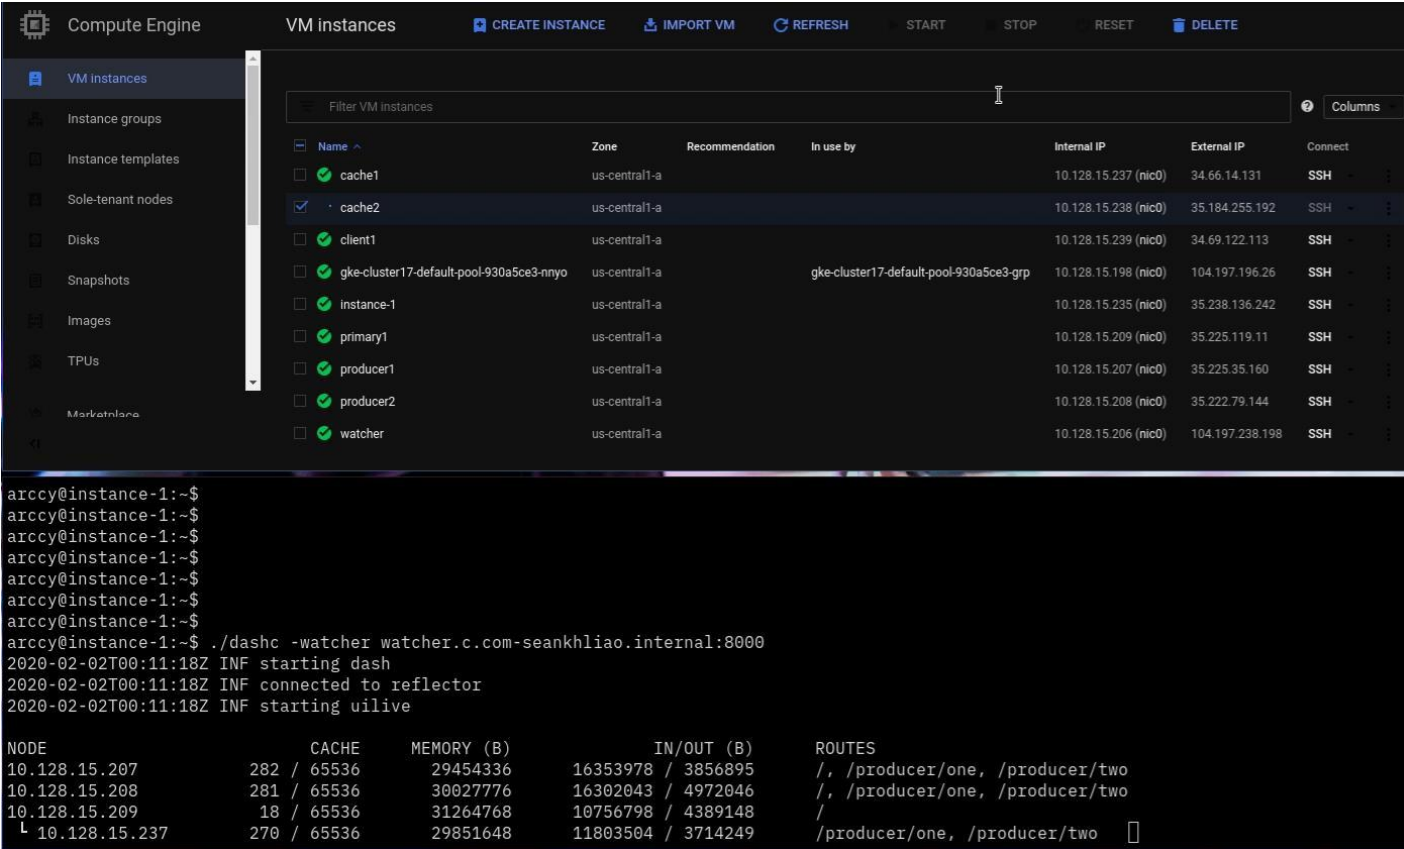
Filter VM instances

Name	Zone	Recommendation	In use by	Internal IP	External IP	Connect
<input type="checkbox"/> cache1	us-central1-a	Stop VM instance		10.128.15.237 (nic0)	None	SSH
<input type="checkbox"/> cache2	us-central1-a	Stop VM instance		10.128.15.238 (nic0)	35.184.255.192	SSH
<input type="checkbox"/> client1	us-central1-a			10.128.15.239 (nic0)	34.69.122.119	SSH
<input type="checkbox"/> gke-cluster17-default-pool-930a5ce3-nnyo	us-central1-a		gke-cluster17-default-pool-930a5ce3-nnyo	10.128.15.198 (nic0)	104.197.196.26	SSH
<input type="checkbox"/> instance-1	us-central1-a			10.128.15.235 (nic0)	35.228.150.212	SSH
<input type="checkbox"/> primary1	us-central1-a			10.128.15.209 (nic0)	35.225.119.11	SSH
<input checked="" type="checkbox"/> primary2	us-central1-a			10.128.15.212 (nic0)	34.66.14.131	SSH
<input type="checkbox"/> producer1	us-central1-a			10.128.15.207 (nic0)	35.225.35.160	SSH
<input type="checkbox"/> producer2	us-central1-a			10.128.15.208 (nic0)	35.222.79.144	SSH
<input type="checkbox"/> watcher	us-central1-a			10.128.15.206 (nic0)	104.197.238.198	SSH

```
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$ ./dashc -watcher watcher.c.com-seankhliao.internal:8000
2020-02-02T00:21:43Z INF starting dash
2020-02-02T00:21:43Z INF connected to reflector
2020-02-02T00:21:43Z INF starting uilive

NODE          CACHE    MEMORY (B)    IN/OUT (B)    ROUTES
10.128.15.207 287 / 65536 29454336      18070730 / 4286053  /, /producer/one, /producer/two
10.128.15.208 286 / 65536 30027776      18008575 / 4263004  /, /producer/one, /producer/two
10.128.15.209 45 / 65536 31264768      10752940 / 4645834  /
10.128.15.238 13 / 65536 25817088      1400604 / 343074    /producer/one, /producer/two
10.128.15.212 11 / 65536 18161664      119308 / 39773     /producer/one, /producer/two
```

Demo: add Cache (44s)



The screenshot shows the Google Cloud Platform interface for VM instances. The left sidebar lists navigation options: VM Instances (selected), Instance groups, Instance templates, Sole-tenant nodes, Disks, Snapshots, Images, TPUs, and Marketplace. The main panel displays a table of VM instances with columns: Name, Zone, Recommendation, In use by, Internal IP, External IP, and Connect. The 'cache2' instance is selected. Below the table, a terminal window shows the command `./dashc -watcher watcher.c.com-seankhliao.internal:8000` and its output, including log messages and a table of node statistics.

Name	Zone	Recommendation	In use by	Internal IP	External IP	Connect
cache1	us-central1-a			10.128.15.237 (nic0)	34.66.14.131	SSH
cache2	us-central1-a			10.128.15.238 (nic0)	35.184.255.192	SSH
client1	us-central1-a			10.128.15.239 (nic0)	34.69.122.113	SSH
gke-cluster17-default-pool-930a5ce3-nnyo	us-central1-a		gke-cluster17-default-pool-930a5ce3-grp	10.128.15.198 (nic0)	104.197.196.26	SSH
instance-1	us-central1-a			10.128.15.235 (nic0)	35.238.136.242	SSH
primary1	us-central1-a			10.128.15.209 (nic0)	35.225.119.11	SSH
producer1	us-central1-a			10.128.15.207 (nic0)	35.225.35.160	SSH
producer2	us-central1-a			10.128.15.208 (nic0)	35.222.79.144	SSH
watcher	us-central1-a			10.128.15.206 (nic0)	104.197.238.198	SSH

```
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$
arccy@instance-1:~$ ./dashc -watcher watcher.c.com-seankhliao.internal:8000
2020-02-02T00:11:18Z INF starting dash
2020-02-02T00:11:18Z INF connected to reflector
2020-02-02T00:11:18Z INF starting uilive

NODE          CACHE      MEMORY (B)      IN/OUT (B)      ROUTES
10.128.15.207 282 / 65536    29454336        16353978 / 3856895  /, /producer/one, /producer/two
10.128.15.208 281 / 65536    30027776        16302043 / 4972046   /, /producer/one, /producer/two
10.128.15.209 18 / 65536     31264768        10756798 / 4389148   /
10.128.15.237 270 / 65536    29851648        11803504 / 3714249   /producer/one, /producer/two
```

Demo: remove Cache (17s)

The screenshot shows the Google Cloud Platform interface for VM instances. A list of VM instances is displayed, including 'cache1', 'cache2', 'client1', 'gke-cluster17-default-pool-930a5ce3-nnyo', 'instance-1', 'primary1', 'producer1', 'producer2', and 'watcher'. A modal dialog is open over 'cache1' with the title 'Stopping VM instance "cache1"'. Below the list, a terminal window shows the command `./dashc -watcher watcher.c.com-seankhliao.internal:8000` and its output, which includes logs for starting dash, connecting to reflector, and starting uilive. The terminal also displays a table of system metrics for various nodes.

NODE	CACHE	MEMORY (B)	IN/OUT (B)	ROUTES
10.128.15.207	285 / 65536	29454336	16627390 / 3928689	/, /producer/one, /producer/two
10.128.15.208	284 / 65536	30027776	16572803 / 5040816	/, /producer/one, /producer/two
10.128.15.209	40 / 65536	31264768	10948720 / 4438910	/
10.128.15.237	275 / 65536	29851648	12052764 / 3779058	/producer/two
10.128.15.238	7 / 65536	16896000	65702 / 24320	/producer/one