splunk> .conf2017

From Zero To 100 In 100 Days

or "How Quickly Can You Drive Splunk Adoption?"

Tom Gerhard | Fellow, priceline.com

Vidhya Ramachandran | Principal Engineer, priceline.com

September 2017 | Washington, DC

Forward-Looking Statements

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 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.

splunk> .conf2017

About Us – priceline.com

priceline.com®

- priceline.com is part of the Priceline Group. The Priceline Group is the world leader in online travel & related services. The Priceline Group's mission is to help people experience the world.

Booking.com priceline.com KAYAK agoda.com Rentalcars.com OpenTable

splunk> .conf2017

How We Got Here…

The problem we set out to solve

- We had a collection of bespoke monitoring systems, evolved over 18+ years; weren't investing enough to get the full value from them.
- Separate systems meant that we sometimes had difficulty seeing data across applications or application layers in the same context

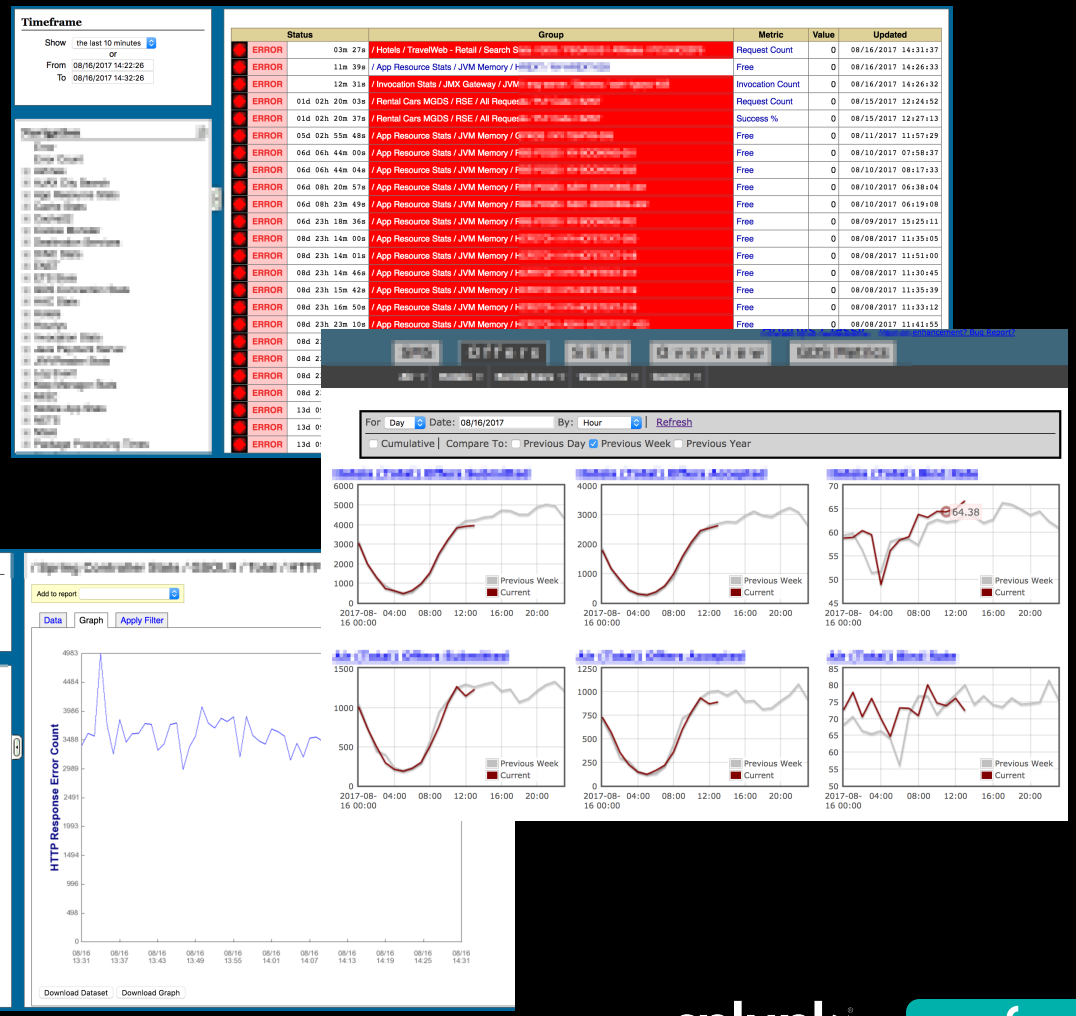

priceline.com®

splunk> .conf2017

Legacy Architecture

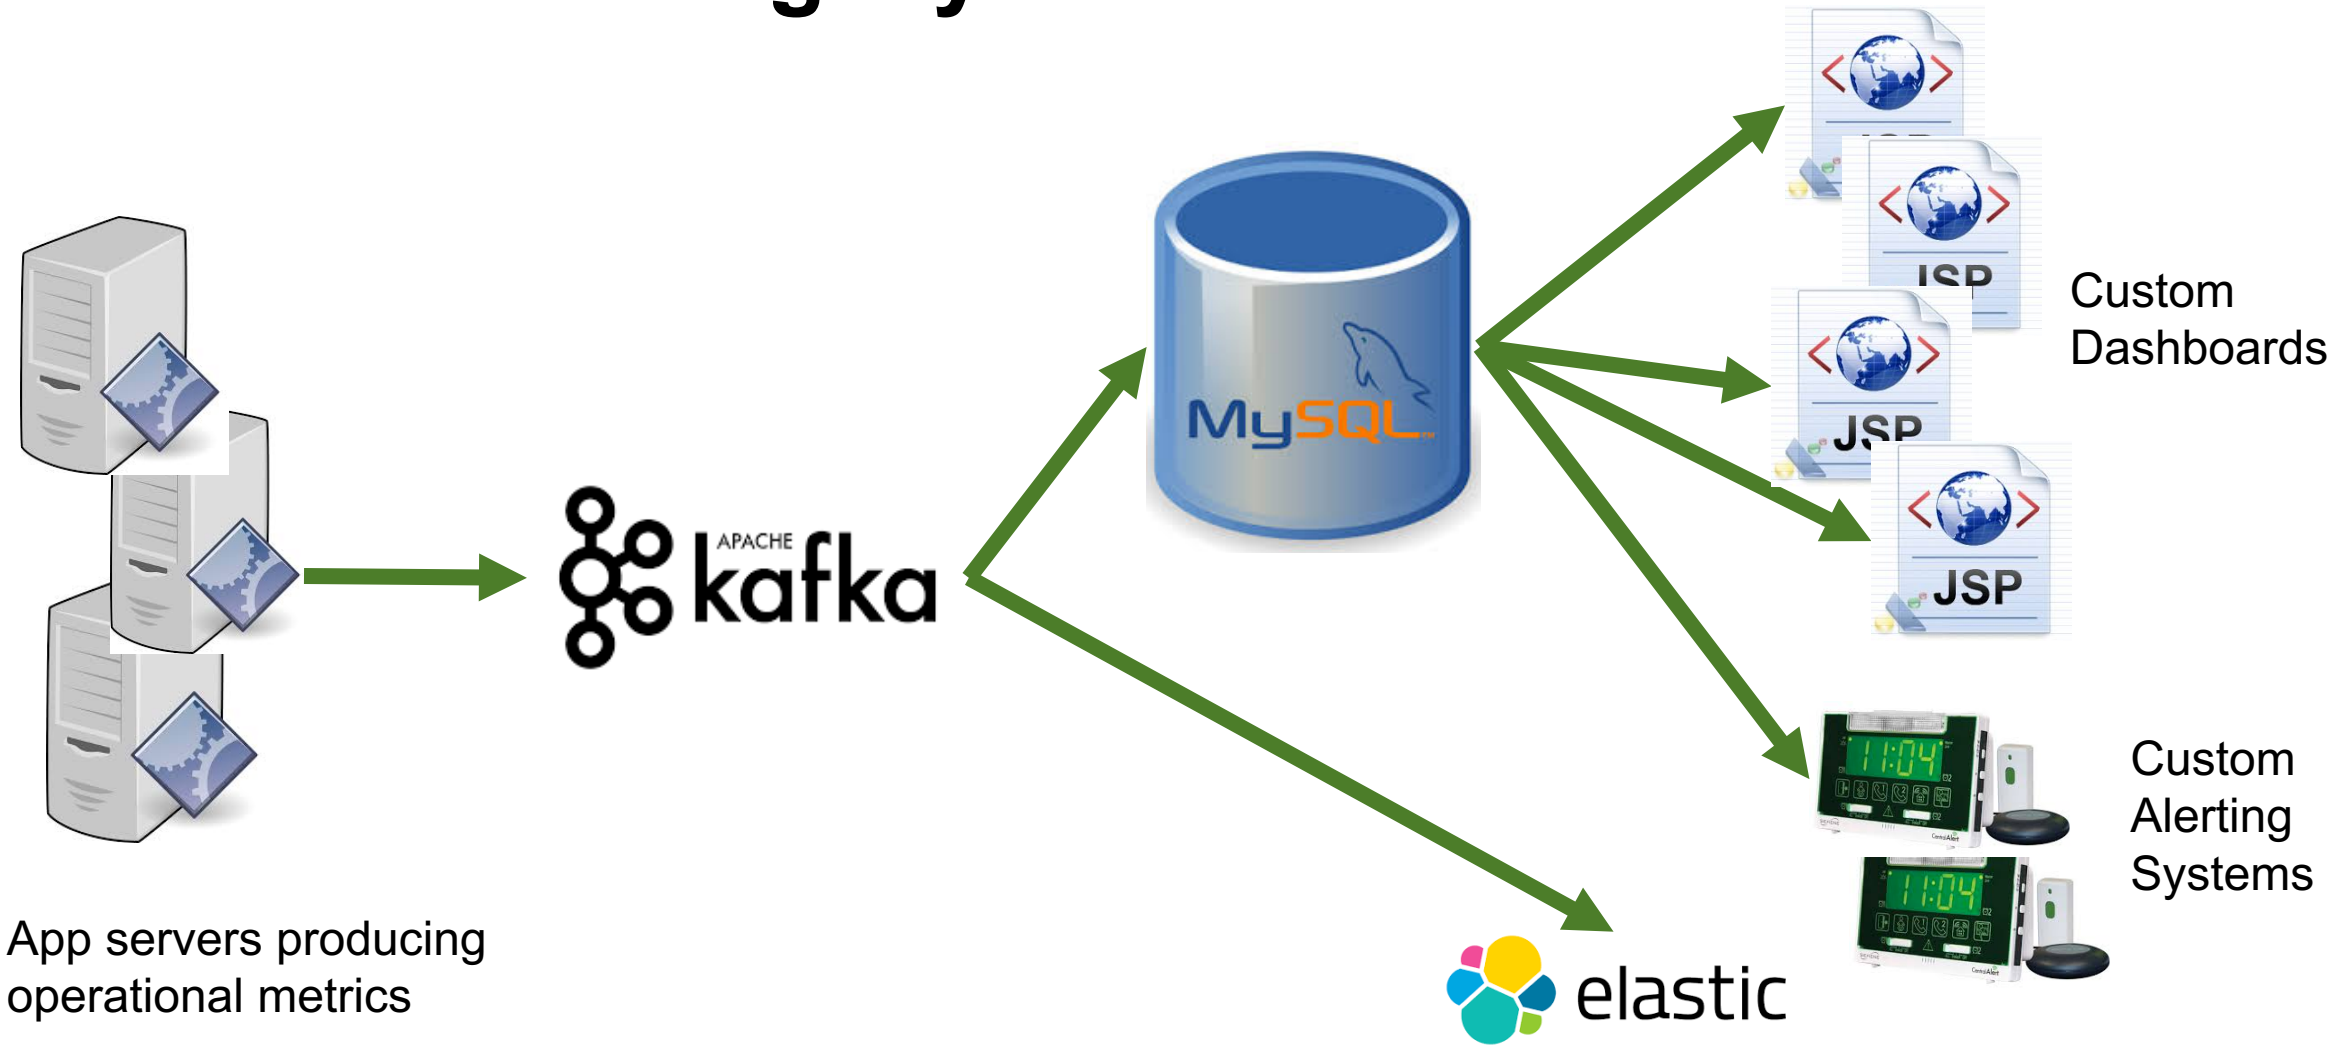

priceline.com

splunk> .conf2017

Our timeline

Team Commissioned
April 2016

POC
May-June

Purchase
July

Launch
Sept 23, 2016

Day 100
Jan 1, 2017

Day 200
Apr 11, 2017

priceline.com®

splunk> .conf2017

Our Timeline

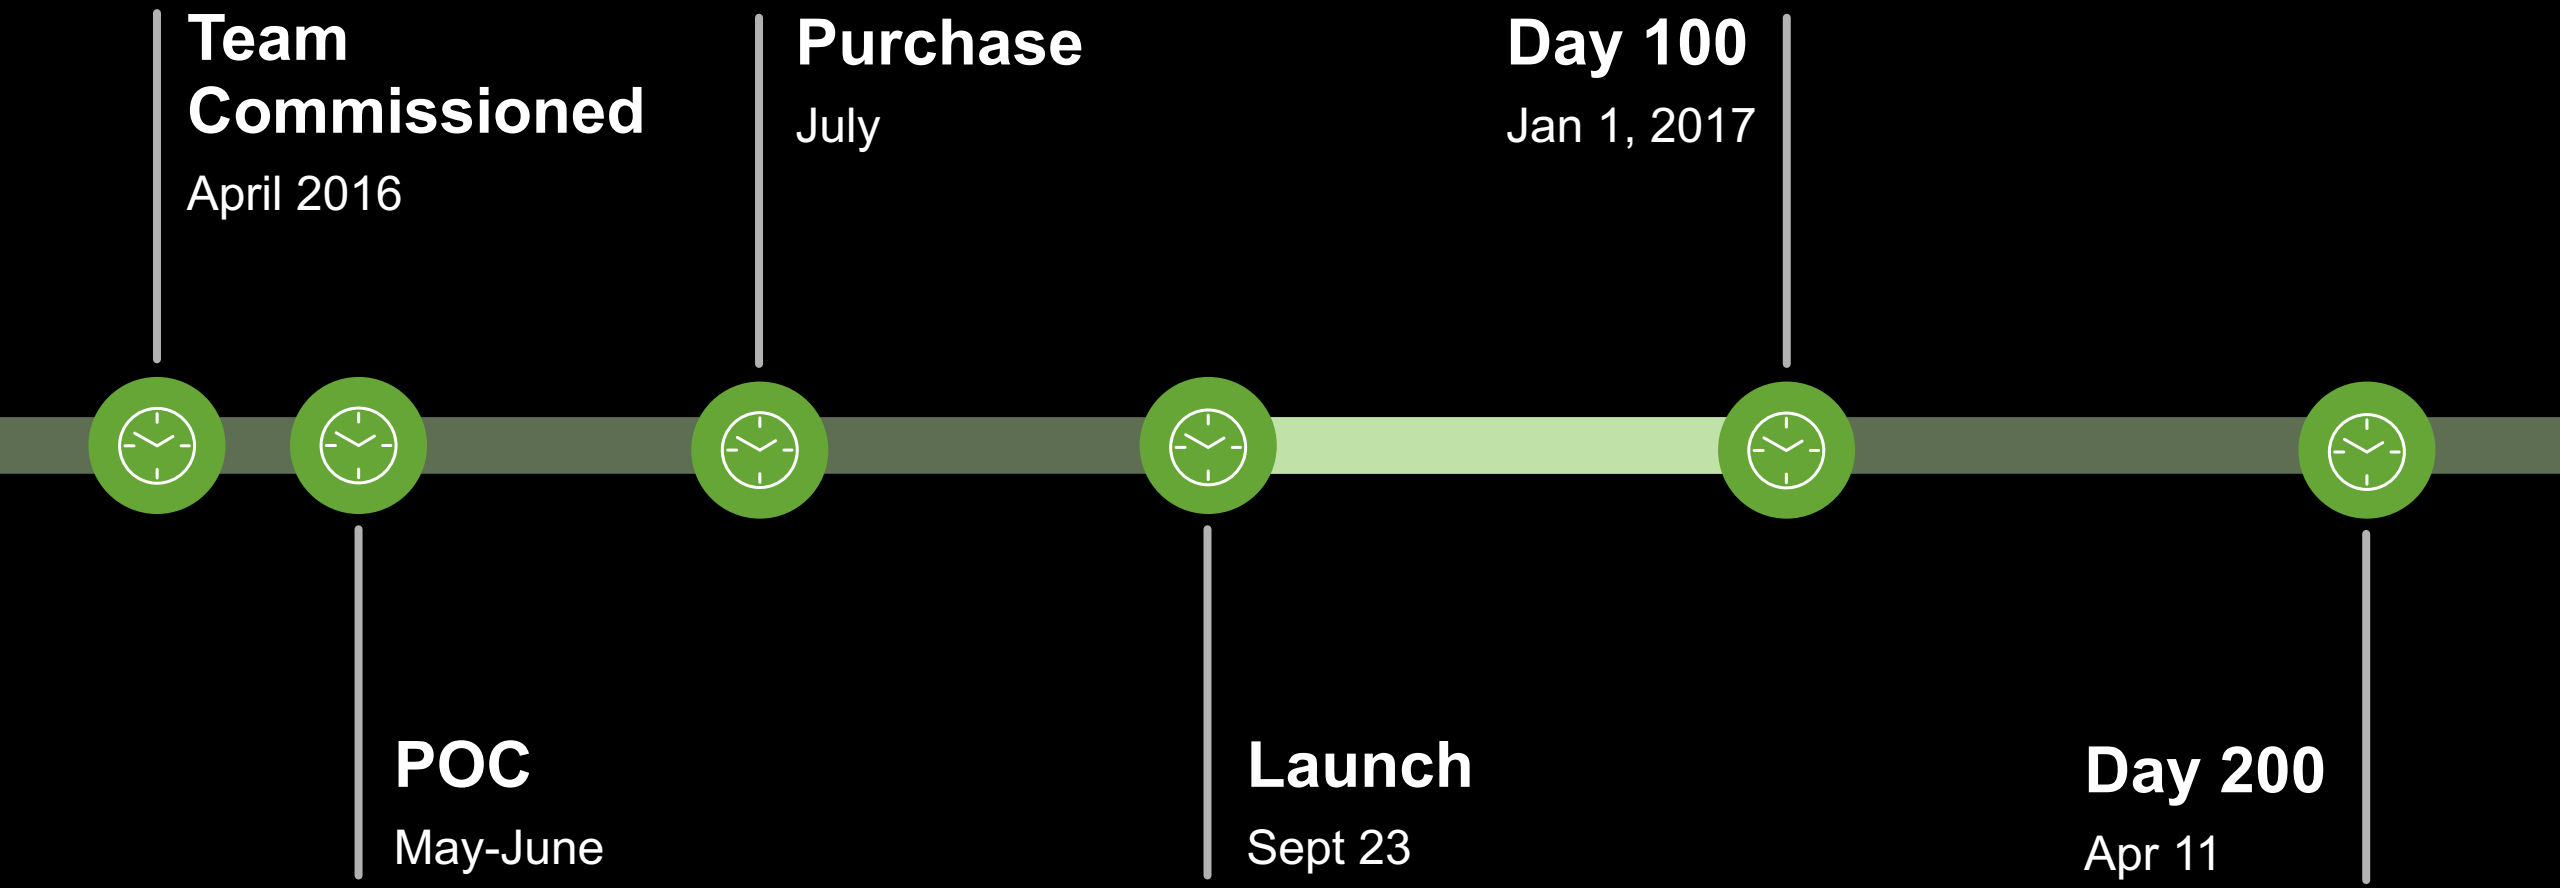

priceline.com®

.conf2017

Why Splunk?

Narrowing the field

- April 2016: Logging & Monitoring team was established to evaluate possible solutions
- Evaluated several open source options, then Splunk
  - Rich ecosystem of 3rd party apps was important - we could leverage vendor and other users contributions, especially in NetOps and SecOps
    - Crowdstrike, F5 BIG-IP, Cisco, *nix, Windows, Palo Alto, Catchpoint, AWS, GCP
  - Flexible data ingestion architecture – HTTP/REST, log scraping, dedicated apps

POC – 60 days

- Ingested live data
- Implemented useful dashboards and alerts
- Started to envision a full rollout

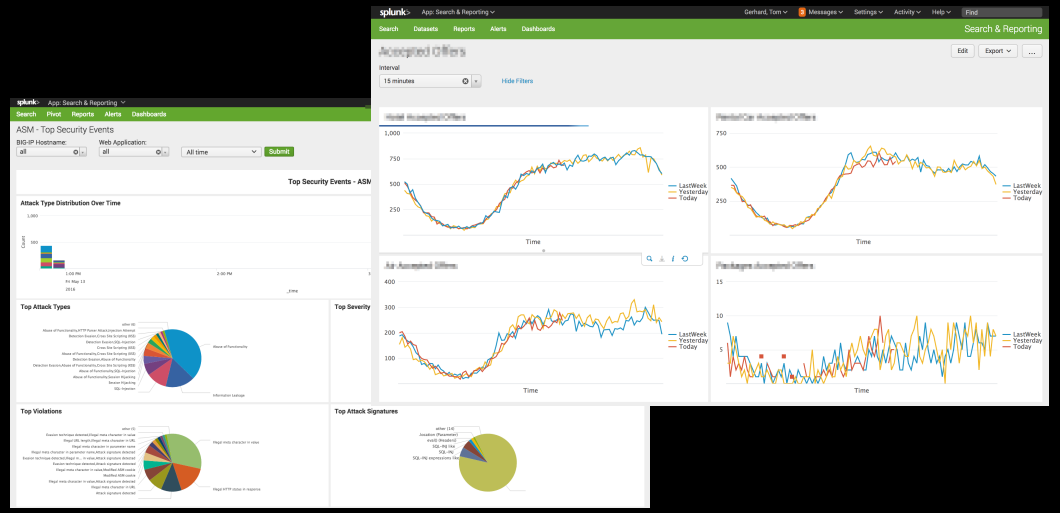

priceline.com®

splunk> .conf2017

Splunk

- Signed 2.5TB contract with Splunk in late July

priceline.com®

splunk> .conf2017

Launch Prep

- System architecture and data ingestion planning ~ 6 weeks
  - Used our POC instance for testing
  - PS engagement to plan cluster configuration
  - Kept index model simple / organized by product line and retention
- Hardware build ~ 2 weeks

Launch Architecture

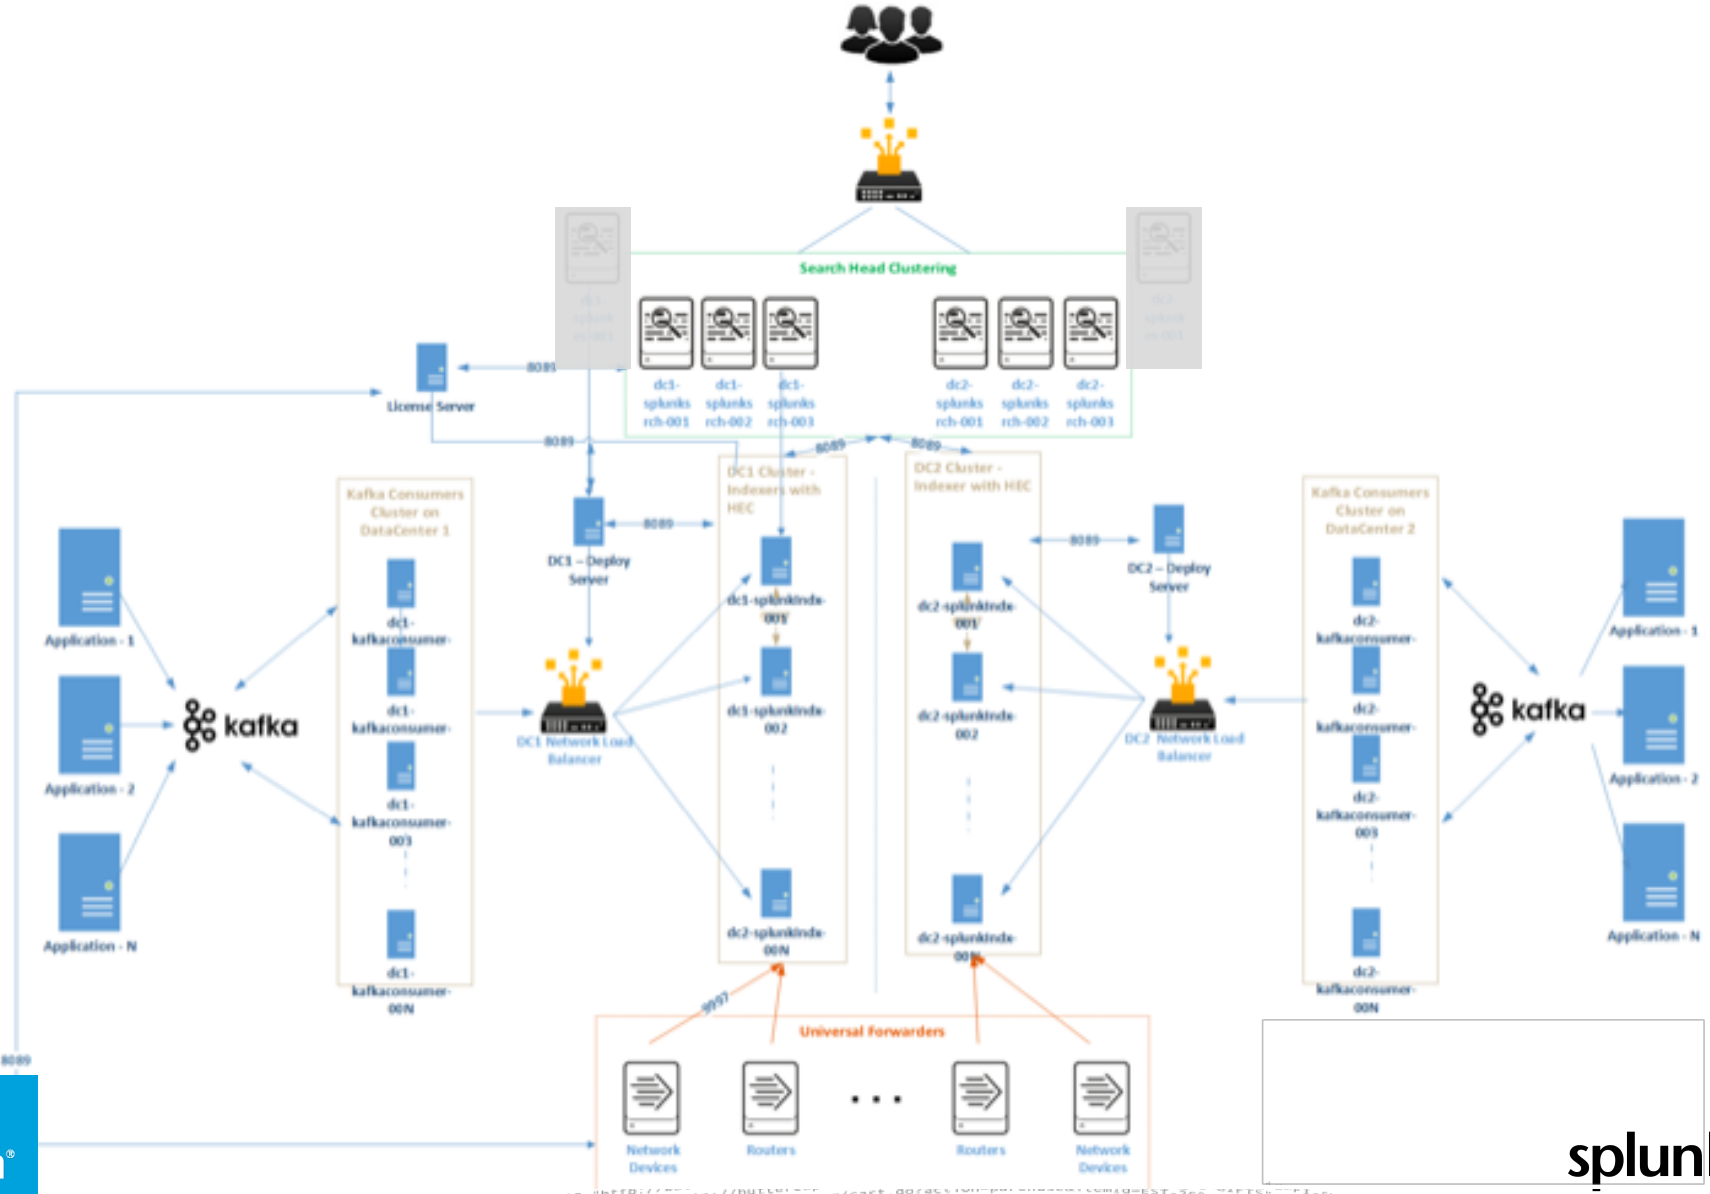

priceline.com®

splunk> .conf2017

Going To Disney World!

- Production cluster opened for business September 23

priceline.com®

splunk> .conf2017

Getting Started

Day 1…

- Enabled all data sources that were available
- Moved key knowledge objects from POC to production
- Trained seed teams – SRO, SRE, NetOps, first application team
- Started to inject Splunk into the conversation

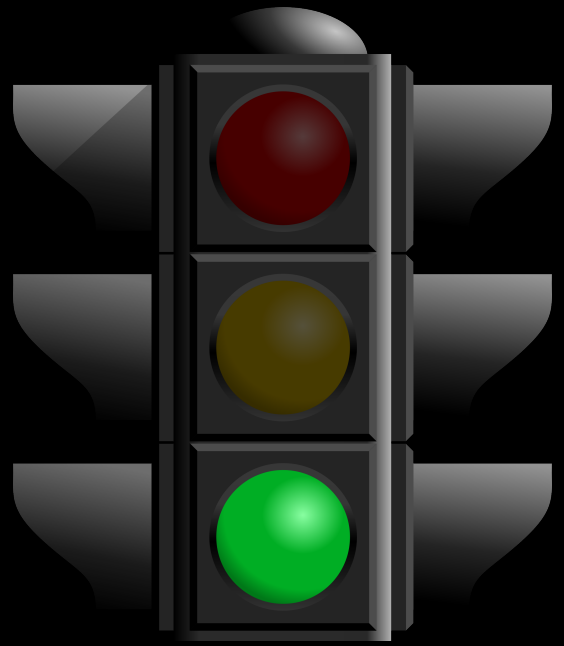

priceline.com®

splunk> .conf2017

Key Takeaways

For viral adoption

1. Make data available early and often
2. Train!
3. Build Community
4. Remove as many barriers as possible

Making Data Available

If you ingest it, users will come…

- On day 1, ingested about 1.5TB/day
- Added data sources that we knew were required as quickly as possible
  - Many of these were Kafka-based, and could be turned on with a simple config change
  - Some required vendor-specific apps or TAs
- By day 100, ingesting over 3TB/day, approaching 200 data sources

Training

Empowering people with knowledge created enthusiastic users

- Offered *Using Splunk* training to anyone interested – had 64 take us up in the first few months, with a backlog still being worked on
  - WebEx format, 12 people together in a room for each class
  - used Splunk Education - Using Splunk/Splunk Fundamentals & S&R/now ...
- At first, we selected people for classes, but soon demand appeared on its own
  - Policy is that anyone requesting training gets it
  - Based on interested, also offered *Searching & Reporting*

priceline.com®

splunk> .conf2017

Building Community

- Created #splunk Slack channel
  - Core team answered quickly, but soon, enthusiastic users were there, too
  - Used the channel to invite people to training
- Splunk office hours held regularly
  - Held in a conference room, with open Zoom meeting for users in other location to easily join
  - Ask questions, but also show off your stuff
  - Core team, and sometimes Splunk PS or Sales Engineer

Removing Barriers

Make it available, then get out of the way

- Everyone has access with their existing credentials – https://splunk/
  - No index restrictions (exception: SecOps)
- Each team has an app space
  - Easy to request access to a team role
- Governance on data ingestions, but applied with a light touch
  - If your familiar data already exists, experimentation follows naturally
  - No enforced data models

Bonus: other things that work well

- Everything (almost) is public - including generous write privileges
- Keeping Splunk in line with our dev culture, i.e., using Bitbucket for configurations

Day 100

110
Weekly active users

175
Data sources

3.4
TB indexed/day

64
Users trained

?

priceline.com®

splunk> .conf2017

Day 100 Recap

- [chart - adoption through early Jan?]
- 10-12 user groups with dedicated apps
- 80+ apps installed
- [how does this slide relate with the prior]
- What did we learn?

Success Story

Even users who haven't (yet) been trained are making useful dashboards and alerts.

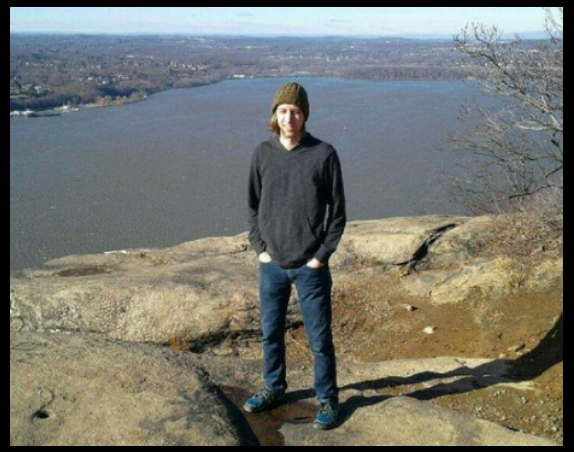

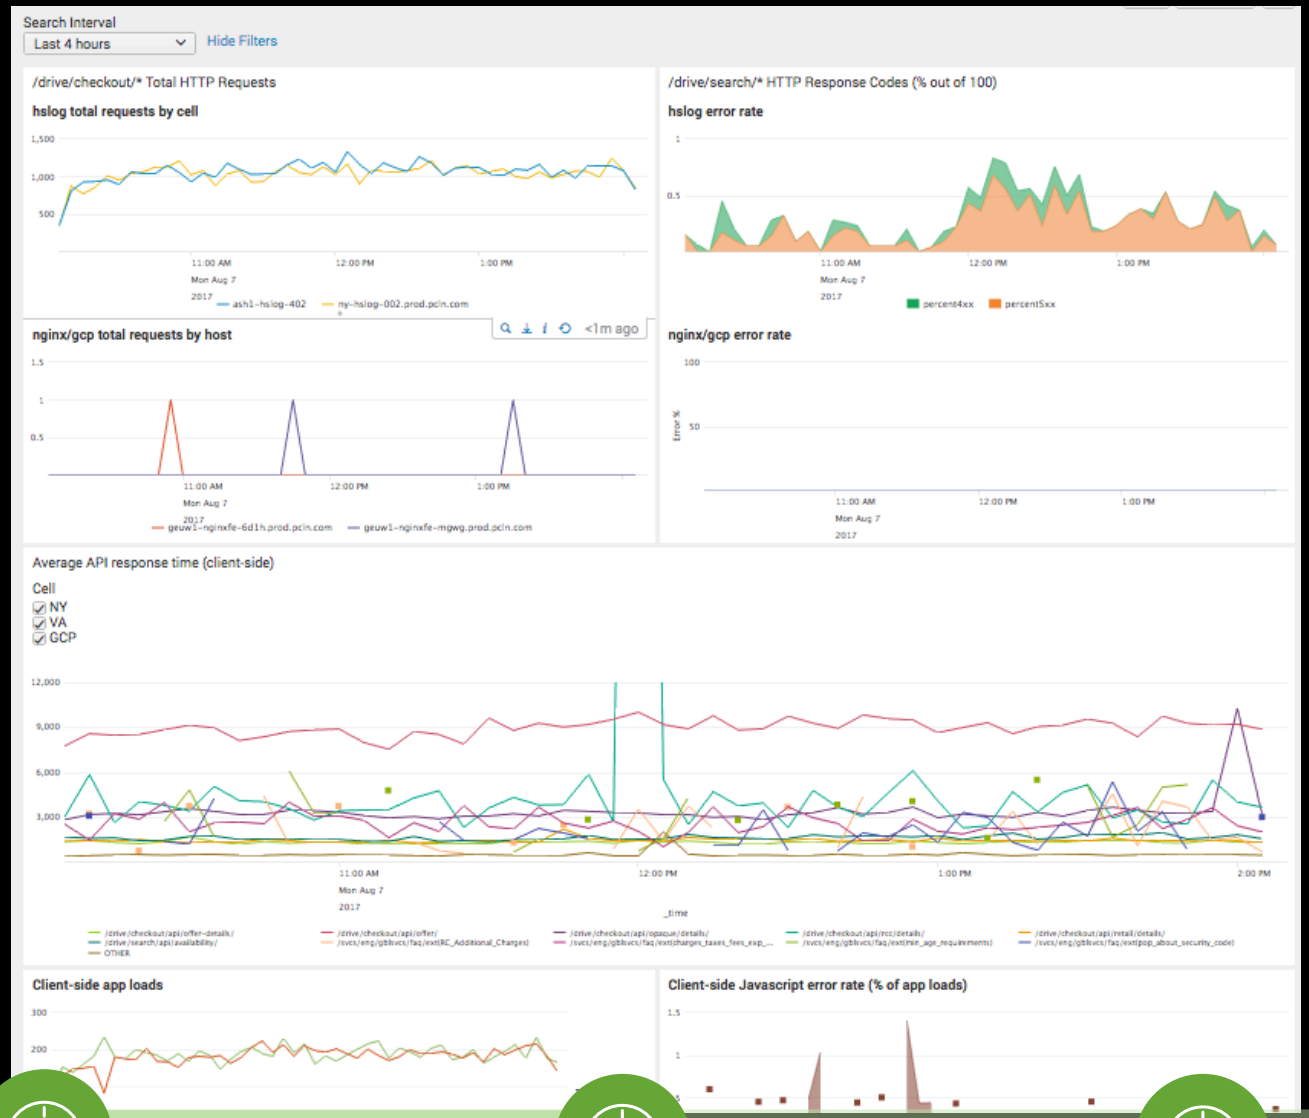

splunk> .conf2017

Day 200

160	200	4.0	110	3
↗50	↗25	↗.6	↗46	↗3
Weekly active users	Data sources	TB indexed/day	Users trained	Certifications

priceline.com®

splunk> .conf2017

Dashboards

priceline.com®

splunk> .conf2017

Day 200 Recap

- SRE and SRO requiring Splunk dashboards & alerts for application turnup or changes
- Maturity in our use starts to inspire ERs
- More stats ???
- ES implemented
- 16 User apps
- First upgrade experience

priceline.com®

splunk> .conf2017

“ Splunk is definitely **indispensable** now, I can not even consider going back! It has saved our team so very many hours and **exposed problems** that were obscured in the **deluge** of logs and data.

Dev manager – Hotels team

Key Takeaways

For viral adoption

1. Make data available early and often
2. Train!
3. Build Community
4. Remove as many barriers as possible

Current Architecture

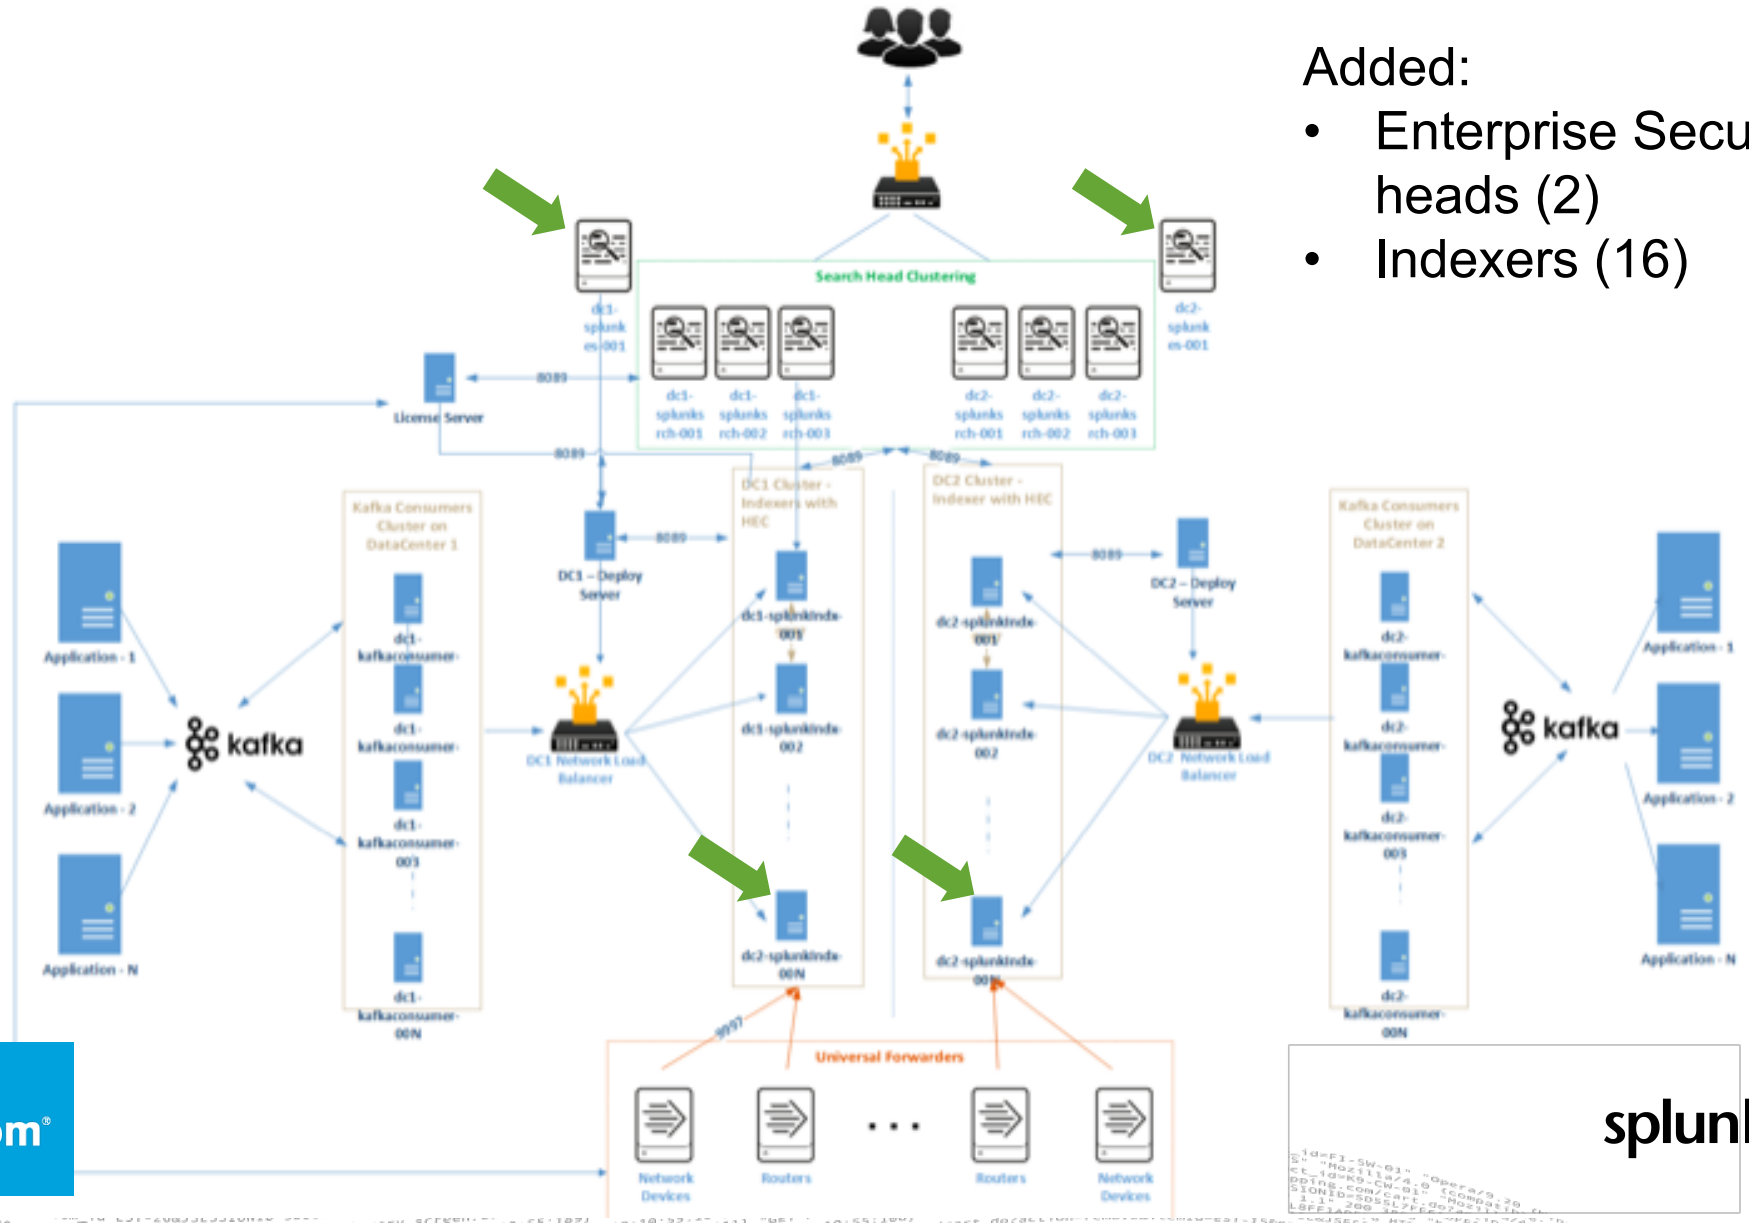

Added:

- Enterprise Security search heads (2)
- Indexers (16)

priceline.com®

splunk> .conf2017

Beyond Day 200

Adoption continues…

- Splunk has become part of our DNA
- Still have teams early in adoption, user community continues to grow
- Moving beyond viral -- adoption moving into team goals
  - Some teams have Q3’17 goals to further their Splunk adoption
- Timelines established for deprecating legacy systems

priceline.com®

splunk> .conf2017

Long Term

Platform continues to evolve

- Smarter ingestion on some statistical data sources
- More use of summary indexes
  - more for long-term retention than performance
- Machine Learning – learn where to use it
- Moving from purely tech ops to biz ops

"Since we launched Splunk,
I haven't had a single reason to use
sed, awk, or grep.

Director, SRE

splunk> .conf2017

?What we'd do differently

priceline.com®

splunk> .conf2017

Q&A

Tom Gerhard | Fellow, priceline.com

Vidhya Ramachandran | Principal Engineer, priceline.com

priceline.com

splunk> .conf2017

Thank You

Don't forget to rate this session in the .conf2017 mobile app