From Servers to Serverless in Ten Minutes

Erik Riedel, PhD
Senior Vice President, Engineering
ITRenew

OPEN PLATINUM™

OCP GLOBAL SUMMIT

Open for All.

OCP
GLOBAL
SUMMIT

The Power of Hyperscale For All

Optimized for your workload, from deskside to data center. No assembly. No guesswork. Just plug them in.

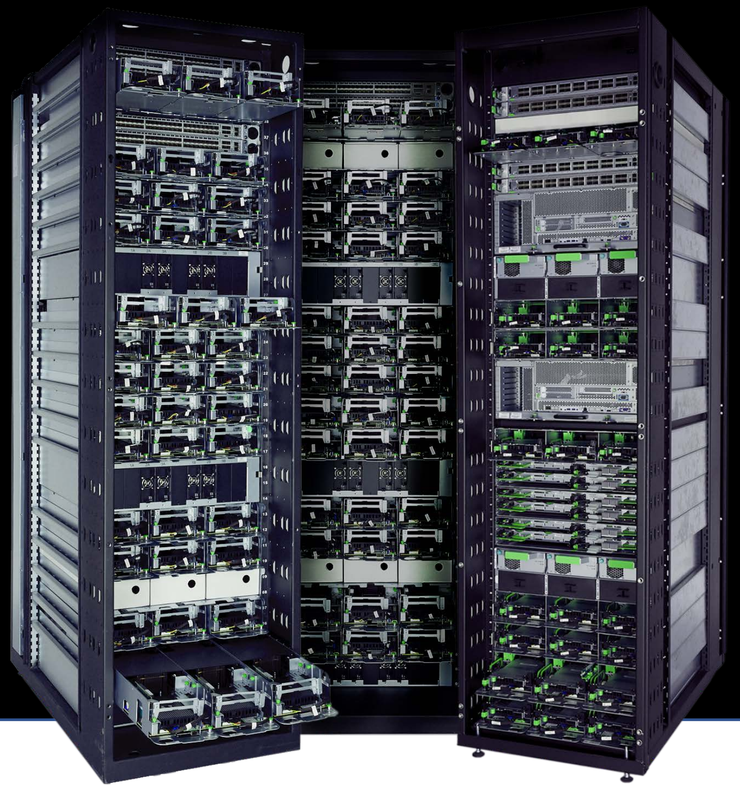

SESAME BY ITRENEW

PROVEN HYPERSCALE TECH
BUILT ON OPEN ARCHITECTURE

CONSISTENT PRODUCT
DEPENDABLE SUPPLY

BETTER-THAN-EVER TCO
FLEXIBLE, SCALABLE

OCP GLOBAL SUMMIT

Servers

OCP GLOBAL SUMMIT

Serverless

in 10 minutes…or so…

OCP GLOBAL SUMMIT

Project goal and genesis

OCP adoption could be accelerated by offering **pre-designed** and **pre-qualified** solutions for key computing use cases.

Many service provider & enterprise SaaS companies are looking for **solutions** to roll onto the floor, plug in, and quickly run workloads.

We are working with **infrastructure software** stacks and **software partners** to pre-design and pre-qualify solutions with OCP equipment.


20
19
18
17
16
15
14
13

Servers

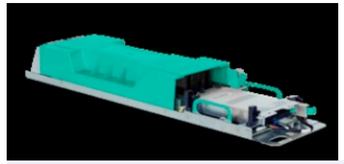

single or 2- socket nodes, 25 GbE connectivity

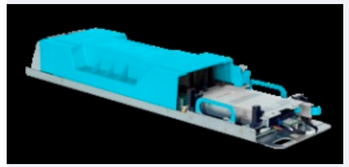

flash-based storage nodes; millions of IOPS and terabytes of capacity

external TOR switches (2x)		
ingress	ingress	ingress
internal TOR switches (2x)		
compute	compute	compute
compute	compute	compute
compute	compute	compute
compute	compute	compute
compute	compute	compute
compute	compute	compute
power zone BB		
compute	compute	compute
compute	compute	compute
compute	compute	compute
storage	storage	storage
storage	storage	storage
storage	storage	storage
mgmt	mgmt	mgmt
infra	infra	infra
power zone AA		

up to 45 nodes

SESAME
for Open Systems

OCP GLOBAL SUMMIT


27
25
23

Servers

compute
compute
compute
compute
infra
power supply + switch

up to 5 nodes

SESAME

Fast-Start

external TOR switches (2x)		
ingress	ingress	ingress
internal TOR switches (2x)		
compute	compute	compute
compute	compute	compute
compute	compute	compute
compute	compute	compute
compute	compute	compute
compute	compute	compute
power zone BB		
compute	compute	compute
compute	compute	compute
compute	compute	compute
storage	storage	storage
storage	storage	storage
storage	storage	storage
mgmt	mgmt	mgmt
infra	infra	infra
power zone AA		

up to 45 nodes

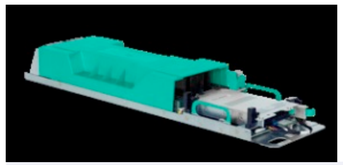

single or 2- socket nodes, 25 GbE connectivity

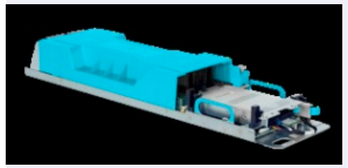

flash-based storage nodes; millions of IOPS and terabytes of capacity

SESAME

for Open Systems

OCP GLOBAL SUMMIT

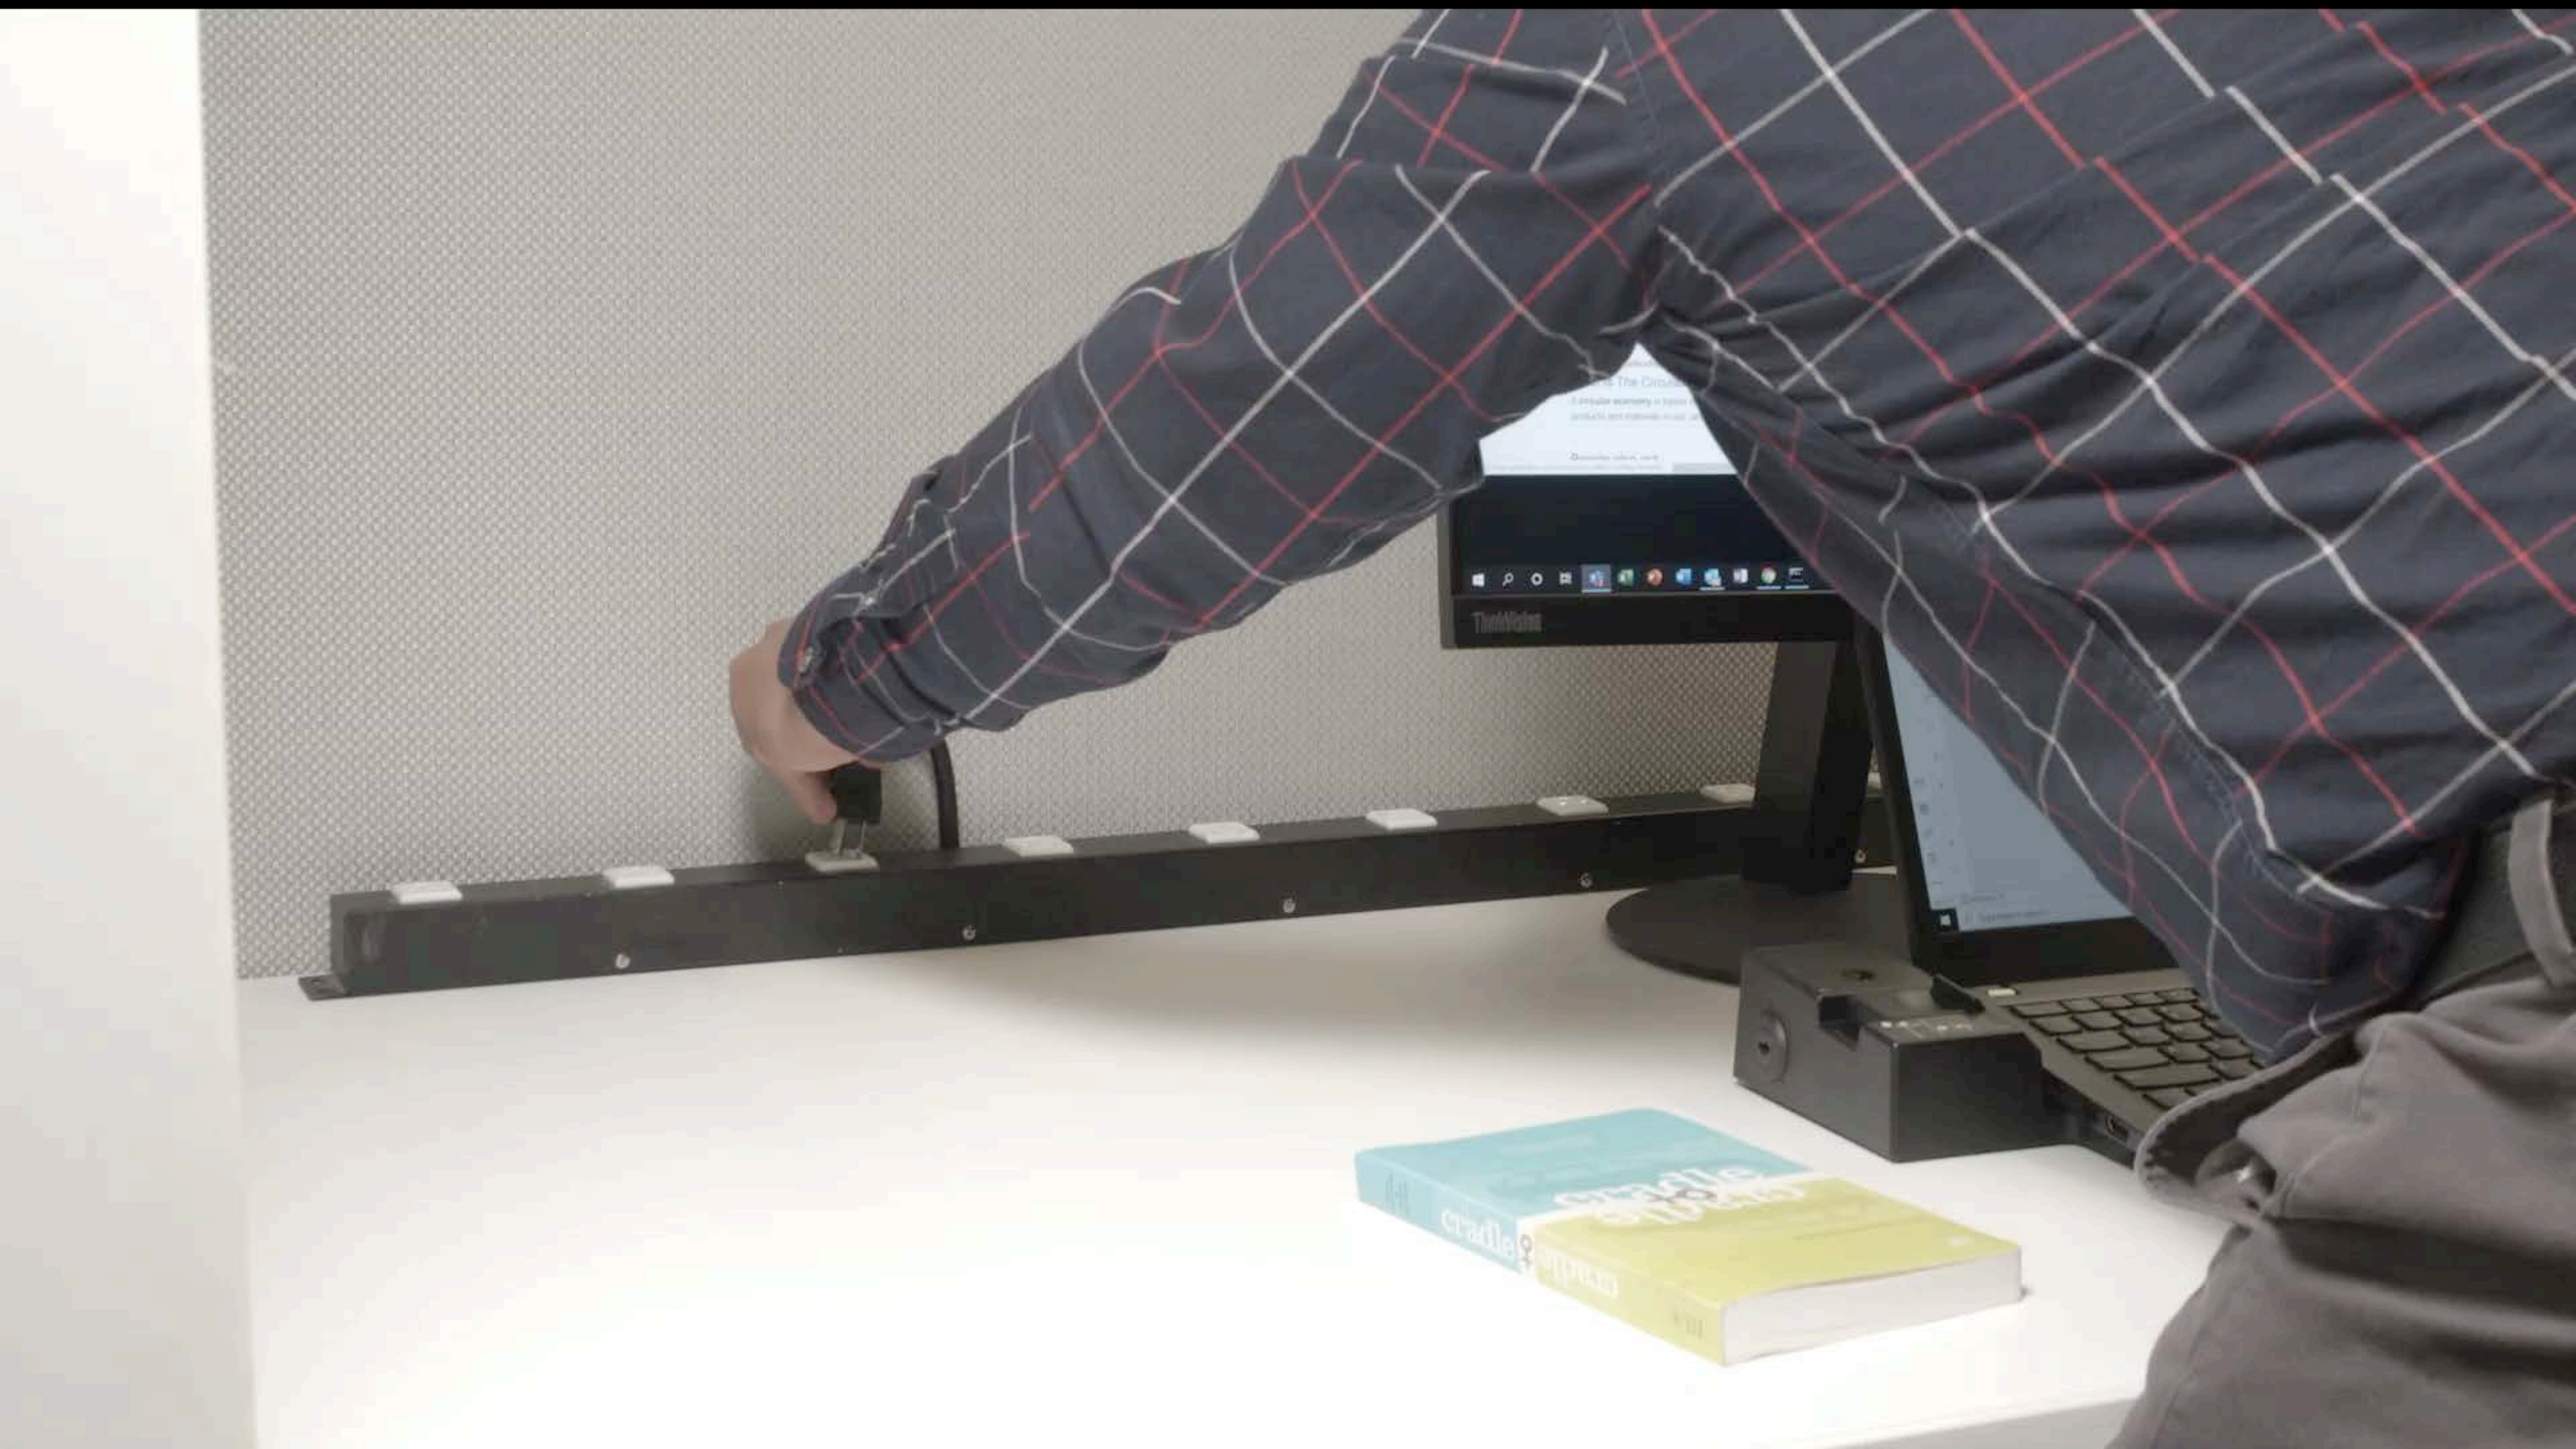
ThinkVision
cradle to cradle

Prerequisites – Before You Start

Provisioning Server	Linux	DevOps	Automation
Our Choices	Ubuntu 19.04	docker, Ansible	Rancher
Deployed Services	DRP (PXE server) DHCP (static IPs)	Digital Rebar docker registry	Rancher

1st Step – PXE Entire Rack

- IPMI power on all the discovered nodes for wipe, install, first boot

  03 min 51 sec Power on -> Disk Erased
  07 min 34 sec Reboot -> Linux Installed (mSATA device)
  04 min 02 sec Reboot -> Boot from installed disk (*)

- (including 2 min of PXE timeout and 30 sec of grub menu)

* caveats for investigation: in some instances, the use of IPMI chassis bootdev at the end of the install, reverts UEFI to traditional BIOS mode and requires additional reboot step, extending the time

Servers

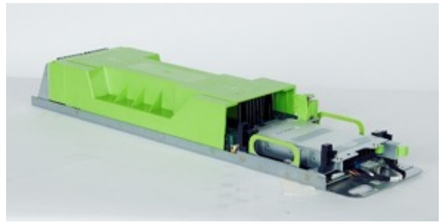

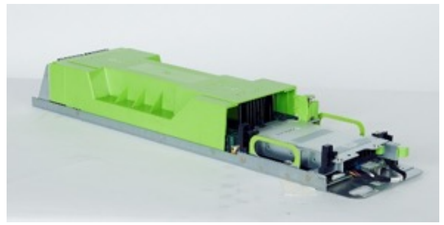

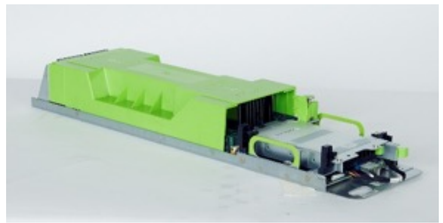

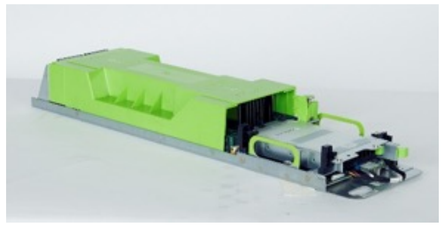

compute
compute
compute
compute
infra
power supply + switch

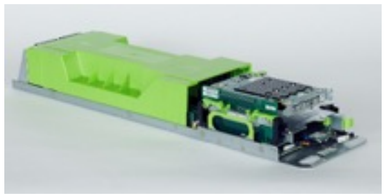

SERVER

SCALE		
compute	**CRIMSON**	2x 12c 256GB 10G
compute	**CRIMSON**	2x 12c 256GB 10G
compute	**CRIMSON**	2x 12c 256GB 10G
compute	**CRIMSON**	2x 12c 256GB 10G
	96 cores	1 TB memory

DRP (PXE server)
DHCP (static IPs)

1st Step – Success ! OSes Deployed

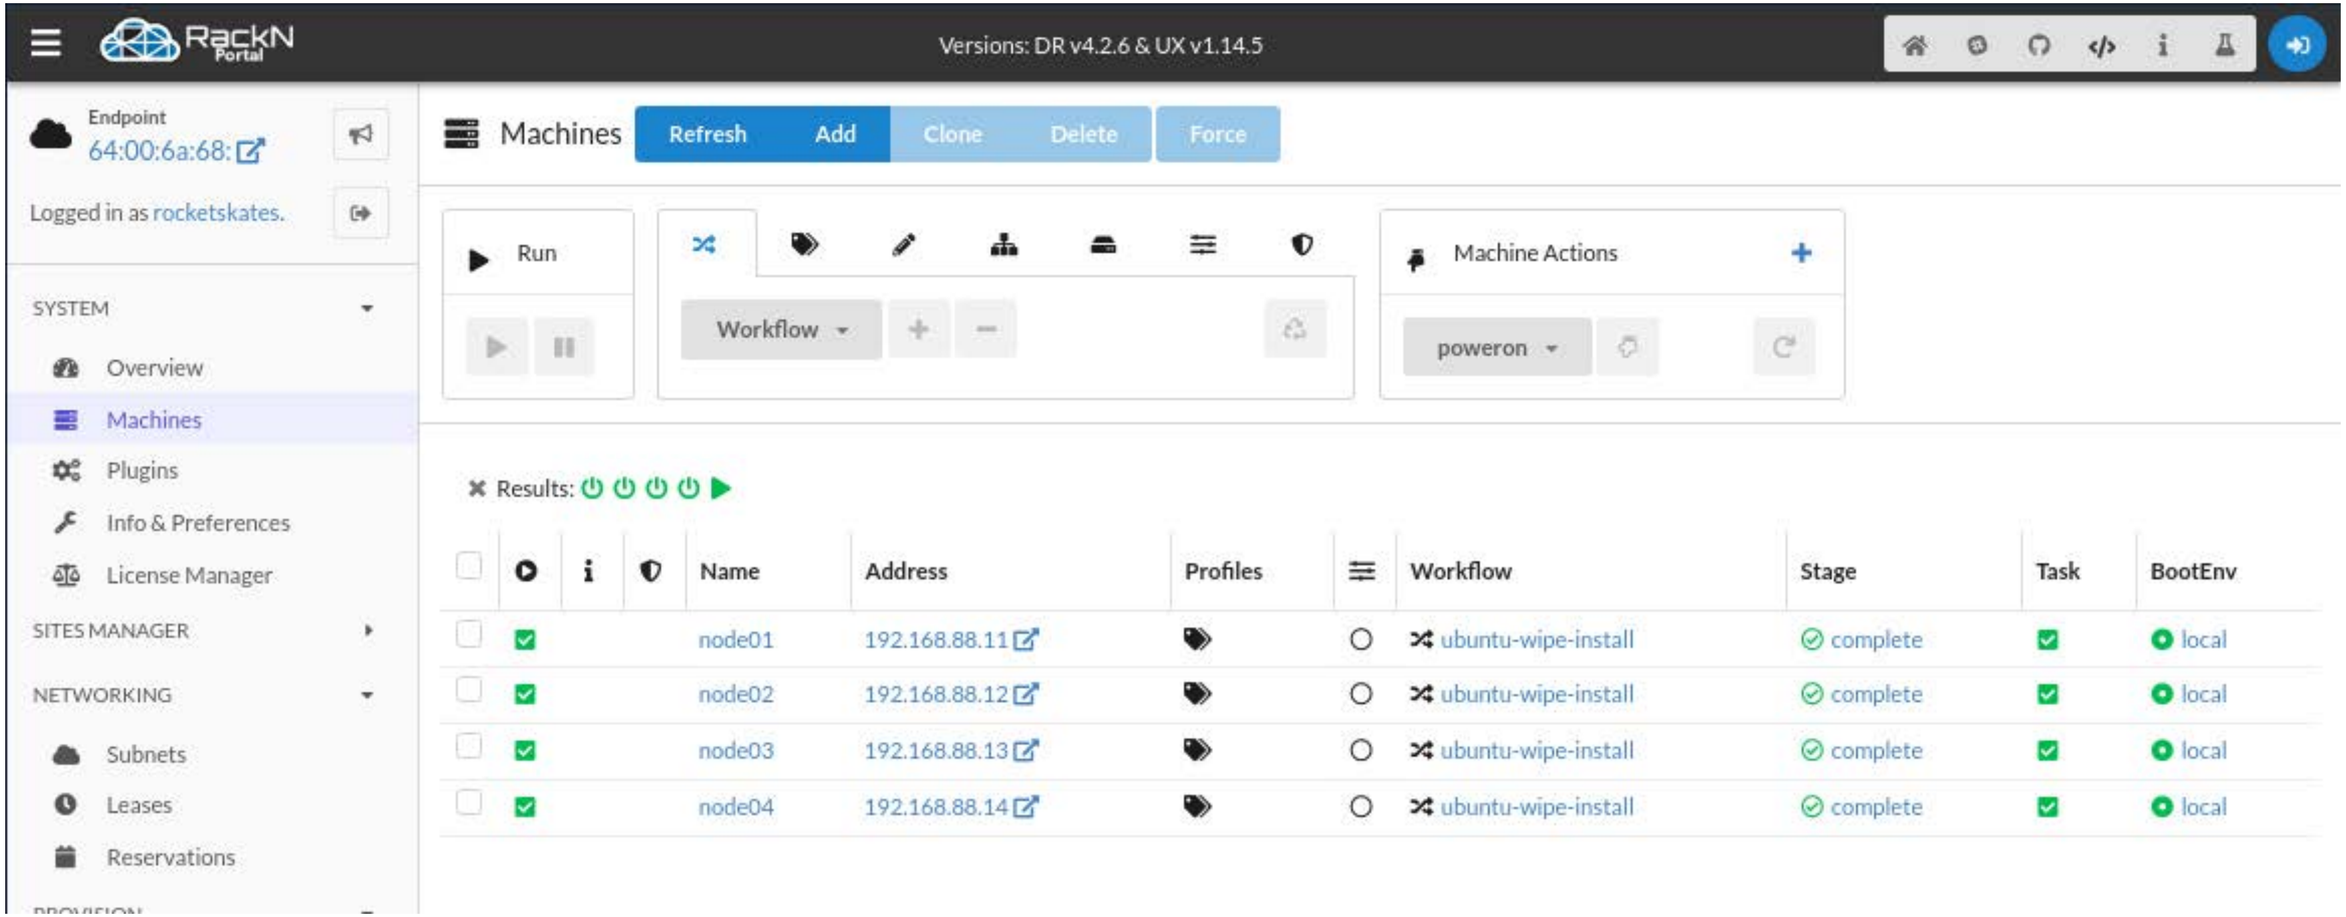

2nd Step – Deploy Rancher

Deploy provisioning server via Ansible

Provisioning server uses Ansible to deploy RKE k8s to all nodes

Provisioning server deploys Rancher and prerequisites into the deployed k8s (*)

* we use one controller/etcd + three workers in a 4-node cluster

2nd Step – Success! – Rancher Deployed

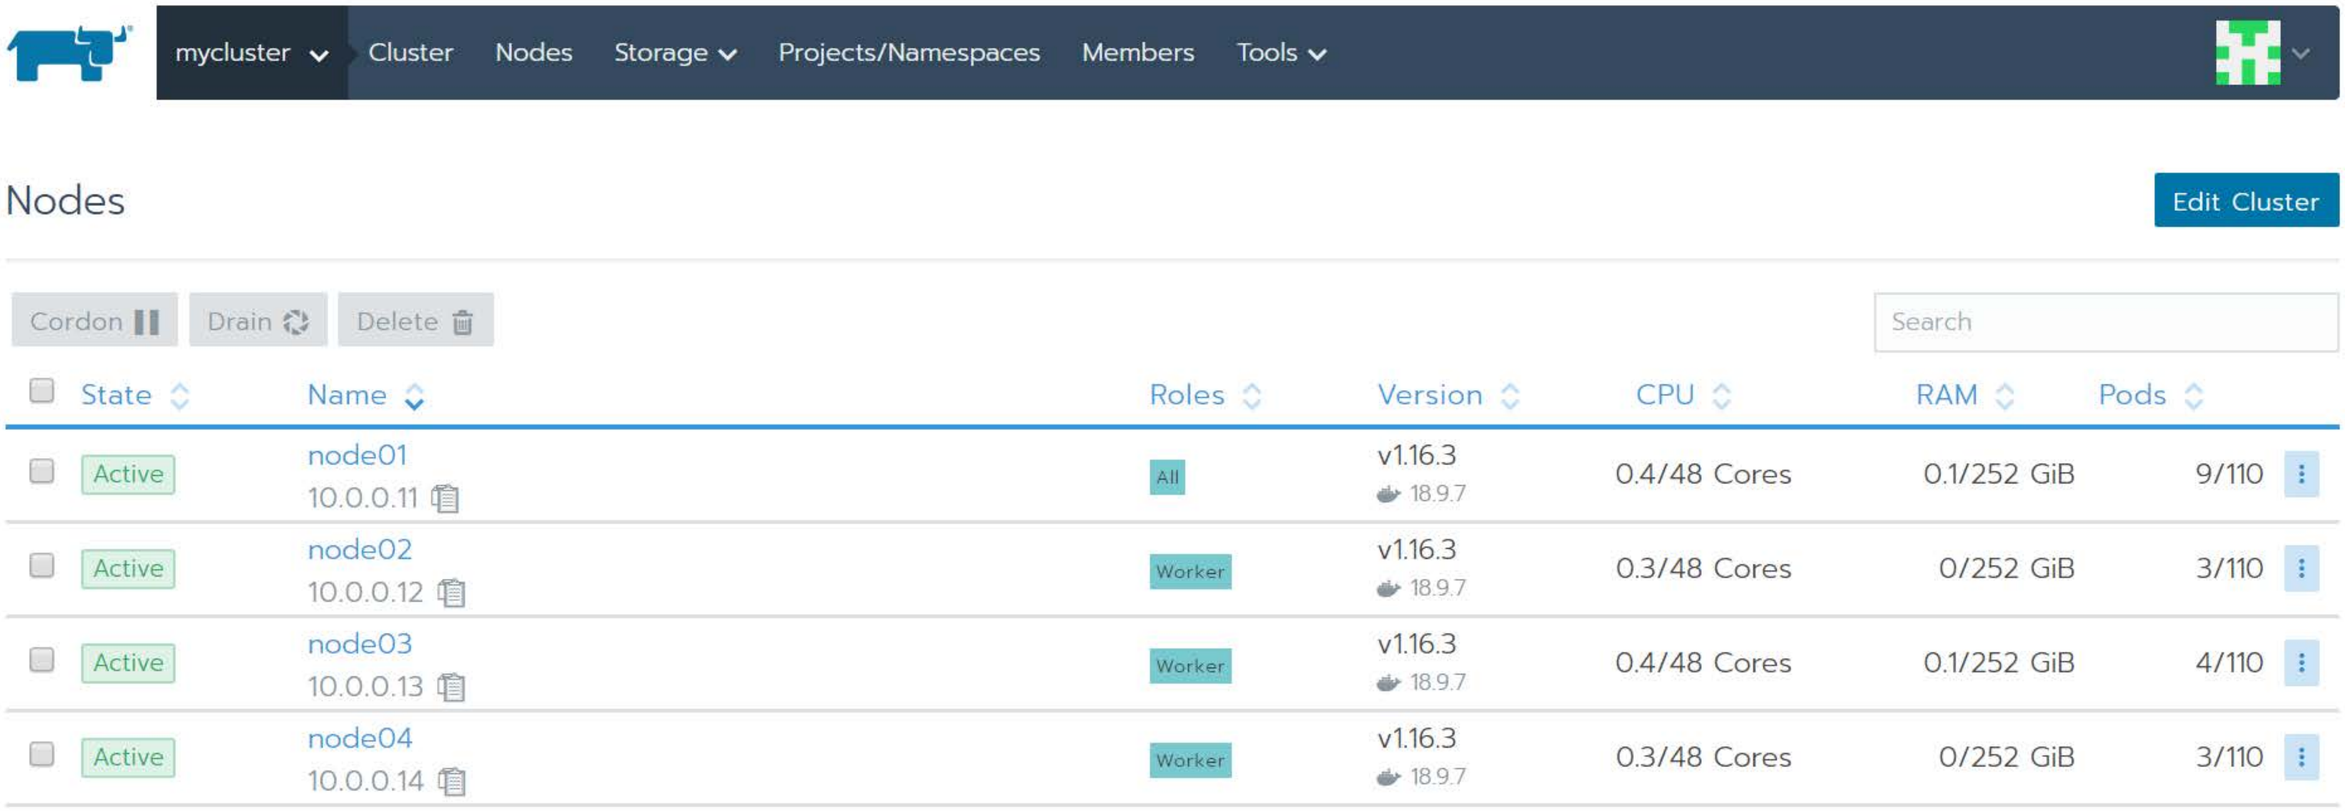

3rd Step – Configure Storage

More details than can be shown here

due to time constraints. See our blog post at link

4th Step – Configure Workloads

Ansible to Rancher to provision pods to exercise cluster

Ansible to Rancher to provision other monitoring tools

Then workloads are visible at the k8s master IP which shows cluster and node health + dashboard + graphs of loads & resource usage

4th Step – Success!

Workloads Configured

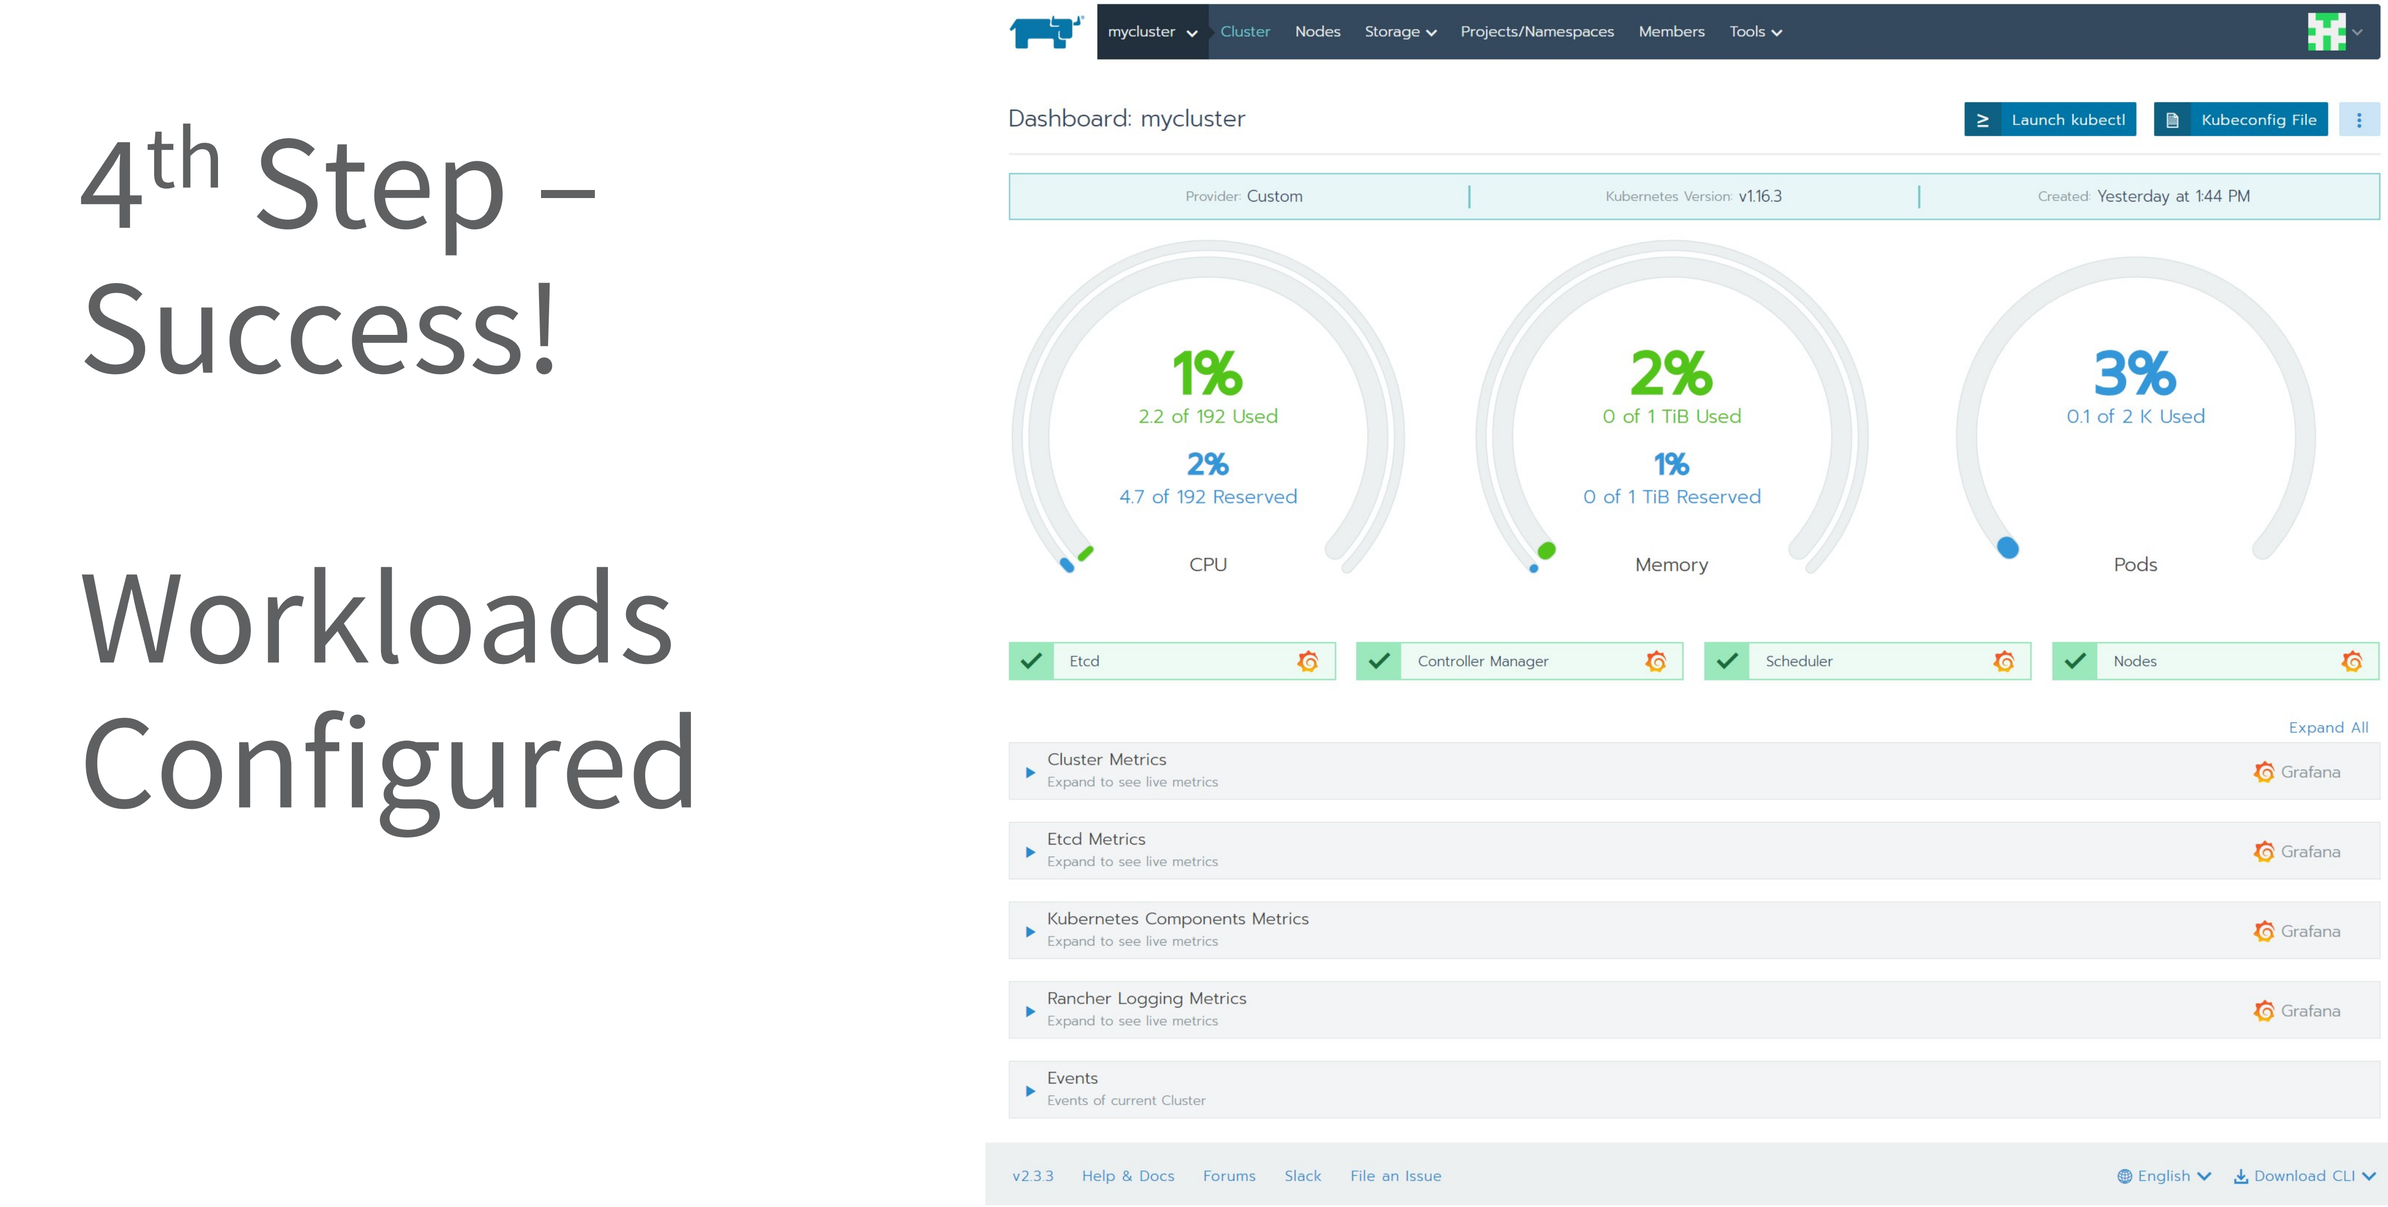

Let’s Review – How Long Did It All Take?

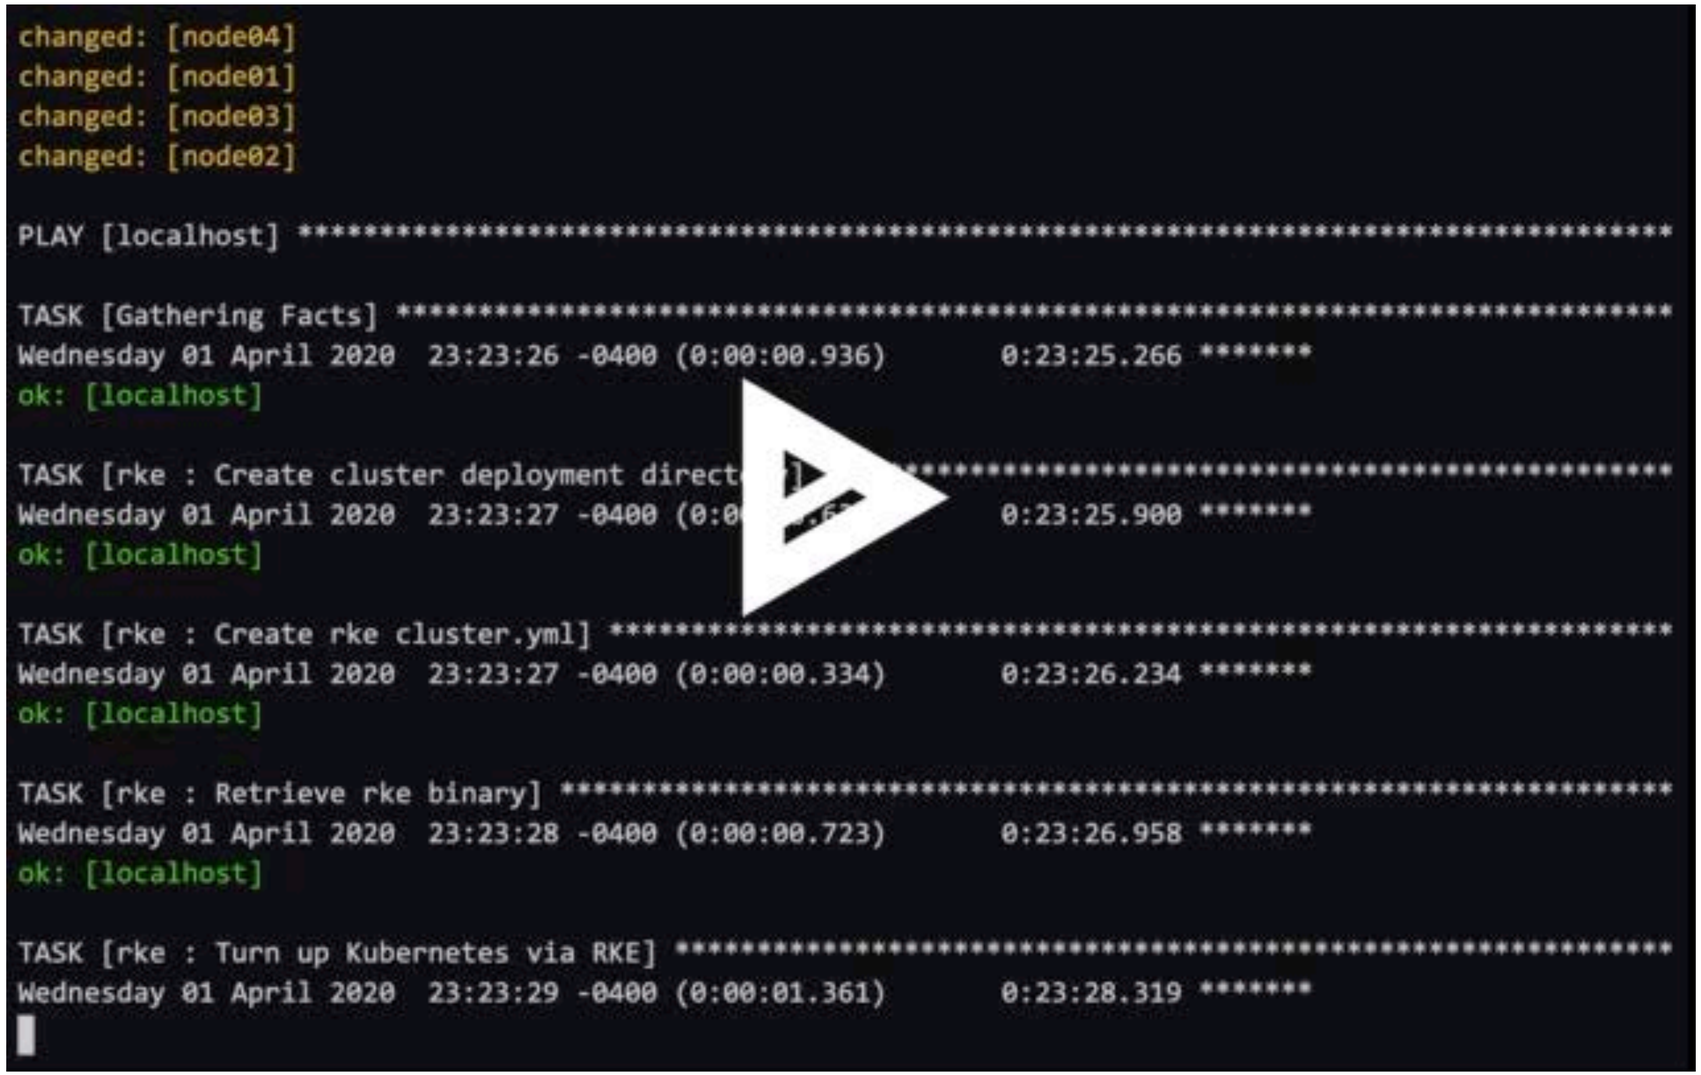

OCP GLOBAL SUMMIT

Case Study

large-scale Sesame customer in the media & entertainment space

compute
compute
compute
compute
infra
power supply + switch

4 workload nodes (PoC platform)

external TOR switches (2x)		
compute	compute	compute
compute	compute	compute
compute	compute	compute
power zone BB		
compute	compute	compute
compute	compute	compute
compute	compute	compute
storage	storage	storage
storage	storage	storage
storage	storage	storage
infra	infra	infra
power zone AA		

27 workload nodes (production platform)

OCP Case Studies

SERVER

SESAME Fast-Start

SESAME for Open Systems

OCP GLOBAL SUMMIT

We Did This With Three Different Stacks

RANCHER	KSPHERE	TALOS	STACKS
20 minutes to full cluster readiness	60 minutes to full cluster readiness	31.5 minutes to full cluster readiness	
1,650 pods	1,800 pods	3,000 pods	SLEEP CONTAINER
500 pods* * docker errors started at 507 pods, crashed at 600 pods	500 pods* * containers began to die, kubelet crashed at 550 pods	500 pods* *health checks began to fail at 500 pods	NGINX CONTAINER

** note that all these tests were done after overriding the default 110 maximum pods per node, as set by Kubernetes

Conclusion

Not quite 10 minutes, as we had hoped…

…but 20 minutes is achievable (on a 4-node cluster), and the gaps to 10 minutes are clear

(for both 4-node and 45-node clusters)

QUESTIONS?

Product/Facility Info

SESAME
Fast-Start

OPEN INSPIRED™

SESAME
for Open Systems

OPEN Compute Project
SOLUTION PROVIDER®

ITRenew Marketplace
https://www.itrenew.com/sesame-solutions

OCP GLOBAL SUMMIT

Call to Action

CHECK US OUT AT THE OCP MARKETPLACE:

https://www.opencompute.org/circular-economy/5/sesame-for-open-systems

QUESTIONS OR COMMENTS, REACH US:

@RiedelAtWork

https://github.com/SesameEngineering

The Power of Hyperscale For All

Fast-Start and Rack-Scale Solutions optimized for your workload, from deskside to data center.

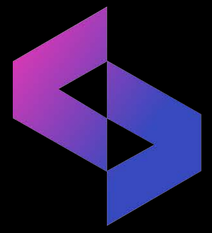

SESAME BY ITRENEW

NO ASSEMBLY. NO GUESSWORK. JUST PLUG THEM IN.

learn more at:

www.ITRenew.com/Sesame

OCP GLOBAL SUMMIT

Open for All.

OCP
GLOBAL
SUMMIT

MARCH 4 & 5, 2020 | SAN JOSE, CA