

Cross-Cloud Connectivity

Diego Casati – Senior Software Engineer Commercial Software Engineering

```
NAME
```

```
dcasati -- Diego Casati: a (curious) Engineer
```

SYNOPSIS

HISTORY

Diego Casati is a Senior Software Engineer for Microsoft focusing on Kubernetes, Networking and Linux and BSDs. Prior to his current role at Microsoft, he spent over a decade working in the Telco and IT industries at various capacities, from Networking Engineering to Systems Engineering and Security Specialist.

He is a strong proponent of free and open source solutions, advocating for the use of BSDs to connect all things.

When not hacking on computers you can catch him spending time playing with his 1-year-old son.

SEE ALSO

- dcasati @github
- diegocasati @Twitter

BUGS

Too many to list here ;)





Cross-Cloud Connectivity

What worked, what failed and lessons learned.

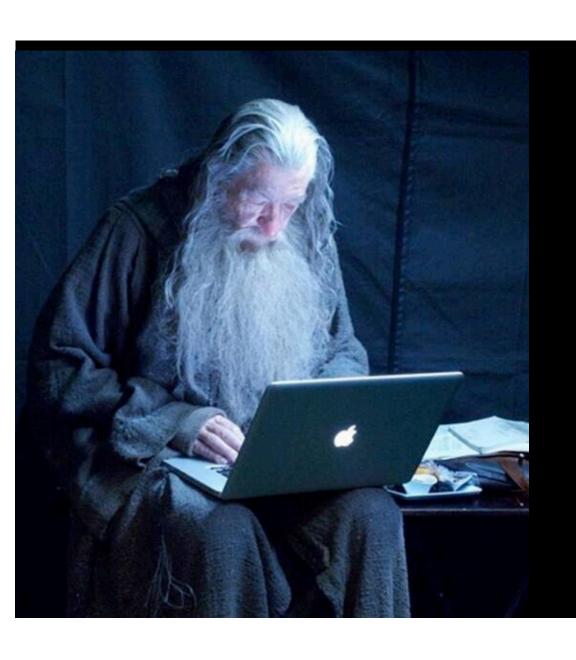
Key objective

Ignite your curiosity by showing a real use case of an engineering exercise.

What?
Why?
Main() Where?
How?

What?

Deploy Cassandra on top of Kubernetes spanning two clouds



Wait? What? Let me check online Why?

Proactive exercise simulating Customers looking to (1) migrate or have their workload (2) spanning in two clouds

Why?

(1) migrate or have their workload

- For replication (e.g.: backup)
- To comply with a scheduled downtime (e.g.: keep a service running by moving data between Data Centers)

Why?

(2) spanning in two clouds

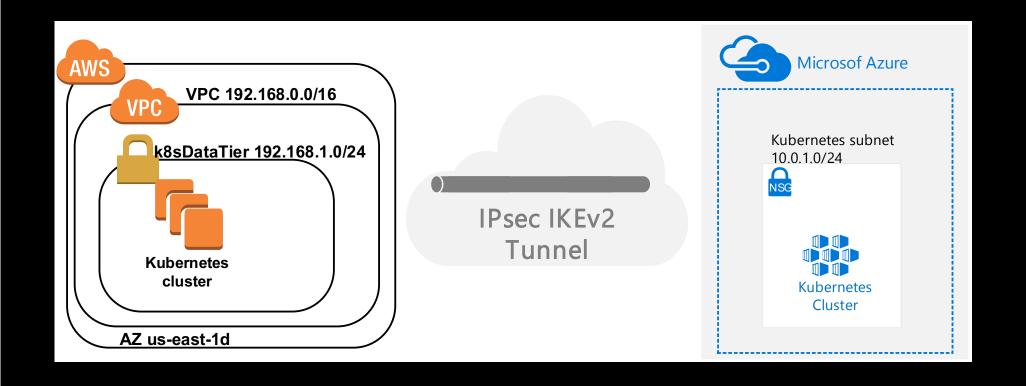
- To increased availability
- To provide load balancing
- To connect to an existing Data Center (e.g.: on premise)

Where? AWS Azure

How?

site-to-site VPN between the clouds using open source components

END-TO-END TOPOLOGY - NETWORK

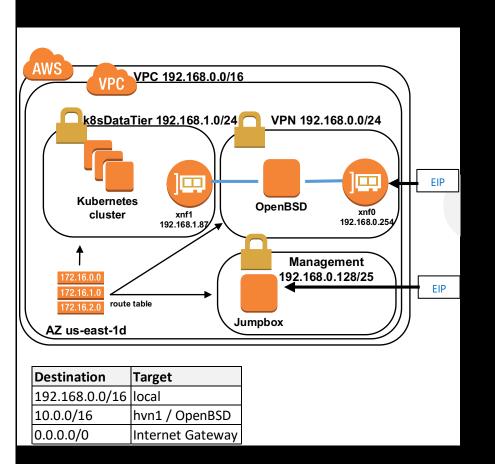




Common Vocabulary

Acronym	Explanation
Jumpbox	A jump server is a hardened and monitored device that spans two dissimilar security zones and provides a controlled means of access between them.
VPN	A virtual private network (VPN) extends a private network across a public network, and enables users to send and receive data across shared or public networks as if their computing devices were directly connected to the private network.

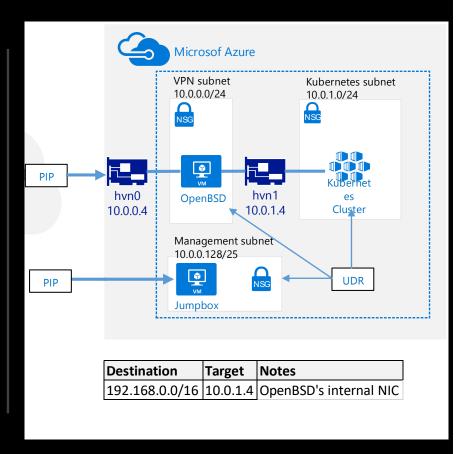
AWS Vocabulary



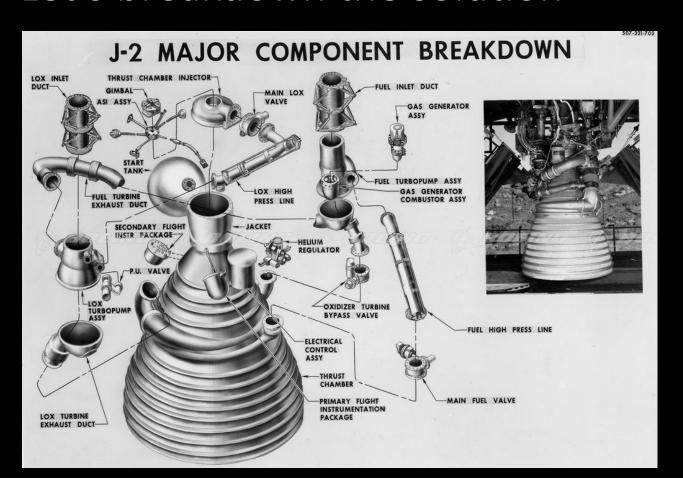
Acronym	Explanation	
VPC	A <i>virtual private cloud</i> (VPC) is a virtual network dedicated to your AWS account.	
EIP	An <i>Elastic IP address</i> is a static, public IPv4 address designed for dynamic cloud computing.	
Security Group	A security group acts as a virtual firewall for your instance to control inbound and outbound traffic	

Azure Vocabulary

Acronym	Explanation	
VNET	Azure Virtual Network enables Azure resources to communicate with each other and the internet.	
PIP	A PIP is a <i>public instance-level IP</i> address associated with the VM in addition to the VIP.	
NSG	A network security group (NSG) contains a list of security rules that allow or deny network traffic to resources connected to Azure Virtual Networks (VNet).	
UDR	User defined routes	



Let's breakdown the solution



The solution – in pieces

Solution	AWS	Azure
Cassandra	Kubernetes YAML files	Helm Chart
Kubernetes	KOPS	ACS-Engine
Network	VPN: OpenVPN? StrongSwan? Other IPSec solution?	

Cassandra



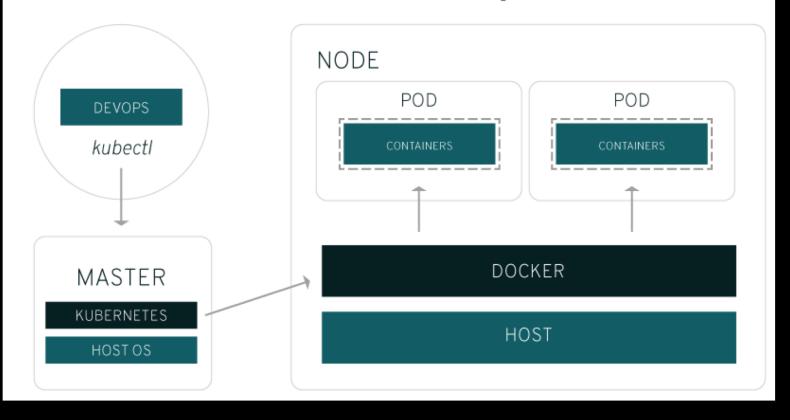
Apache Cassandra is a free and open-source distributed NoSQL database management system designed to handle large amounts of data across many commodity servers, providing high availability with no single point of failure



Kubernetes is an open-source system for automating deployment, scaling, and management of containerized applications.



A look at how Kubernetes fits into your infrastructure



https://www.redhat.com/en/topics/containers/what-is-kubernetes

Microsoft Azure Container Service Engine - Builds Docker Enabled Clusters



Overview

The Azure Container Service Engine (acs-engine) generates ARM (Azure Resource Manager) templates for Docker enabled clusters on Microsoft Azure with your choice of DC/OS, Kubernetes, Swarm Mode, or Swarm orchestrators. The input to the tool is a cluster definition. The cluster definition is very similar to (in many cases the same as) the ARM template syntax used to deploy a Microsoft Azure Container Service cluster.

The cluster definition file enables the following customizations to your Docker enabled cluster:

- · choice of DC/OS, Kubernetes, Swarm Mode, or Swarm orchestrators
- · multiple agent pools where each agent pool can specify:
- · standard or premium VM Sizes,
- node count,
- · Virtual Machine ScaleSets or Availability Sets,
- Storage Account Disks or Managed Disks (under private preview)
- Docker cluster sizes of 1200
- Custom VNET

ACS-Engine https://github.com/Azure/acs-engine

kops - Kubernetes Operations

build passing go report A+ godoc reference

The easiest way to get a production grade Kubernetes cluster up and running.

What is kops?

We like to think of it as kubectl for clusters.

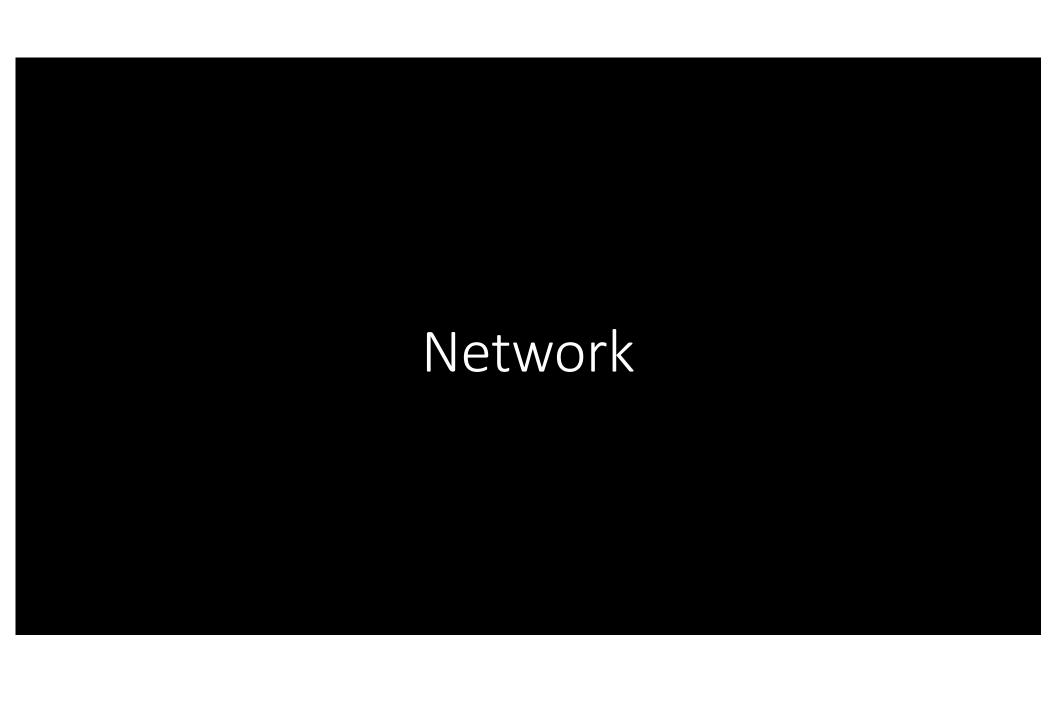
kops helps you create, destroy, upgrade and maintain production-grade, highly available, Kubernetes clusters from the command line. AWS (Amazon Web Services) is currently officially supported, with GCE in beta support, and VMware vSphere in alpha, and other platforms planned.

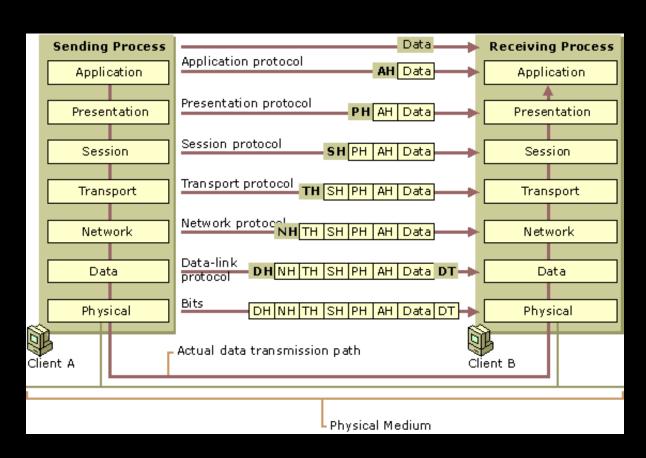
Can I see it in action?

```
Autoscaling/roug/modes.example.nivenly.com
Nindize
Nin
```

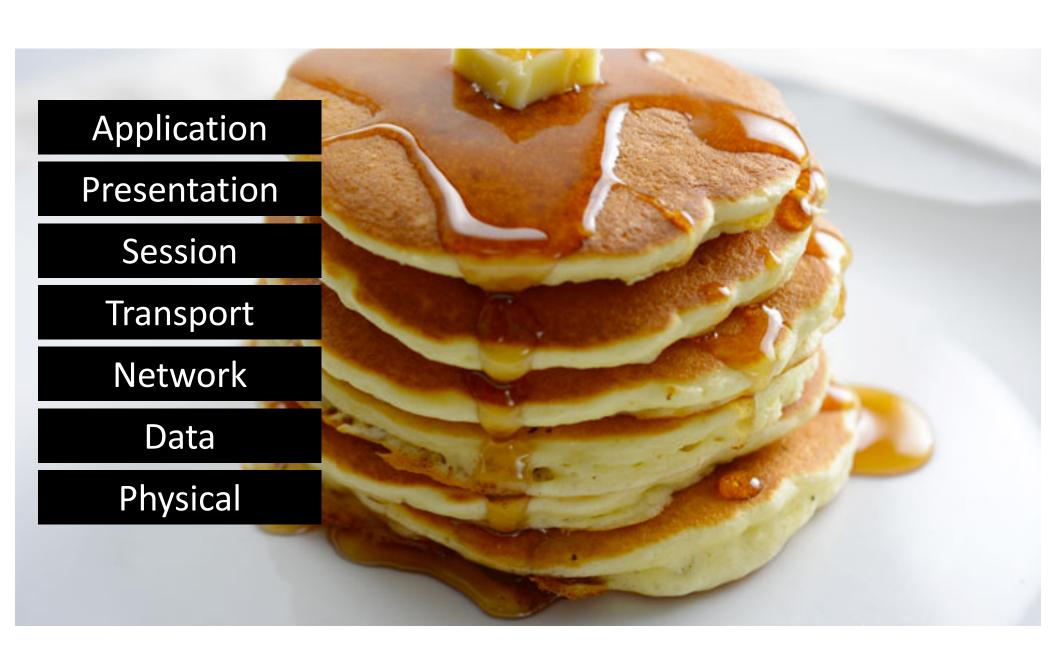
KOPS

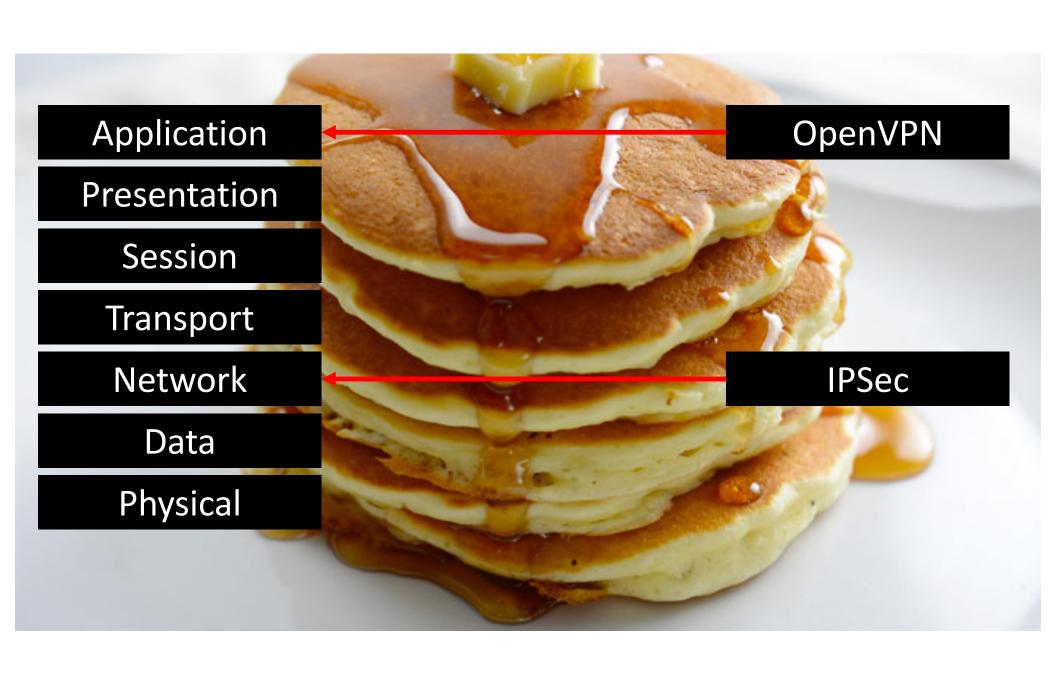
https://github.com/kubernetes/kops

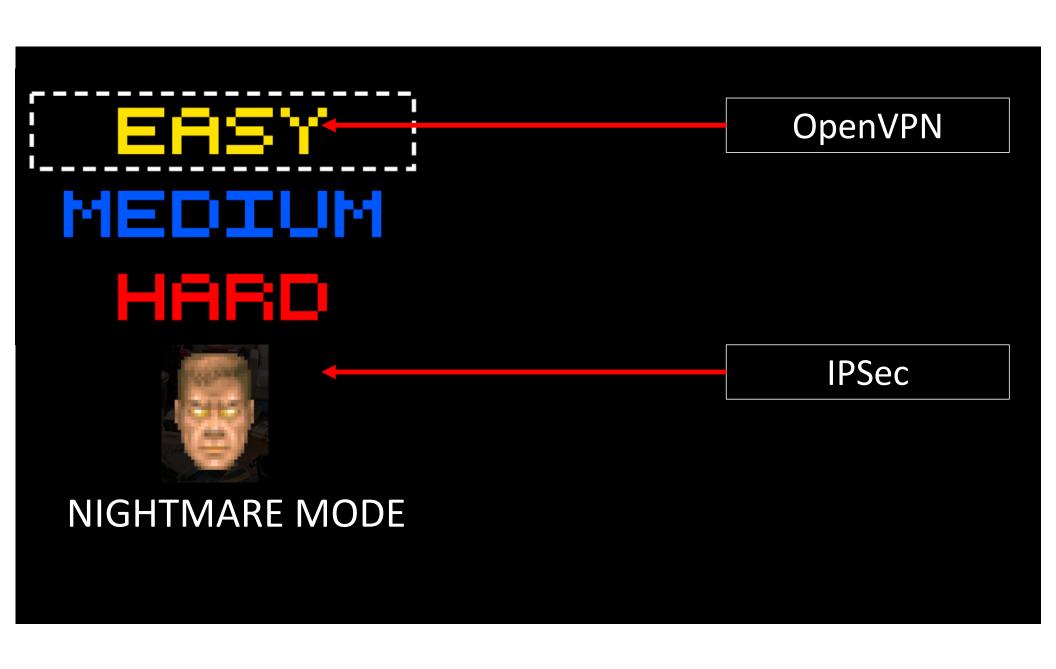




OSI Model – It's 1999 all over again







IPsec and its RFCs — This is HARD!

IETF documentation [edit]

Standards Track [edit]

- . RFC 1829@: The ESP DES-CBC Transform
- . RFC 2403@: The Use of HMAC-MD5-96 within ESP and AH
- RFC 2405∉: The ESP DES-CBC Cipher Algorithm With Explicit IV
- . RFC 24106: The NULL Encryption Algorithm and Its Use With IPsec
- . RFC 2451 €: The ESP CBC-Mode Cipher Algorithms
- . RFC 2857@: The Use of HMAC-RIPEMD-160-96 within ESP and AH
- . RFC 3526€: More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE)
- RFC 3686€: Using Advanced Encryption Standard (AES) Counter Mode With IPsec Encapsulating Security Payload (ESP)
- . RFC 3947@: Negotiation of NAT-Traversal in the IKE
- RFC 3948∉: UDP Encapsulation of IPsec ESP Packets
- . RFC 41066: The Use of Galois/Counter Mode (GCM) in IPsec Encapsulating Security Payload (ESP)
- RFC 4301 €: Security Architecture for the Internet Protocol
- . RFC 4302 €: IP Authentication Header
- RFC 4303@: IP Encapsulating Security Payload
- RFC 4304@: Extended Sequence Number (ESN) Addendum to IPsec Domain of Interpretation (DOI) for Internet Security Association and Key Management Protocol (ISAKMP)
- . RFC 4307€: Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)
- . RFC 4308@: Cryptographic Suites for IPsec
- RFC 4309⊕: Using Advanced Encryption Standard (AES) CCM Mode with IPsec Encapsulating Security Payload (ESP)
- . RFC 4555€: IKEv2 Mobility and Multihoming Protocol (MOBIKE)
- . RFC 4806@: Online Certificate Status Protocol (OCSP) Extensions to IKEv2
- . RFC 4868€: Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec
- RFC 5280⊗: Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile
- RFC 52829: Using Authenticated Encryption Algorithms with the Encrypted Payload of the Internet Key Exchange version 2 (IKEv2) Protocol
- RFC 5386€: Better-Than-Nothing Security: An Unauthenticated Mode of IPsec
- . RFC 5529€: Modes of Operation for Camellia for Use with IPsec
- RFC 56856: Redirect Mechanism for the Internet Key Exchange Protocol Version 2 (IKEv2)
- RFC 57236: Internet Key Exchange Protocol Version 2 (IKEv2) Session Resumption
- RFC 58576: IKEv2 Extensions to Support Robust Header Compression over IPsec
 RFC 58586: IPsec Extensions to Support Robust Header Compression over IPsec
- RFC 7296∉: Internet Key Exchange Protocol Version 2 (IKEv2)
- RFC 7321€: Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)
- . RFC 7383∉: Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation
- RFC 7427⊕: Signature Authentication in the Internet Key Exchange Version 2 (IKEv2)
- RFC 7634€: ChaCha20, Poly1305, and Their Use in the Internet Key Exchange Protocol (IKE) and IPsec

Experimental RFCs [edit

Informational RFCs [edit]

- . RFC 2367ø: PF_KEY Interface
- . RFC 2412 €: The OAKLEY Key Determination Protocol
- RFC 3706⊕: A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers
- RFC 37156: IPsec-Network Address Translation (NAT) Compatibility Requirements
- . RFC 4621€: Design of the IKEv2 Mobility and Multihoming (MOBIKE) Protocol
- RFC 4809 Requirements for an IPsec Certificate Management Profile
- RFC 5387∉: Problem and Applicability Statement for Better-Than-Nothing Security (BTNS)
- RFC 58569: Integration of Robust Header Compression over IPsec Security Associations
- RFC 5930Ø: Using Advanced Encryption Standard Counter Mode (AES-CTR) with the Internet Key Exchange version 02 (IKEv2) Protocol
- RFC 6027∉: IPsec Cluster Problem Statement

Fear not OpenIKED is here!



https://www.openiked.org

OpenIKED is

a FREE implementation of the Internet Key Exchange (IKEv2) protocol which performs mutual authentication and which establishes and maintains IPsec VPN security policies and associations (SAs) between peers.

The solution – in pieces

Solution	AWS	Azure
Cassandra	Kubernetes YAML files	Helm Chart
Kubernetes	KOPS	ACS-Engine
Network		VPN: OpenIKED

What failed?

#FAIL!

Attempt	Approach	Why it failed?
#1	Site-to-site VPN between AWS's Virtual Gateway and Azure's VPN Gateway	Tunnels did not sync Hard to troubleshoot
#2	GRE tunnel between FreeBSD VMs	Site-to-site VPN between AWS's Virtual Gateway and Azure's VPN Gateway
#3	GIF tunnel between FreeBSD VMs	
#4	IPsec tunnel on FreeBSD (need either GRE or GIF)	

What worked?

#IT_WORKS!

Attempt	Approach	Notes
#1	SSH-based VPN	For testing only. Overhead is too high. Refer to ssh(1) for details
#2	VM to Gateway service (both AWS and Azure)	Works for NVAs such as Cisco ASA's
#3	IKEv2 tunnel with OpenBSD	IPsec. Easiest OSS option

```
$ ssh -o "VerifyHostKeyDNS ask" host.example.com
[...]
Matching host key fingerprint found in DNS.
Are you sure you want to continue connecting (yes/no)?
```

See the VerifyHostKeyDNS option in ssh_config(5) for more information.

SSH-BASED VIRTUAL PRIVATE NETWORKS

ssh contains support for Virtual Private Network (VPN) tunnelling using the tun(4) network pseudo-device, allowing two networks to be joined securely. The sshd_config(5) configuration option PermitTunnel controls whether the server supports this, and at what level (layer 2 or 3 traffic).

The following example would connect client network 10.0.50.0/24 with remote network 10.0.99.0/24 using a point-to-point connection from 10.1.1.1 to 10.1.1.2, provided that the SSH server running on the gateway to the remote network, at 192.168.1.15, allows it.

On the client:

```
# ssh -f -w 0:1 192.168.1.15 true
# ifconfig tun0 10.1.1.1 10.1.1.2 netmask 255.255.255.252
# route add 10.0.99.0/24 10.1.1.2
```

On the server:

```
# ifconfig tun1 10.1.1.2 10.1.1.1 netmask 255.255.255.252
# route add 10.0.50.0/24 10.1.1.1
```

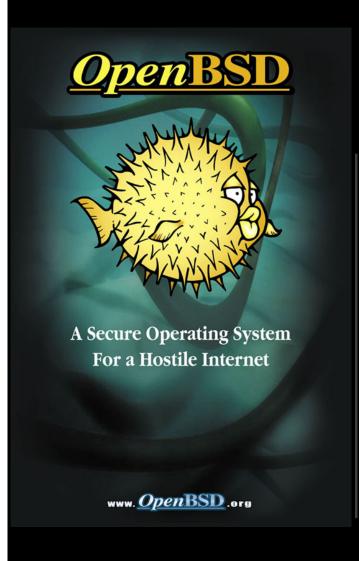
Client access may be more finely tuned via the <u>/root/.ssh/authorized_keys</u> file (see below) and the **PermitRootLogin** server option. The following entry would permit connections on tun(4) device 1 from user ``jane'' and on tun device 2 from user ``john'', if **PermitRootLogin** is set to ``forced-commands-only'':

```
tunnel="1",command="sh /etc/netstart tun1" ssh-rsa ... jane tunnel="2",command="sh /etc/netstart tun2" ssh-rsa ... john
```

Since an SSH-based setup entails a fair amount of overhead, it may be more suited to temporary setups, such as for wireless VPNs. More permanent VPNs are better provided by tools such as ipsecctl(8) and isakmpd(8).

The solution – in pieces

Solution	AWS	Azure
Cassandra	Kubernetes YAML files	Helm Chart
Kubernetes	KOPS	ACS-Engine
Network	VPN: OpenIKED on (<mark>OpenBSD</mark>



Why OpenBSD?

- 1. A FREE, multi-platform 4.4BSD-based UNIX-like operating system.
- **2. Proactive security** (strlcpy, strlcat, W^X, privilege separation, ...)
- 3. Many of our day-today tools in the base install (by default):
 - tmux, nvi, OpenSSH, mg (emacs-like), tcpdump, ...
- 4. Infrastructure tools:
 - OpenIKED IKEv2 Daemon
 - OpenBGPD BGP daemon
 - OpenNTPD NTP daemon
 - OpenOSPFD OSPF daemon
 - OpenSMTPD SMTP daemon
 - Relayd L3/7 Load balancer
 - httpd HTTP daemon
 - ACME-client ACME certificate client
 - Rebound DNS proxy
 - PF firewall

Why OpenBSD?

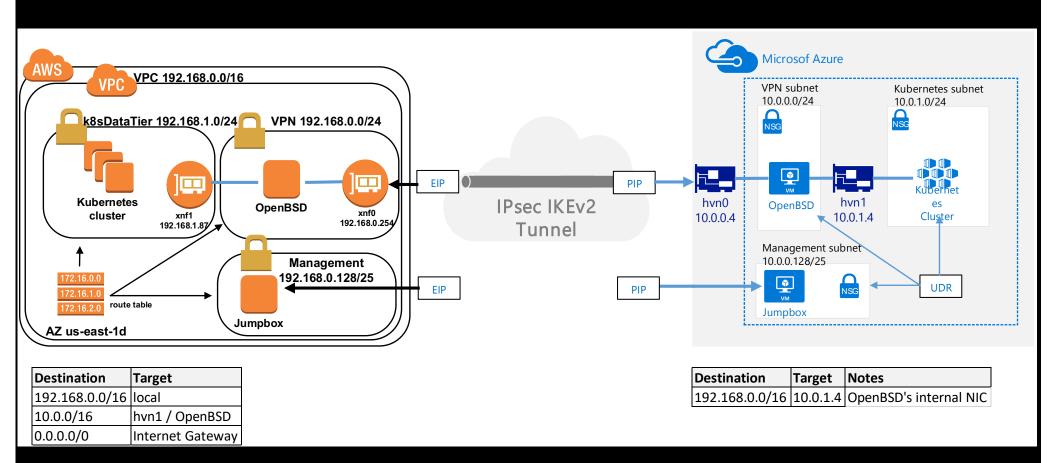
Problem:

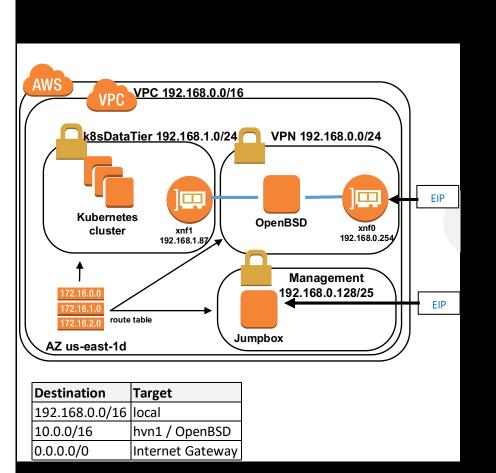
- Vanilla image available on AWS or Azure.
- Option to use Esdenera Firewall 3 (NVA based on OpenBSD).
- Our current documentation is not working – Needs TLC.

Solution:

- Bake your own image.
- Based on scripts from core OpenBSD devs (folks from Esdenera): qemu,
 Makefile et al.
- While doing this, 3 PRs were opened to fix issues (AWS and documentation).
- All files on Github:
 - https://github.com/dcasati/cloudopenbsd
 - https://github.com/dcasati/portsazure

END-TO-END TOPOLOGY - NETWORK



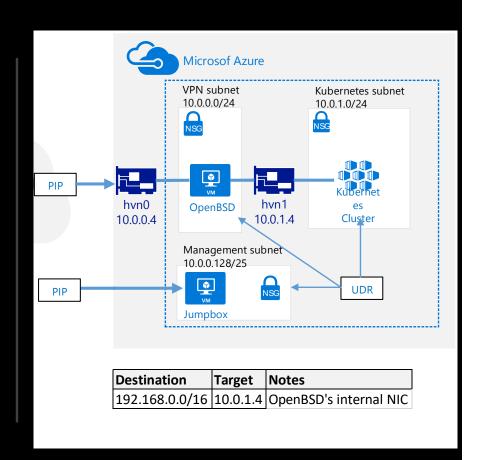


AWS

- 1. Create a VPC with a large network (e.g.: 192.168.0.0/16)
- 2. Carve 3 subnets (k8sDataTier, Management, VPN)
- Create the OpenBSD VM with two NICs
 a) For each NIC disable source/dest check
- 4. Add the route to Azure (e.g.: 10.0.0.0/8)
- 5. Allow traffic on the Security Groups (ports 500, 4500 UDP)
- 6. Attach an Elastic IP to the OpenBSD interface on the VPN subnet.
- 7. Configure OpenIKED.

Azure

- 1. Create a VNet with a large network (e.g.: 10.0.0.0/8)
- 2. Carve 3 subnets (k8sDataTier, Management, VPN)
- 3. Create the OpenBSD VM with two NICs (via Azure CLI)
 - a) For each NIC enable IP forwarding
- 4. Add the route to Azure (e.g.: 192.168.0.0/16) on the UDR
- 5. Allow traffic on the Security Groups (ports 500, 4500 UDP)
- 6. Attach a Public IP to the OpenBSD interface on the VPN subnet.
- 7. Configure OpenIKED.



OpenBSD

Configuration

/etc/iked.conf

```
local_gw = "51.143.95.27"
remote_gw = "34.233.91.14"
local_net = "10.0.0.0/16"
remote_net = "192.168.0.0/16"
kops_net = "100.64.0.0/10"
state = "active"

ikev2 $state ipcomp esp \
    from $local_gw to $remote_gw \
    from $local_net to $remote_net peer $remote_gw \
    psk "1BigSecret"

ikev2 $state ipcomp esp \
    from $local_gw to $remote_net peer $remote_gw \
    psk "1BigSecret"
```



OpenBSD



Commands

rcctl enable iked

rcctl start iked

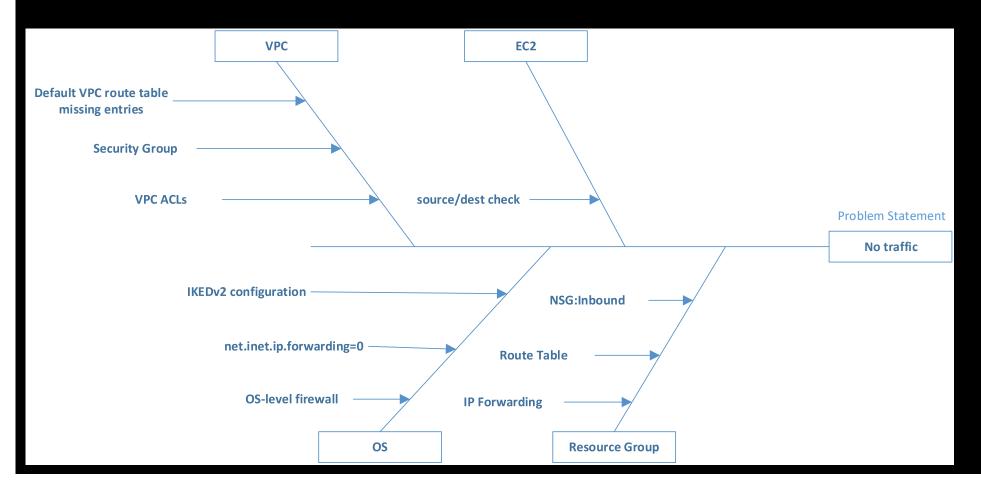
sysctl -w net.inet.ip.forwarding=1

To check if Ipsec is working:

ipsecctl –sa

Lessons Learned

Pressing all of the right buttons



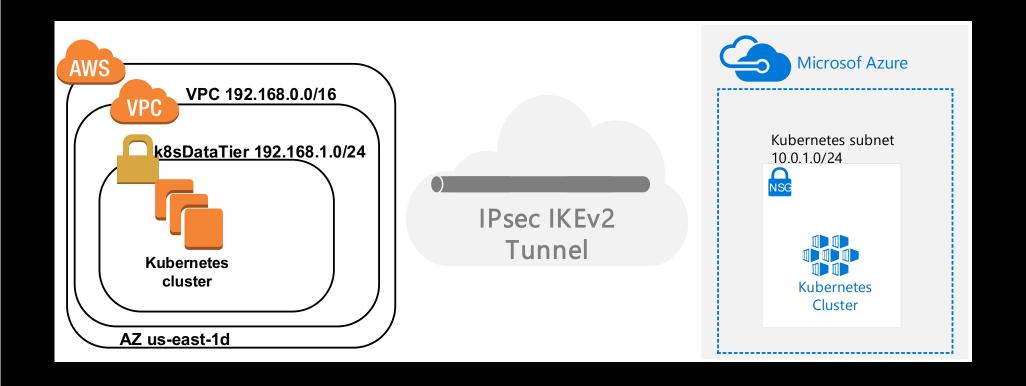


Our goal:

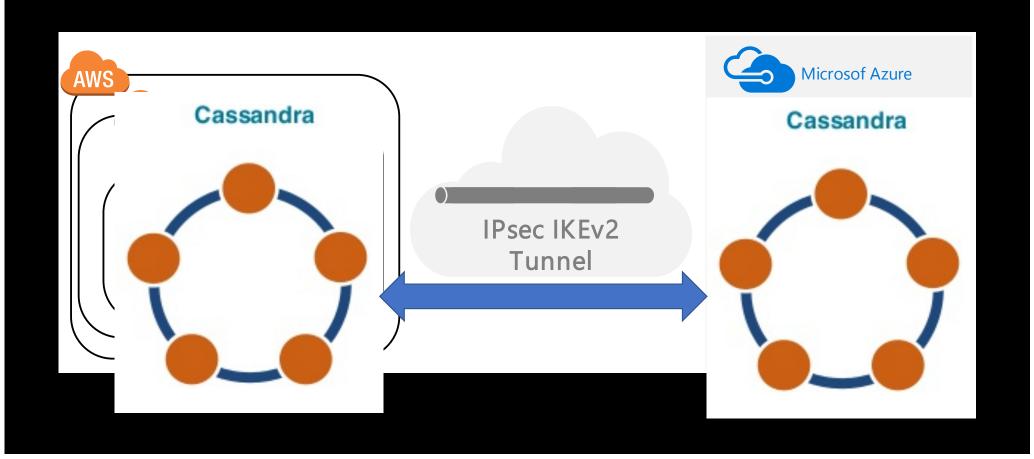
Deploy Cassandra on top of Kubernetes spanning two clouds



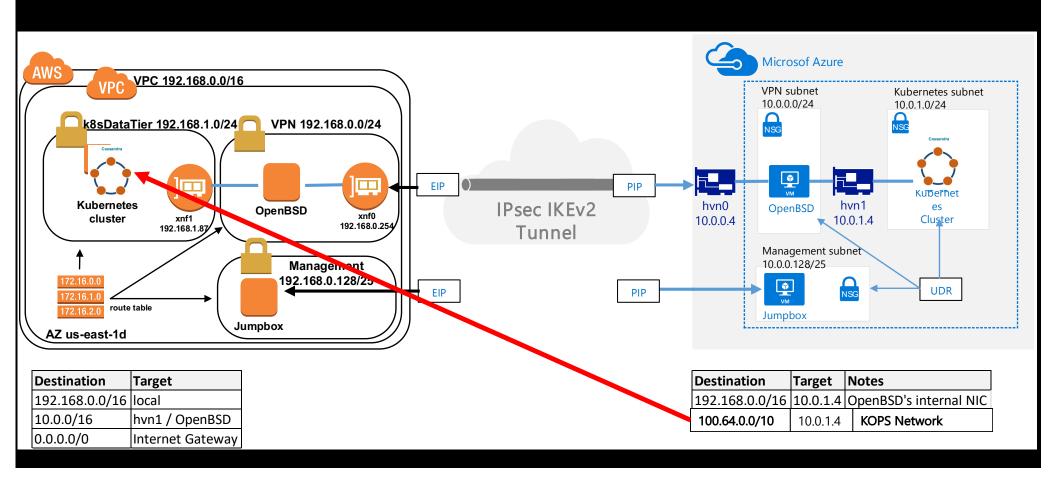
END-TO-END TOPOLOGY - KUBERNETES



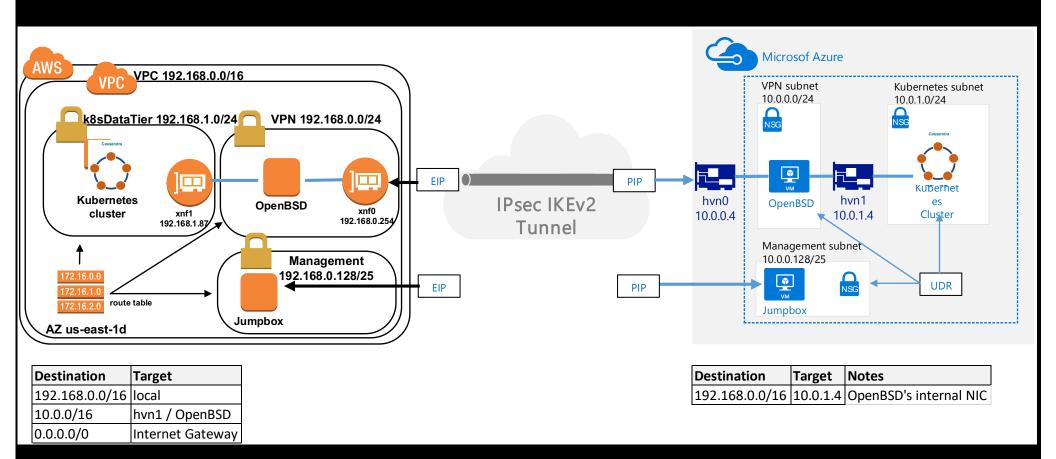
END-TO-END TOPOLOGY WITH CASSANDRA!



END-TO-END TOPOLOGY - NETWORK



END-TO-END TOPOLOGY - NETWORK



Demo

Next Steps

Hashicorp's Terraform Templates

Hashicorp's Packer to build the images

Highly Available Solutions

Key takeaways

A better view of the design choices when connecting public clouds

Understand that failures will happen – and that's ok

Reach out for help. Talk to the community

A&D

Let's talk



diegocasati @twitter

https://github.com/dcasati/cross-cloud-vpn