How To Run Splunk As A Docker Image?

Marc Chéné

IT Markets Product Manager , Splunk

Denis Gladkikh

Principal Dev Engineer (aka outcoldman), Splunk

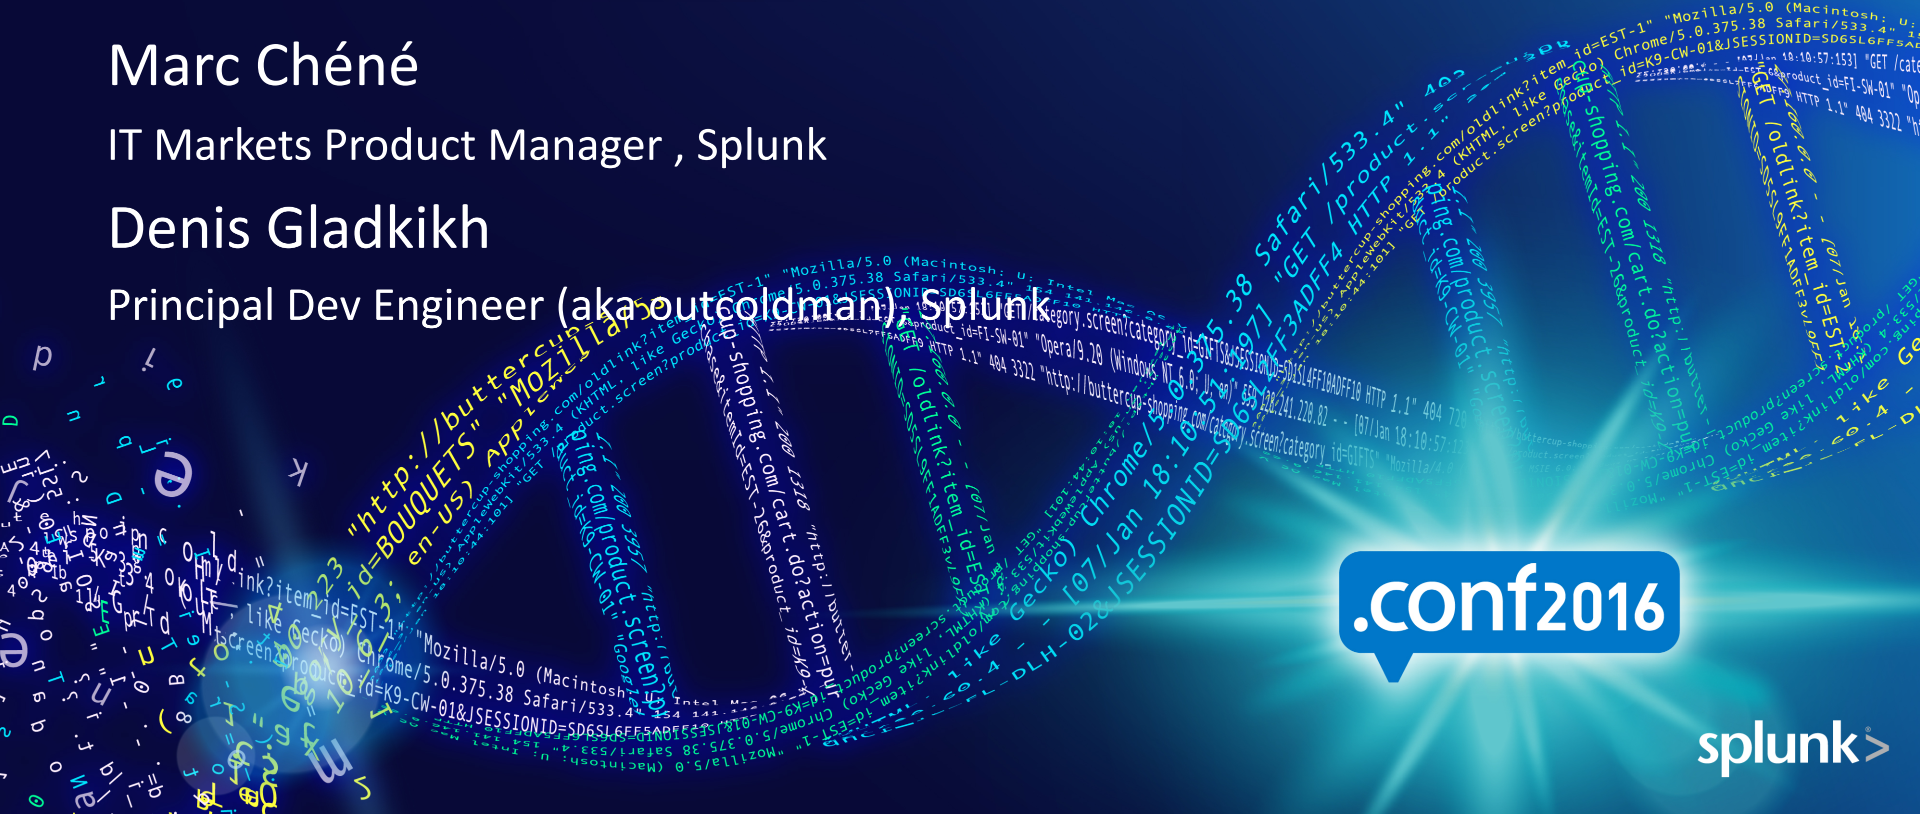

.conf2016

splunk>

Disclaimer

During the course of this presentation, we may make forward looking statements regarding future events or the expected performance of the company. We caution you that such statements reflect our current expectations and estimates based on factors currently known to us and that actual events or results could differ materially. For important factors that may cause actual results to differ from those contained in our forward-looking statements, please review our filings with the SEC. The forward-looking statements made in the this presentation are being made as of the time and date of its live presentation. If reviewed after its live presentation, this presentation may not contain current or accurate information. We do not assume any obligation to update any forward looking statements we may make. In addition, any information about our roadmap outlines our general product direction and is subject to change at any time without notice. It is for informational purposes only and shall not, be incorporated into any contract or other commitment. Splunk undertakes no obligation either to develop the features or functionality described or to include any such feature or functionality in a future release.

Who we are?

- Marc Chéné
- Denis Gladkikh

Agenda

- What is Docker?
- Why run Splunk in Docker?
- Demo Scenarios
  1. Setup Splunk Cluster in Docker
  2. Scaling Up Splunk in Docker
  3. Cluster Upgrade: 6.4.1 to 6.4.2
- Guidance & Best Practices

What is Docker?

First, a bit about containers…

Docker, in one Slide

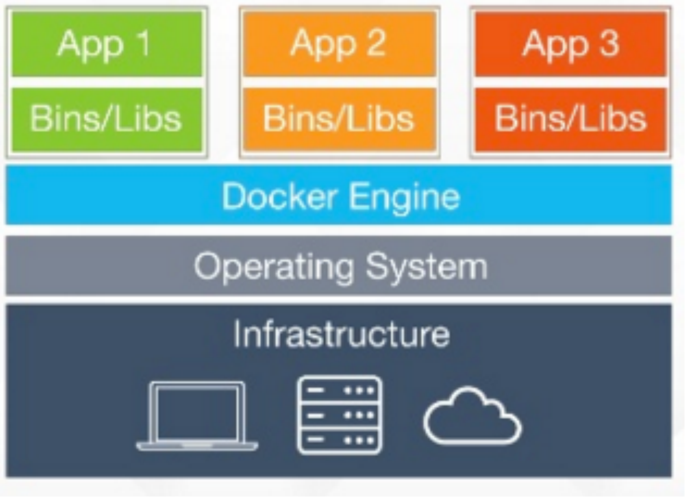

- Build - Ship - Run your applications
  - ”Infrastructure as code”
  - Enables microservices architectures
  - Portable – Enables Cloud Migration
- Open Source and Community Minded
  - Docker Engine is Open Source
  - Thousands of apps can be “pulled” in Dockerhub
  - Your developers LOVE Docker

Docker – It's not Virtualization

- VMs – focus on OS
- Docker – focus on applications
- Docker – lightweight and FAST
- NOT mutually exclusive with VMs

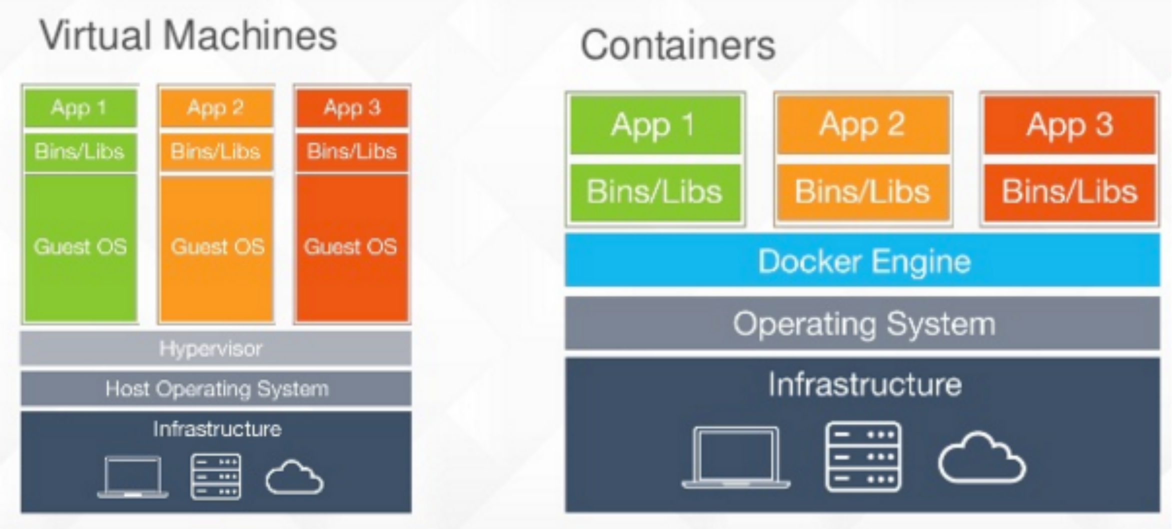

Docker – it’s a big deal

460K
Dockerized Apps
3,000% Growth in 2 years

Open Source driven ecosystem

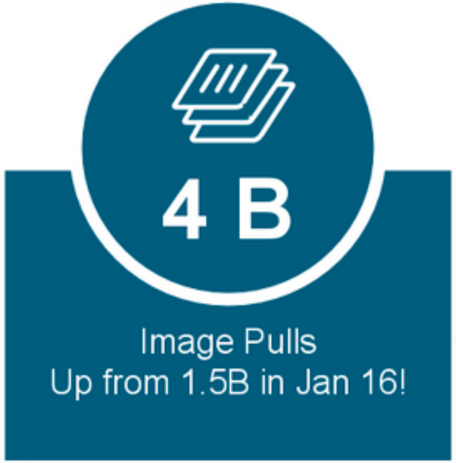

Massive increase in adoption…

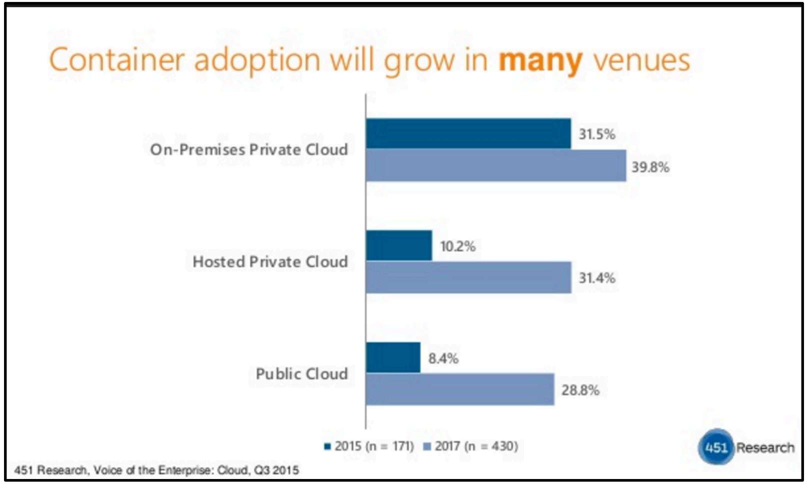

451 Research, Voice of the Enterprise: Cloud, Q3 2015

…But the growth is just getting started

3 Primary Container Use Cases

- Application Modernization
- DevOps – New Apps
- Cloud Migration

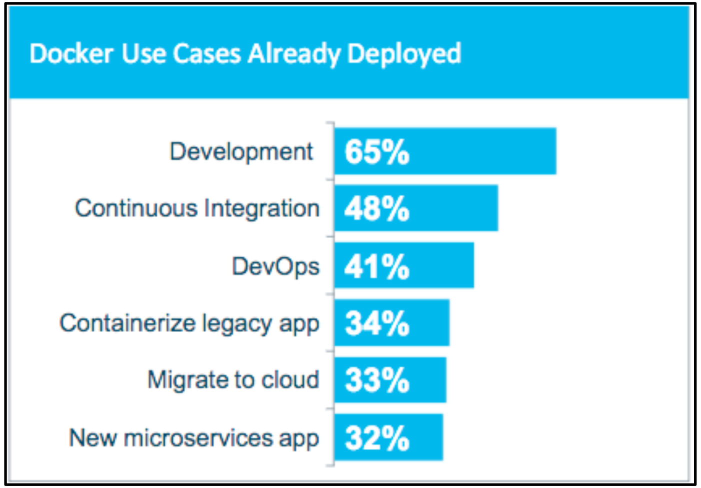

Why run Splunk in Docker?

- Goals & Benefits
- Splunk Reference Architecture
- Target Docker Environment

Goals & Benefits

- Reduce Management Costs
- Time to Value
- High Available
- Reduce time to Upgrade
- Simplified Rollback
- Standard Configurations
- Easier to Support

Splunk Reference Architecture in Docker

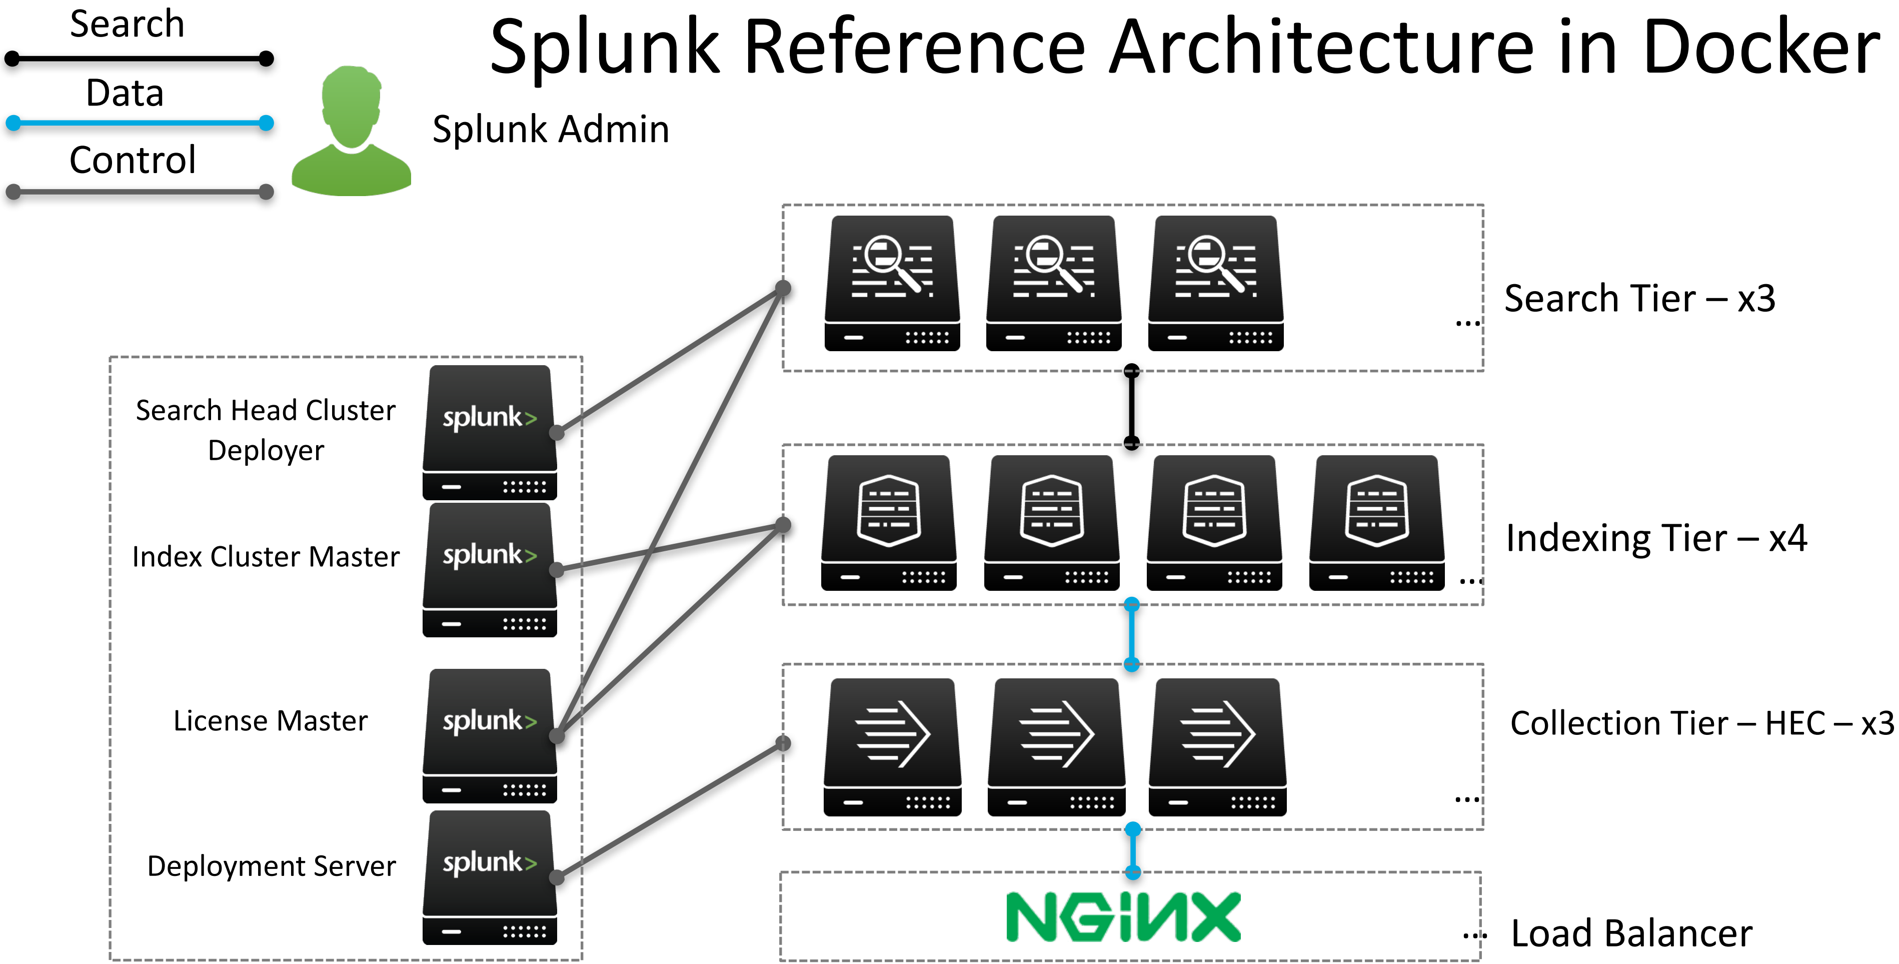

Target Environment Architecture

- Docker v1.12
- Docker for AWS
  – https://beta.docker.com/docs/aws/#upgrading-docker-and-changing-instance-sizes
- Docker SWARM

Splunk Sizing Guidelines in AWS

- Storage
  - EBS
    - High Available
    - Reliable
    - Grow up to 16TB
  - EBS General Purpose (SSD): consistent performance
  - EBS Provisioned IOPS (SSD): consistent performance up to 4K IOPS

  * EBS volumes can be deployed in a RAID architecture
- Compute Requirements per Splunk Components
  - 4 vCPUs
  - 8 GB RAM

Splunk Enterprise deployment on AWS

Workload = Searching + Indexing

Storage

- Ephemeral or EBS
- Data Retention Dependent

Compute

- Best Available

Archiving

- S3

Best Practices for Sizing

[Splunk on AWS Tech Brief]

[Splunk Cloudformation Templates]

[Splunk Admin Docs]

Search Heads (8+ users)

c4.4xlarge	16 vCPU, 30 GB RAM
c4.8xlarge	36 vCPU, 60 GB RAM

Indexers (50-250GB/day/indexer)

c4.4xlarge	16 vCPU, 30 GB RAM
d2.4xlarge	16 vCPU, 122 GB RAM
c4.8xlarge	36 vCPU, 60 GB RAM

CloudFormation Templates

Consistent, repeatable deployments for Splunk

Abstract away details of configuring distributed Splunk

Extensible and **customizable** to fit any need

Cloudformation Templates On GitHub

What is docker for AWS?

- Announced at dockercon16
- Autoscaling groups
- CloudFormation Updates

Delivering Splunk as a Container Image

- Splunk container images
  - Splunk Enterprise 6.4.1
  - Splunk Universal Forwarder 6.4.1
- Includes configuration and Docker Add-On for container monitoring out-of-the-box
- Certified image
- Coming soon to the Docker Store (http://store.docker.com)

```
docker run splunk/enterprise:6.4.1-monitor
docker run splunk/universalforwarder:6.4.1-monitor
```

Demo Apps

- Splunk for AWS
- Splunk Docker App
  - UF for logs
  - cAdvisor for metrics

Demo Time!

- Setup Splunk Cluster in Docker in AWS

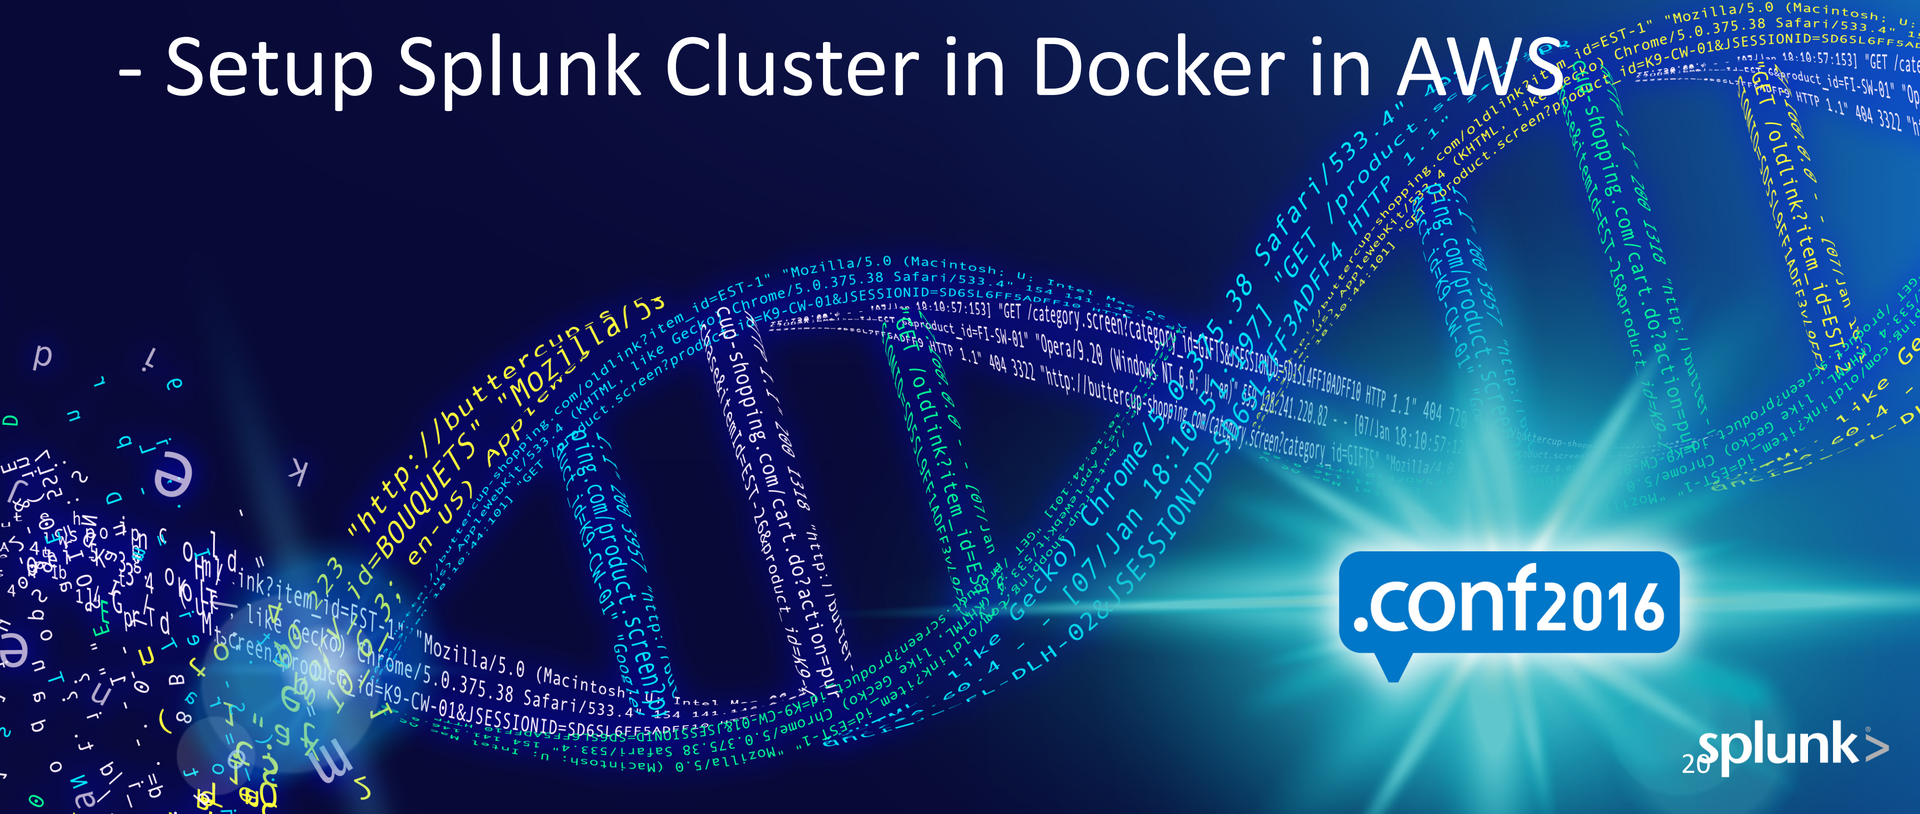

Splunk Scale Up

- Goal
  - Scale up the Search Heads by 3 → Total: 5
  - Scale up the Indexers by 4 → Total: 8
  - Scale up the Collection layer by 2 → Total: 4

Splunk in Docker – Scale Up

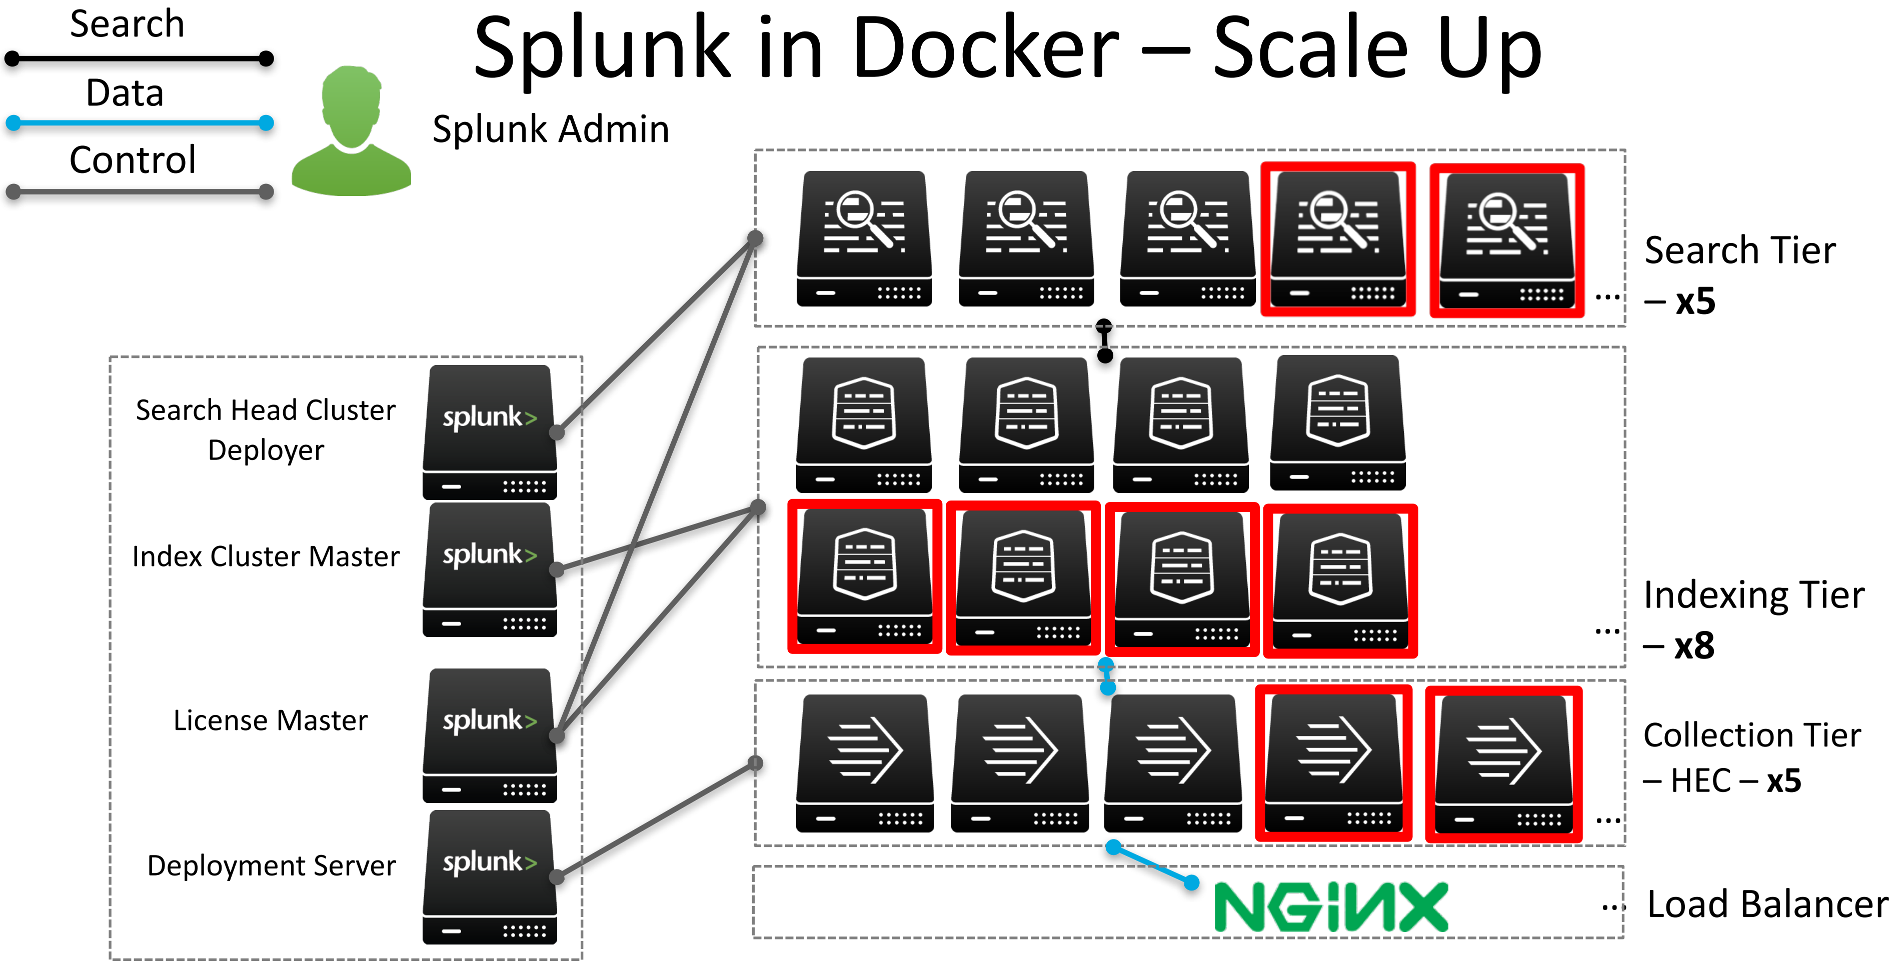

Demo Time!

- Scale Up Splunk Cluster

Splunk Upgrade

- Order of Upgrade
  - Search Head
  - License Manager
  - Cluster Master
  - Indexers
- Recommendations
  - Backup your configurations
  - Backup your data
- Goal
  - Upgrade 6.4.1 to 6.4.2

Demo Time!

- Upgrade Splunk Cluster

Guidance and Best Practices

Guidance and Best Practices

- Understand the sizing factors
  - How much data (raw sizes)? Daily, Peak, Retained (archive size), Future
  - How much searching? Use Cases, # of people, Apps
  - Jobs: Summarization, alerting, reporting
- Standard Operation Procedures
- Data volume
- Search volume

What are the Pitfalls of docker?

- TBD

What Now?

Related breakout sessions and activities…

- Architecting Splunk for High Availability and Disaster Recovery, Session ID: 74762
- Observations and Recommendations on Splunk Performance, Session ID: 74765
- Monitoring and Troubleshooting Docker across Cloud and On-Prem Environments, Session ID: IT88095
- Splunking AWS for End-to-end Visibility, Session ID: 87942
  - Track: Splunk Platform for Operational Intelligence

Call to Action…

```
# 1. Come visit us at our booth
docker run splunk/visitourbooth
visitourbooth_1 | Booth IT Markets

# 2. Try out our docker images in Docker Store
docker run splunk/enterprise:6.4.1-monitor
docker run splunk/universalforwarder:6.4.1-monitor

# 3. Demos will all be available on GitHub under Splunk!
git clone https://github.com/splunk/docker-gettingstarted-conf2016-
sf88089.git

# 4. Visit our site to learn more about containers
curl http://www.splunk.com/containers
```

Resources

Splunk Education

- Architecting and Deploying Splunk 6.4 – Virtual

Splunk Docs

- Upgrade Guide, http://docs.splunk.com/Documentation/Splunk/6.4.2/installation/Upgradeto6.4onUNIX
- Capacity Planning Manual, http://docs.splunk.com/Documentation/Splunk/6.4.1/Capacity/Referencehardware
- DEPLOYING SPLUNK® ENTERPRISE ON AMAZON WEB SERVICES, http://www.splunk.com/pdfs/technical-briefs/deploying-splunk-enterprise-on-amazon-web-services-technical-brief.pdf

Docker

- Docker for AWS, https://beta.docker.com/docs/
- Docker Store, http://store.docker.com

THANK YOU

.conf2016

splunk>