

splunk>

.conf2017

© 2017 SPLUNK INC.

Indexer Clustering Fixups

Cluster recovery process

Da Xu | Engineering | Splunk

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, Splunk>, Listen to Your Data, The Engine for Machine Data, Splunk Cloud, Splunk Light and SPL are trademarks and registered trademarks of Splunk Inc. in the United States and other countries. All other brand names, product names, or trademarks belong to their respective owners. © 2017 Splunk Inc. All rights reserved.

Clustering Fixup Cycle

Clustering Overview - Indexer Failure



130.60.4 - - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GIFTS&SESSIONID=SD1SLAFF10ADFF10 HTTP 1.1" 404 720 "http://buttercup-shopping.com/cart.do?action=view&itemId=EST-6&product_id=FI-5W-03"
128.241.220.82 - - [07/Jan 18:10:57:123] "GET /product.screen?product_id=FL-D5H-01&SESSIONID=SD35L7FF6ADFF9 HTTP 1.1" 404 3322 "http://buttercup-shopping.com/category.screen?category_id=GIFTS&SESSIONID=SD1SLAFF10ADFF10 HTTP 1.1"
10.317.27.160.0.0 - - [07/Jan 18:10:56:156] "GET /product.screen?product_id=FL-D5H-01&SESSIONID=SD35L7FF6ADFF9 HTTP 1.1" 200 1318 "http://buttercup-shopping.com/cart.do?action=purchase&itemId=EST-26&SESSIONID=SD5L9FF1ADFF3 HTTP 1.1"
10.317.27.160.0.0 - - [07/Jan 18:10:56:156] "GET /product.screen?product_id=FL-D5H-01&SESSIONID=SD35L7FF6ADFF9 HTTP 1.1" 200 1318 "http://buttercup-shopping.com/cart.do?action=purchase&itemId=EST-26&SESSIONID=SD5L9FF1ADFF3 HTTP 1.1"
10.317.27.160.0.0 - - [07/Jan 18:10:56:156] "GET /product.screen?product_id=FL-D5H-01&SESSIONID=SD35L7FF6ADFF9 HTTP 1.1" 200 1318 "http://buttercup-shopping.com/cart.do?action=purchase&itemId=EST-26&SESSIONID=SD5L9FF1ADFF3 HTTP 1.1"
10.317.27.160.0.0 - - [07/Jan 18:10:56:156] "GET /product.screen?product_id=FL-D5H-01&SESSIONID=SD35L7FF6ADFF9 HTTP 1.1" 200 1318 "http://buttercup-shopping.com/cart.do?action=purchase&itemId=EST-26&SESSIONID=SD5L9FF1ADFF3 HTTP 1.1"
10.317.27.160.0.0 - - [07/Jan 18:10:56:156] "GET /product.screen?product_id=FL-D5H-01&SESSIONID=SD35L7FF6ADFF9 HTTP 1.1" 200 1318 "http://buttercup-shopping.com/cart.do?action=purchase&itemId=EST-26&SESSIONID=SD5L9FF1ADFF3 HTTP 1.1"
10.317.27.160.0.0 - - [07/Jan 18:10:56:156] "GET /product.screen?product_id=FL-D5H-01&SESSIONID=SD35L7FF6ADFF9 HTTP 1.1" 200 1318 "http://buttercup-shopping.com/cart.do?action=purchase&itemId=EST-26&SESSIONID=SD5L9FF1ADFF3 HTTP 1.1"
10.317.27.160.0.0 - - [07/Jan 18:10:56:156] "GET /product.screen?product_id=FL-D5H-01&SESSIONID=SD35L7FF6ADFF9 HTTP 1.1" 200 1318 "http://buttercup-shopping.com/cart.do?action=purchase&itemId=EST-26&SESSIONID=SD5L9FF1ADFF3 HTTP 1.1"
10.317.27.160.0.0 - - [07/Jan 18:10:56:156] "GET /product.screen?product_id=FL-D5H-01&SESSIONID=SD35L7FF6ADFF9 HTTP 1.1" 200 1318 "http://buttercup-shopping.com/cart.do?action=purchase&itemId=EST-26&SESSIONID=SD5L9FF1ADFF3 HTTP 1.1"

Fixup Lifecycle – service



#1 bucket
added to
to_fix list

#2 service()
schedules
Jobs

- ▶ Periodically, the CM will run through its service() loop
 - Defaults to every second, see the setting of “service_interval”
- ▶ During the service() loop, the CM will iterate through ALL the bucket on all the to_fix lists, and schedule any jobs that it needs to for repairs.

```
function service()
{
  for bucket in to_fix_lists:
    schedule_jobs_to_fix(bucket) // maybe schedule a job?
}
```

- (not actually taken from splunk/clustering.cpp)
- ▶ If the bucket doesn't need any jobs OR we've scheduled all jobs to fix a bucket, its removed from the to_fix list

Fixup Lifecycle – jobs

- ▶ Once service() schedules jobs, the CM has a threadpool running and picking up all the scheduled jobs.
- ▶ These jobs are generally REST POSTs to the cluster Indexers, telling them to run some command

Job name	Function
CMRepJob	Replicate a bucket (includes recipient info)
CMChangeBucketJob_build	Make bucket searchable (RF copy -> SF copy)
CMChangeMasks	Change bucket primaries (batch job, can change maybe buckets)
CMRollHotBucket	Roll a bucket from hot->warm
CMTruncJob	Truncate a bucket's size (remove bucket will use this with size=0)

Fixup Lifecycle – jobs



#1 bucket added to to_fix list

#2 service() schedules Jobs

#3 jobs run separately

- ▶ All the jobs are ran asynchronously on separate threads
 - "executor_workers" specifies how many separate threads run these jobs (defaults to 10)
- ▶ These jobs generally POST to a specific Peer endpoint
 - ex :
indexer01:8089/services/cluster/slave/buckets/main~1~.../replicate
- ▶ If the job runs successfully
 - All is well! The bucket replicated, or got its primary changed, or etc.
- ▶ If the job encounters an error
 - The bucket is placed BACK into the to_fix list, so that we can fix it up again on the next service loop!

Clustering Fixup Cycle - examples

to_fix lists metrics

New Search

Save As Close

```
index=_internal source=*metrics.log name=cmmaster_service | timechart max(to_fix_gen) max(to_fix_rep_factor) max(to_fix_search_factor) span=30s
```


Date time range

18 events (9/5/17 1:41:00.000 AM to 9/5/17 1:45:30.000 AM) No Event Sampling

Job Stop Refresh Download Smart Mode

Events Patterns Statistics (9) Visualization

Column Chart Format



Searchable

RF+SF Met



jobs metrics

New Search

Save As Close

index=_internal source=*metrics.log group=jobs name=cmmaster | timechart max(*Job*) span=10s


Date time range

8 events (9/5/17 1:41:00.000 AM to 9/5/17 1:46:00.000 AM) No Event Sampling

Job [] [] [] [] [] Smart Mode

Events Patterns Statistics (30) Visualization

Column Chart Format



max(CMChangeBucketJob_build)
max(CMChangeMasksJob)
max(CMRepJob)

searchable

RF+SF Met

130.60.4 - - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GIFTS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 404 3322 "http://buttercup-shopping.com/cart.do?action=view&itemId=EST-6&product_id=FI-SW-03" "Mozilla/5.0 (Windows NT 5.1; SV1; .NET CLR 1.1.4322)" 468 125.17 14...
128.241.220.82 - - [07/Jan 18:10:57:123] "GET /product.screen?product_id=FL-DSH-01&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 200 1318 "http://buttercup-shopping.com/cart.do?action=changequantity&itemId=EST-18&product_id=AV-CB-01&JSESSIONID=5D55L9FF1ADFF3" "Mozilla/5.0 (Windows NT 5.1; SV1; .NET CLR 1.1.4322)" 468 125.17 14...
137.27.160.0 - - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GIFTS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 404 3322 "http://buttercup-shopping.com/cart.do?action=view&itemId=EST-6&product_id=FI-SW-03" "Mozilla/5.0 (Windows NT 5.1; SV1; .NET CLR 1.1.4322)" 468 125.17 14...
137.27.160.0 - - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GIFTS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 404 3322 "http://buttercup-shopping.com/cart.do?action=view&itemId=EST-6&product_id=FI-SW-03" "Mozilla/5.0 (Windows NT 5.1; SV1; .NET CLR 1.1.4322)" 468 125.17 14...
137.27.160.0 - - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GIFTS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 404 3322 "http://buttercup-shopping.com/cart.do?action=view&itemId=EST-6&product_id=FI-SW-03" "Mozilla/5.0 (Windows NT 5.1; SV1; .NET CLR 1.1.4322)" 468 125.17 14...
137.27.160.0 - - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GIFTS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 404 3322 "http://buttercup-shopping.com/cart.do?action=view&itemId=EST-6&product_id=FI-SW-03" "Mozilla/5.0 (Windows NT 5.1; SV1; .NET CLR 1.1.4322)" 468 125.17 14...
137.27.160.0 - - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GIFTS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 404 3322 "http://buttercup-shopping.com/cart.do?action=view&itemId=EST-6&product_id=FI-SW-03" "Mozilla/5.0 (Windows NT 5.1; SV1; .NET CLR 1.1.4322)" 468 125.17 14...
137.27.160.0 - - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GIFTS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 404 3322 "http://buttercup-shopping.com/cart.do?action=view&itemId=EST-6&product_id=FI-SW-03" "Mozilla/5.0 (Windows NT 5.1; SV1; .NET CLR 1.1.4322)" 468 125.17 14...
137.27.160.0 - - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GIFTS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 404 3322 "http://buttercup-shopping.com/cart.do?action=view&itemId=EST-6&product_id=FI-SW-03" "Mozilla/5.0 (Windows NT 5.1; SV1; .NET CLR 1.1.4322)" 468 125.17 14...

jobs by time

New Search

Save As Close

index=_internal source=*master*splunkd.log CMRepJob | timechart span=15s count by job

Date time range


✓ 3,264 events (9/5/17 1:41:00.000 AM to 9/5/17 1:46:00.000 AM) No Event Sampling

Job | | | | | Smart Mode

Events Patterns Statistics (20) Visualization

Column Chart Format

CMChangeBucketJob_build CMChangeMasksJob CMRepJob



searchable

RF+SF Met

splunk> .conf2017

Thanks - Q&A

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

splunk® **.conf2017**

