Build-Deploy-Run large scale logging infrastructure for SAP Cloud Platform and Cloud Applications

Istvan Zoltan Ballok Hariharan Gandhi





GridKa School 2018 August 30, 2018, SCC at KIT Karlsruhe, Germany



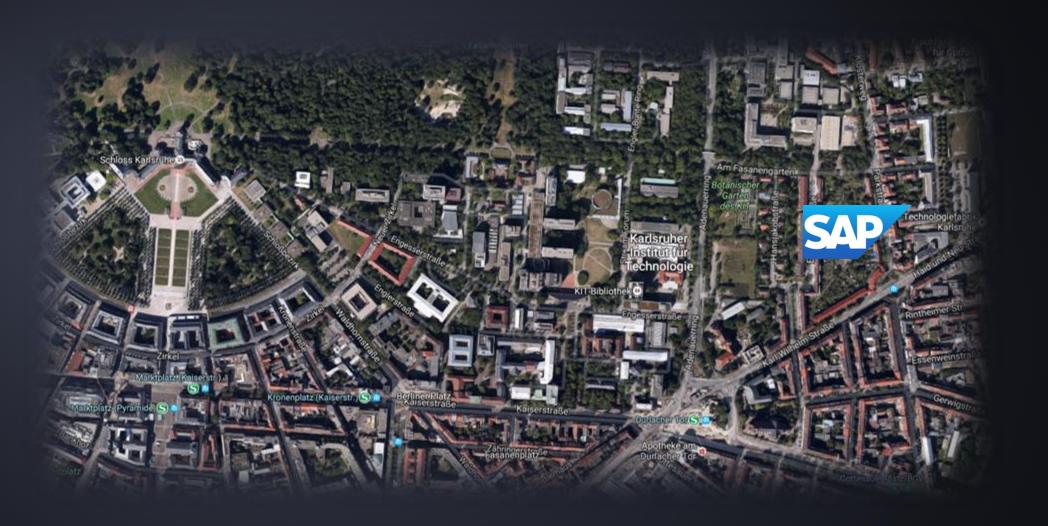
Istvan Zoltan Ballok Developer



Hariharan Gandhi
Developer

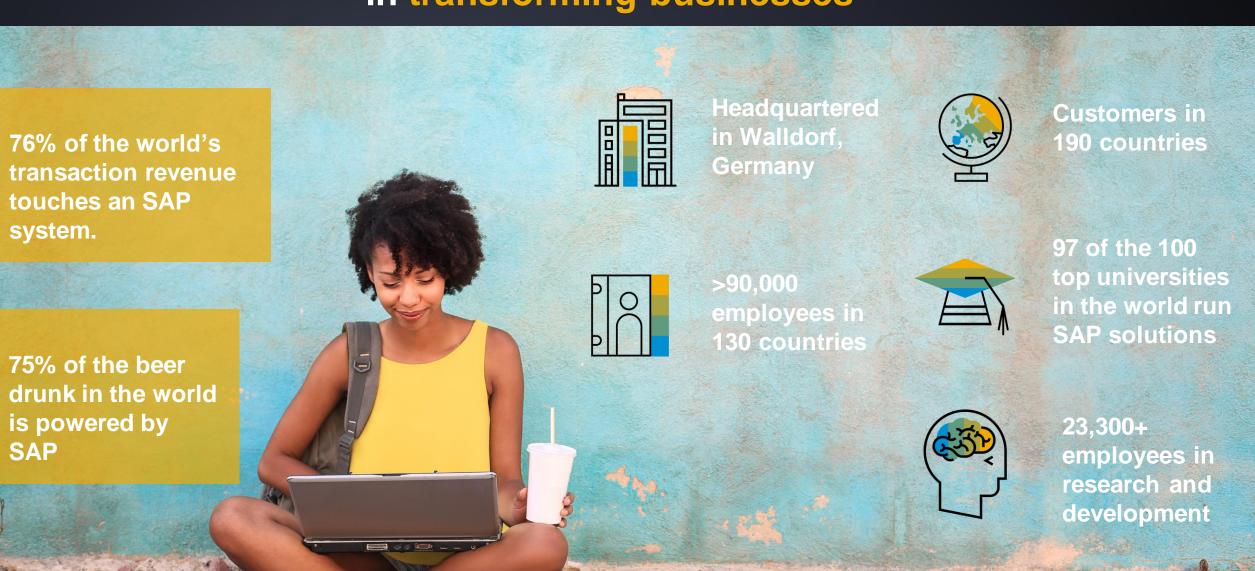


Jörg Bischof Development Manager



SAP is...

a long time trusted leader in transforming businesses





Our Customers

Produce more than **79% of chocolates** produced in the world



Brew more than **77% of beer** produced in the world



Our Customers

Produce more than **86% of sport shoes** produced in the world



Outline

What SAP does

- Business applications
- Technology Platform

What we do

- From Open Source to an Enterprise-ready Cloud Platform
- Cloud Services at Karlsruhe

Technical focus

- Cloud Foundry
- Application Logging pipeline

Infrastructure

Applications

Runtimes

Security

Databases

Server

Virtualization

Storage

Networking

Platform (as a Service)

nok Managed by

Applications

Runtimes

Databases

Server

Virtualization

Storage

Networking

Managed by Vendor

Software (as a Service)

Applications

Databases

Managed by

Vendor

Server

Virtualization

Storage

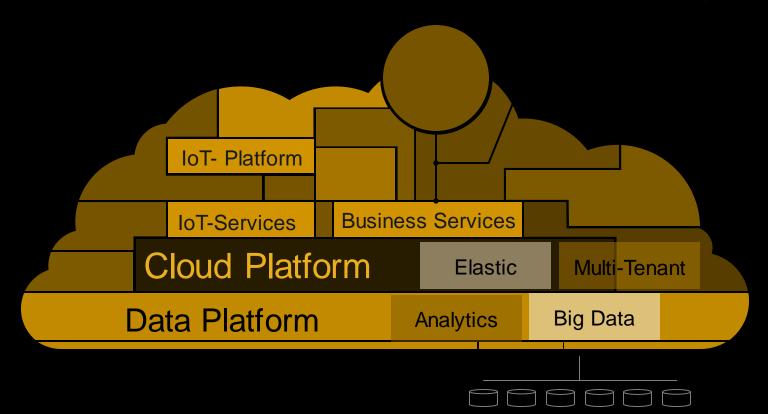
Networking

Managed by Vendor

Source: http://cloudacademy.com/blog/cloud-foundry-benefits/

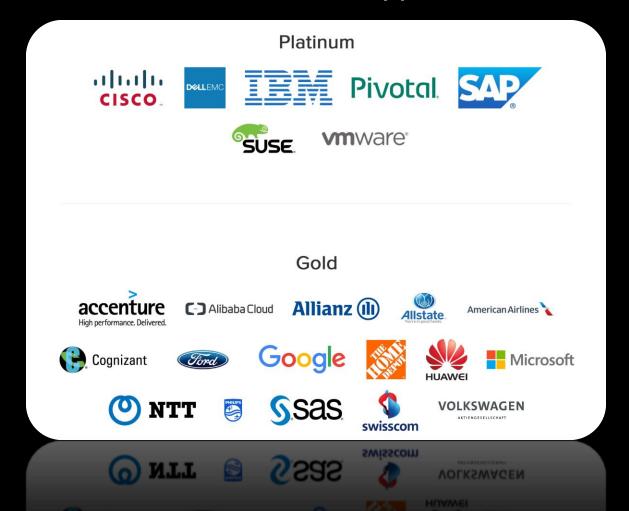
"Cloud is about how you do computing, not where you do computing."

- Paul Maritz



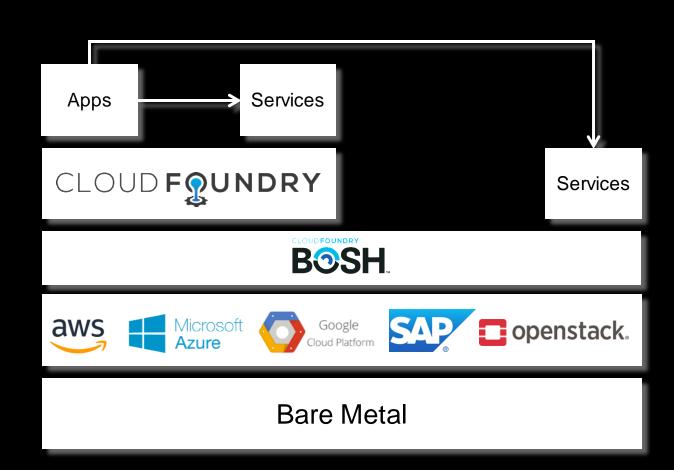
Cloud Foundry

is an Open Source PaaS driven **by 60+** strong partners providing choice of clouds, frameworks, and application services



Cloud Foundry: PaaS

- "Here is my source code, run it on the cloud for me, I do not care how"
 `cf push` experience
- Security updates for operating system, runtimes and "app servers" (e.g. Tomcat) are handled by CF



Cloud Foundry – from the developer's perspective

- Create the first version and add a manifest
- Deploy new versions rapidly

```
cf push myApp
```

Assemble services to create integrated applications

```
cf create service dbService medium myDB cf bind-service myApp myDB
```

Scale applications

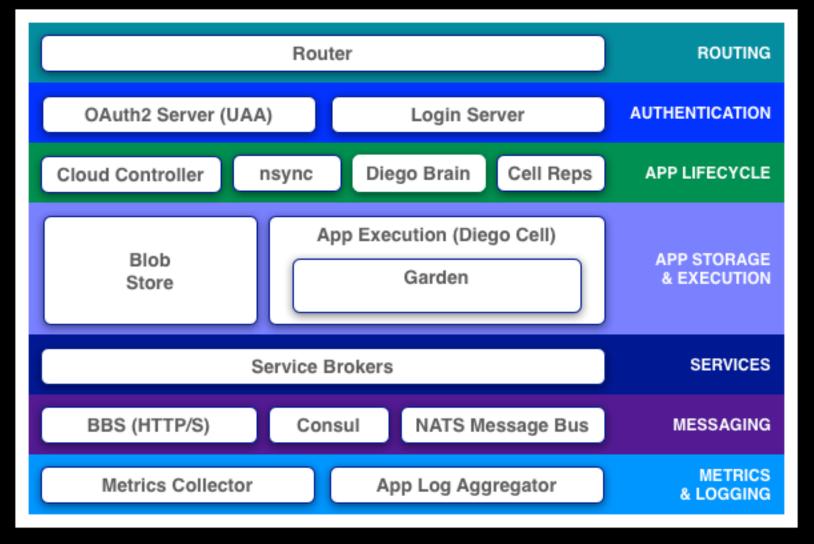
```
cf scale myApp -i 100
```

Don't worry about underlying operating systems. Easier backup and recovery strategies using platform functionality.

http://bit.ly/gridka-chat

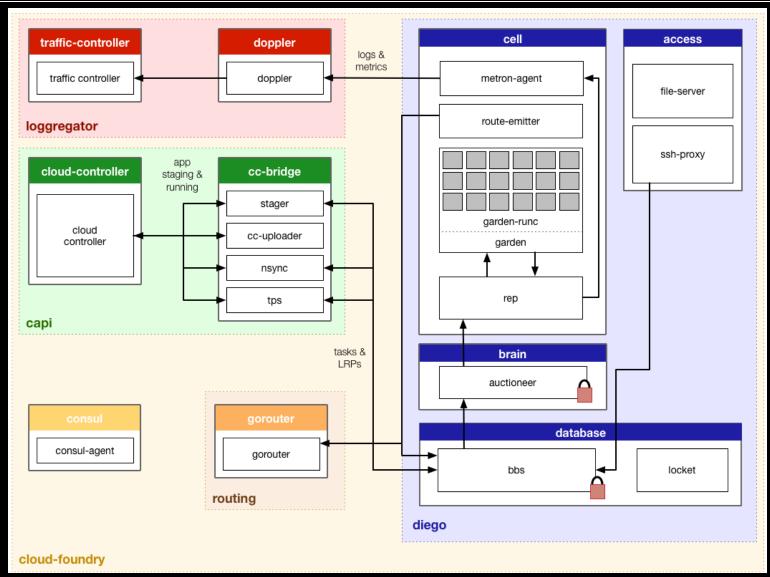


Cloud Foundry Architecture

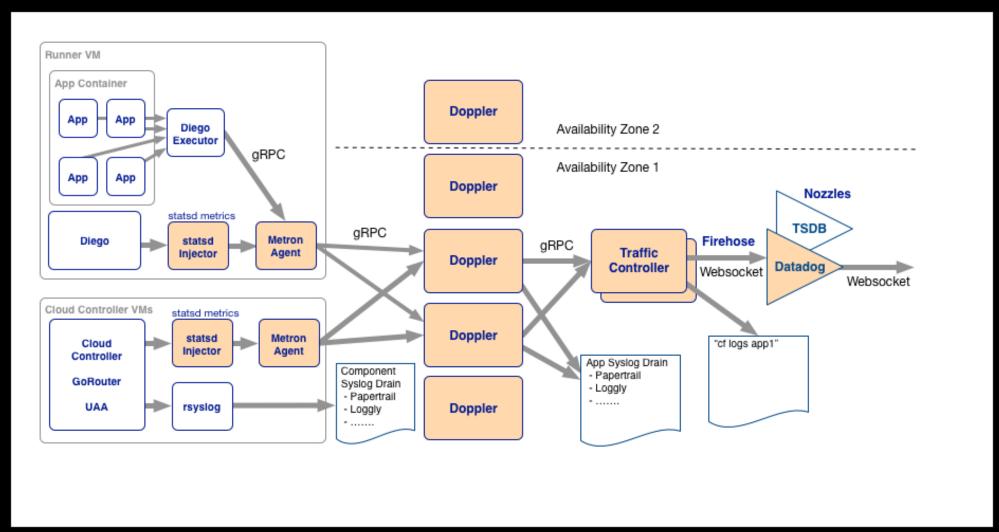


https://docs.cloudfoundry.org/concepts/architecture/

Cloud Foundry, Diego Architecture



CF Loggregator Architecture



https://docs.cloudfoundry.org/loggregator/architecture.html



Application Logging

making lives of developers easier

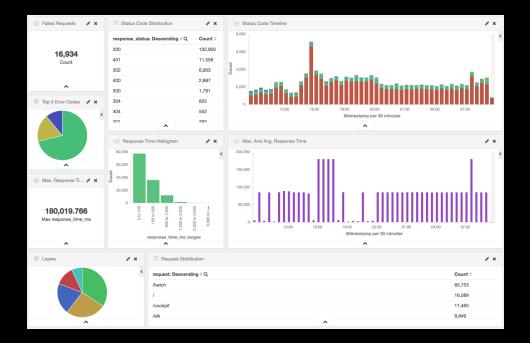
6, "component_type": "application", "component_id": "9924b8dc-1cbb-4a60-88e1-f020e70d85b8", "space_name": "test", "component_id": "9924b8dc-1cbb-4a60-88e1-f020e70d85b8", "space_name": "test", "component_id": "9924b8dc-1cbb-4a60-88e1-f020e70d85b8", "space_name": "test", "component_id": "9924b8dc-1cbb-4a60-88e1-f020e70d85b8", "space_name": "test", "component_id": "0.045f-e3d1a54208ec", "organization_name": "-", "space_id": "67bbd273-cc25-48e8-9d93-c279ef7dab4e", "request_id": "bab4f76f51-4a00-645f-e3d1a54208ec", "container_id": "10.1.7.66", "type": "request", "request": "foo/bar", "referer": "-", "responsent_id": "2016-11-18T08:36:17.325Z", "response_status": 404, "method": "GET", "response_size_b":-1, "request_size_b":-1, "rete_port": "11376", "layer": "[SERVLET]", "remote_host": "172.18.74.8", "x_forwarded_for": "-", "remote_user": "-", "protocol' TTP/1.1", "remote_ip": "172.18.74.8", "response_content_type": "-", "request_received_at": "2016-11-18T08:36:17.324Z", "response_time_ms": 0.877432, "direction": "IN"}

2016-11-18T09:36:17.32+0100 [RTR/0] OUT d048888-logging-sample-app.cfapps.staging.hanavlab.ondemand.com - [18/ 2016:08:36:17.226 +0000] "GET /foo/bar HTTP/1.1" 404 0 0 "-" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10.12.0) Apple Kit/537.36 (KHTML, like Gecko) Chrome/54.0.2840.98 Safari/537.36" 192.168.0.107:37654 x_forwarded_proto: "172.18.74.8, 2.168.0.107" x_forwarded_proto: "https" vcap-request_id:bab4f702-bf51-4a00-645f-e3d1a54208ec response_time:0.1013211 app_id:9924b8dc-1cbb-4a60-88e1-f020e70d85b8 app_idex:0

2016-11-18T09:36:21.38+0100 [APP/PROC/WEB/0]OUT { "written_at":"2016-11-18T08:36:21.387Z", "written_ts":195451869114 8, "component_type":"application", "component_id":"9924b8dc-1cbb-4a60-88e1-f020e70d85b8", "space_name":"test", "component_mame":"d048888-logging-sample-app", "component_instance":00", "organization_id":"-", "correlation_id":"2d45c241-9357-4-595a-92d9b39b9cfa", "organization_name":"-", "space_id":"67b273-cc25-48e8-9a93-c279ef7dab4e", "request_id":"2d45c24-357-48af-595a-92d9b39b9cfa", "container_id":"10.1.7.66", "type":"log", "logger":"com.sap.hcp.cf.sample.PickGreeting", ead":"http-nio-8080-exec-7", "level":"INFO", "categories": [], "msg": "Picked greeting: \times off }

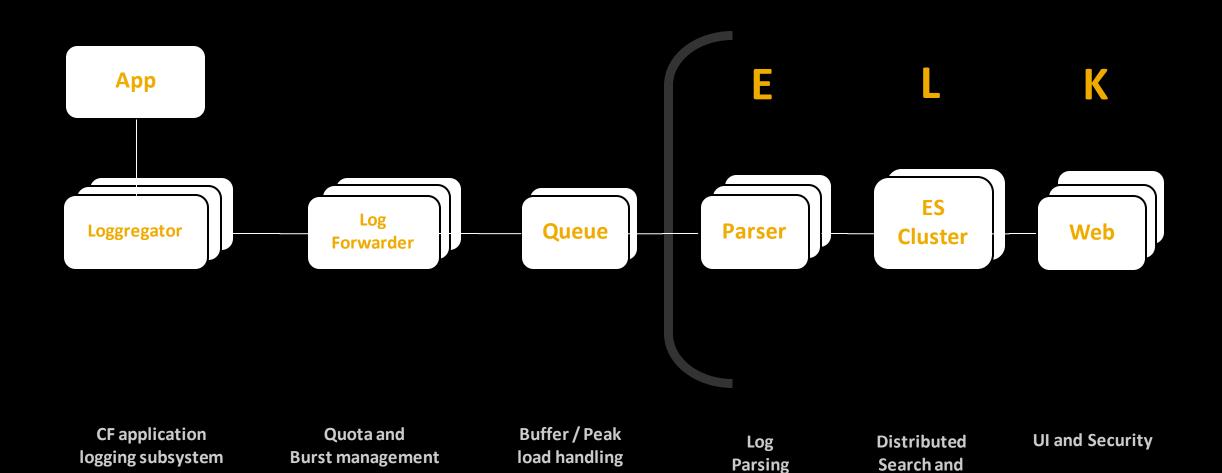
2016-11-18T09:36:21.38+0100 [RTR/0] OUT d048888-logging-sample-app.cfapps.staging.hanavlab.ondemand.com - [18/ 2016:08:36:21.305 +0000] "GET / HTTP/1.1" 200 0 3123 "-" "Mozilla/5.0 (Mactintosh; Intel Mac OS X 10.12_0) AppleWebb 537.36 (KHTML, like Gecko) Chrome/54.0.2840.98 Safari/537.36" 192.168.0.106:53120 x_forwarded_fror:"172.18.74.8, 192 8.0.106" x_forwarded_proto: "https" vcap_request_id:2d45c241-9357-48af-595a-92d9b39b9cfa response_time:0.083701217 c id:9924b8dc-1cbb-4a60-88e1-f020e70d85b8 app_index:0

2016-11-18T09:36:21.38+0100 [APP/PROC/WEB/0]OUT { "written_at":"2016-11-18T08:36:21.389Z", "written_ts":19545186929: 1, "component_itype": "application", "component_it":"9924b8dc-1cbb-4c66-88e1-f020e70d85b8", "space_name": "test", "component_ine":"d048888-logging-sample-app", "component_instance":"0", "organization_id":"-", "correlation_id":"2d45c241-9357-4-595a-92d9b39b9cfa", "organization_name":"-", "space_id":"67bbd273-cc25-48e8-9a93-c279ef7dab4e", "request_id":"2d45c24357-48af-595a-92d9b39b9cfa", "container_id":"10.1.7.66", "type":"request":"/", "referrer":"-", "response_sent:"2016-11-18T08:36:21.388Z", "response_status":200, "method":"GET", "response_size_b":3119, "request_size_b":-1, "remote rt":"60030", "layer":"[SERVLET]", "remote_lost":"172.18.74.8", "x_forwarded_for":"-", "remote_user":"-", "protocol":"HTI



Application Logging

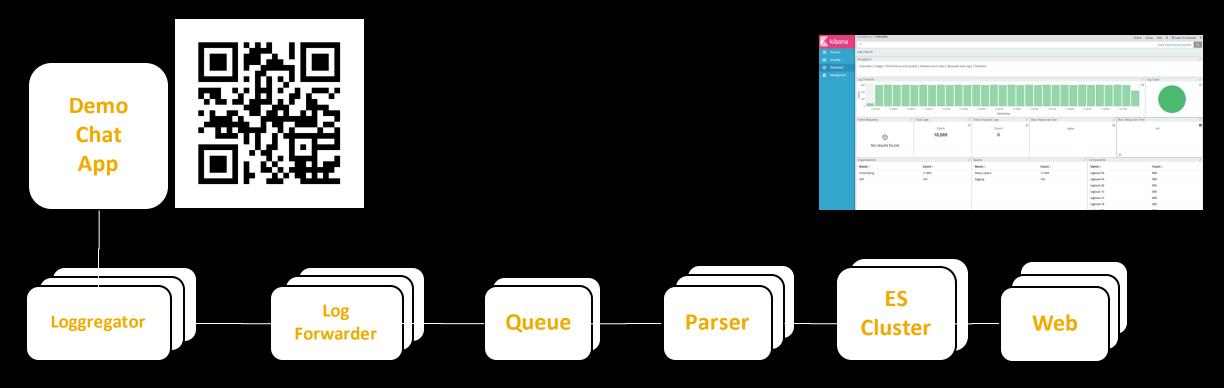
Pipeline



Analytics

Let's check if the demo works?

http://bit.ly/gridka-chat



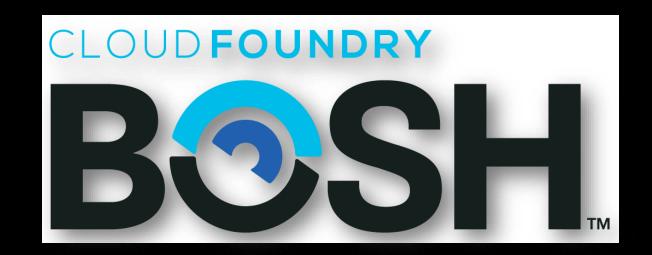
Try it! https://cloudplatform.sap.com

How do we deploy the logging stack?

- We can not use Cloud Foundry Application Runtime to deploy stateful services
- Bosh
 - The same tool that is used to deploy and manage the lifecycle of Cloud Foundry itself

BOSH is a project that unifies release engineering, deployment, and lifecycle management of small and large-scale cloud software.

- Immutable infrastructure (pet vs. cattle)
- Infrastructure agnostic (CPI)
- Bosh release, jobs, packages
- Deployment manifest
- Rolling upgrade
- Health checks



Why is parsing needed?

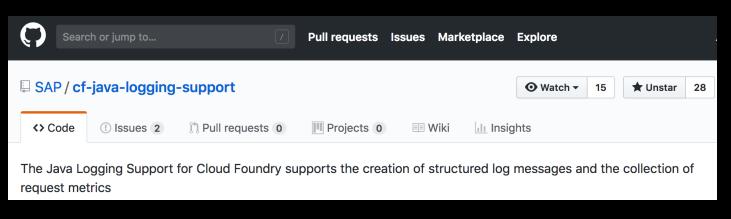
```
cf-logging-monitor.cf.sap.hana.ondemand.com - [2018-08-29T12:58:52.421+0000] "GET /health HTTP/1.1" 200 0 129 "-" "Apache-HttpClient/4.5.3 (Java/1.8.0_172)" "-" "10.0.137.45:62962" x_forwarded_for:"-" x_forwarded_proto:"https" vcap_request_id:"72fdd51d-eb6d-48cb-4001-69681d02573e" response_time:0.00129951 app_id:"26f8ad18-eddf-41df-9f38-6f7d44036867" app_index:"0" x_b3_traceid:"9ca7e3e7b875fca7" x b3 parentspanid:"-"
```

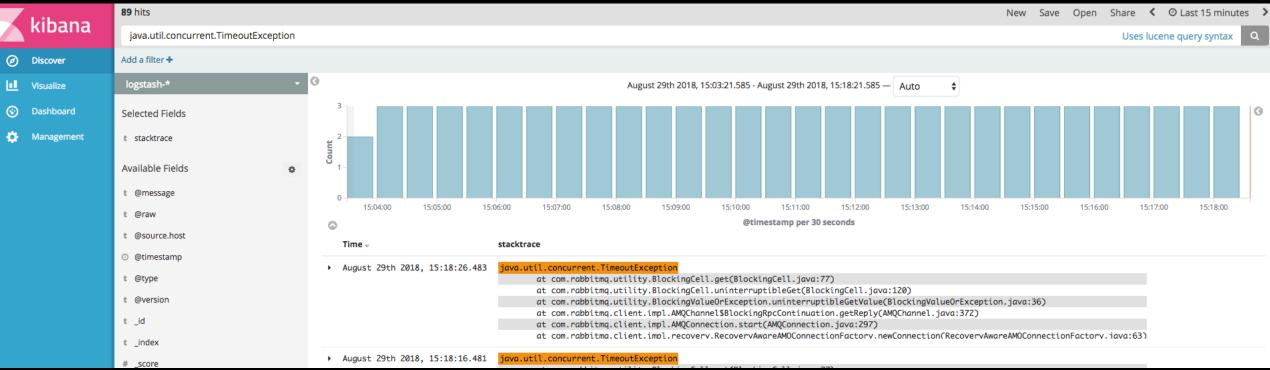


```
{
    "msg": "cf-logging-monitor.cf.sap.hana.ondemand.com - [2018-08-29T12:58:52.421+0000] \"GET /health HTTP/1.1\" 200 0 129 \"-\"
    (Java/1.8.0_172)\" \"-\" \"10.0.137.45:62962\" x_forwarded_for:\"-\" x_forwarded_proto:\"https\" vcap_request_id:\"72fdd51d-eb
    response_time:0.00129951 app_id:\"26f8ad18-eddf-41df-9f38-6f7d44036867\" app_index:\"0\" x_b3_traceid:\"9ca7e3e7b875fca7\"
    x_b3_spanid:\"9ca7e3e7b875fca7\" x_b3_parentspanid:\"-\"",
        "referer": "-",
        "request_size_b": 0,
        "channel": "0UT",
        "type": "request",
        "remote_user": "-",
        "protocol": "HTTP/1.1",
        "remote_ip": "\"-\"",
        "request_received_at": "2018-08-29T12:58:52.4212",
        "component_id": "26f8ad18-eddf-41df-9f38-6f7d44036867",
        "method": "GET",
        "level": "INFO",
```

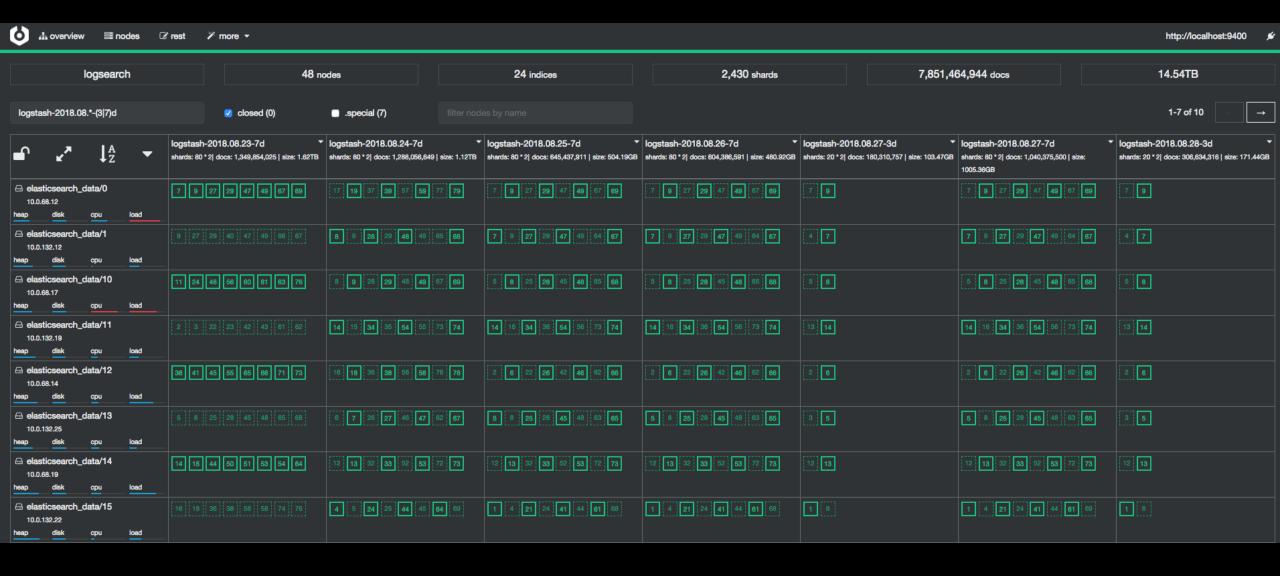
Emit structured logs

https://cloudplatform.sap.com/enterprise-paas/cloudfoundry/cf-contributions.html



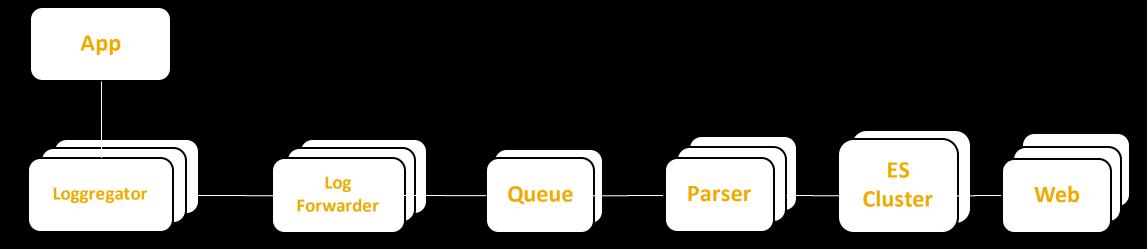


Inside Elasticsearch



Pain Points, Challenges

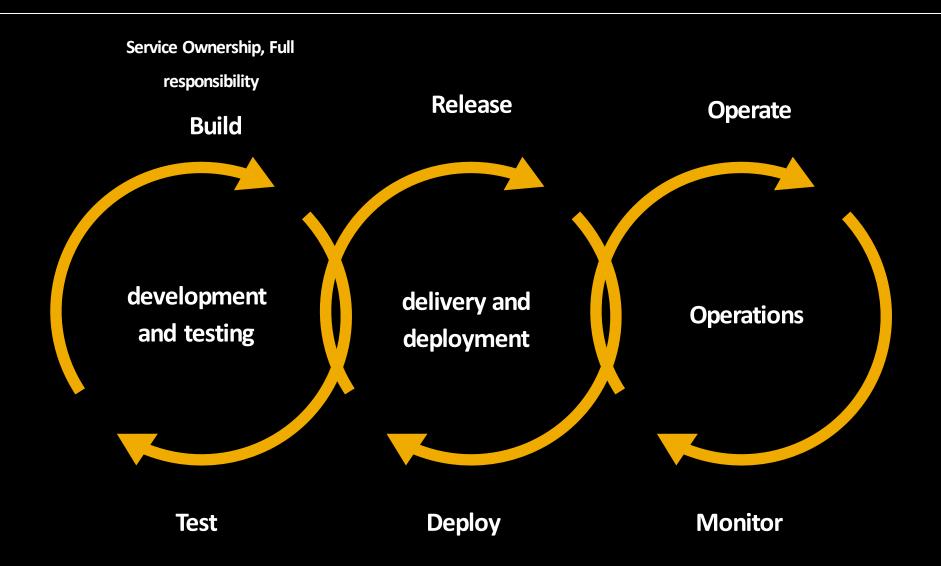
- Regular Expr performance during parsing
- Rolling update of Elasticsearch and all components staying ZDM
- Resizing persistent disks used by Elasticsearch attached volumes cannot be expanded by Bosh
- Tuning for indexing performance in Elasticsearch (async flush to disk)
- Quota to protect the stack







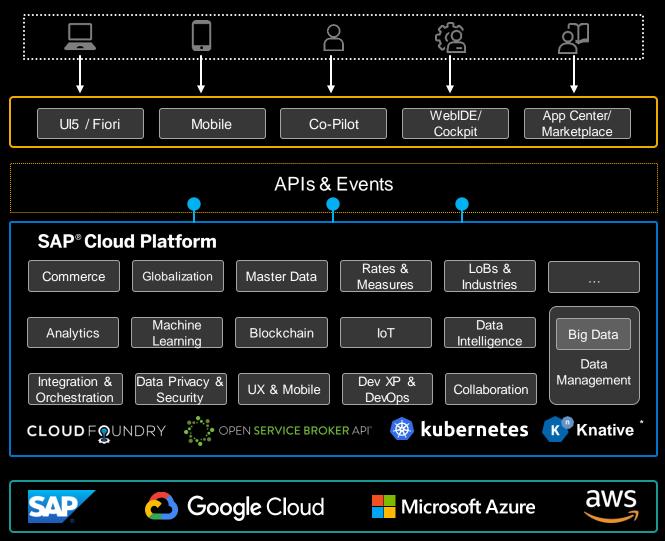
DevOps



SAP Cloud Platform: Centerpiece of the Intelligent Enterprise



SAP Cloud Platform: An ongoing journey, especially around Open Source technologies







SAP® Cloud Platform

DEVELOPERS WANTED

Cloud Cockpit

Logging

Big Data

Persistence



Kubernetes, Knative
Cloud Foundry
Security
Performance