

# An Unexpected Open Source Win

---

Amy Tobey  
Database Reliability Engineering Manager

# Tenable Cassandra

---

- I joined Tenable in September 2017
- We have a multi-cloud mandate
- DynamoDB → Cassandra
- 34 clusters

# I ♥ (some) Pickles

---

- Salty
- Spicy
- Sweet
- ~~Vinegar~~
- ~~import pickle~~

@MissAmyTobey

# I ♥ (many) Pickles

---


- Salty
- Spicy
- Sweet
- ~~Vinegar~~
- ~~import pickle~~
- Cassandra Consultants

# THE LAST PICKLE

[@MISSAMYTOBEY](#)



# THE LAST PICKLE



# More on Pickles

---

- Configuration Analysis
- Cluster analysis
- Data Model Review
- ALL FOR ONE ~~LOW LOW~~ MODERATELY GOOD PRICE!





[CASSANDRA-13080] x

Secure | https://issues.apache.org/jira/browse/CASSANDRA-13080

APACHE SOFTWARE FOUNDATION  
http://www.apache.org/

Dashboards Projects Issues

Search

Log In

Cassandra / CASSANDRA-13080

# Use new token allocation for non bootstrap case as well.

Export

### Details

Type:	Improvement	Status:	RESOLVED
Priority:	Major	Resolution:	Fixed
Component/s:	Core	Fix Version/s:	4.0
Labels:	None		

### Description

There are couple reasons I think we should use the new token allocation for non bootstrap case as well.

### People

Assignee: Dikang Gu

Reporter: Dikang Gu

Reviewer: Branimir Lambov

Votes:

# Cassandra

---

- Cassandra 4.0 isn't even in beta yet
- And even if it was, it's a large complicated OSS project...
- Furthermore it's Cassandra...
- Not getting near prod until 4.0.10 \*minimum\*





# Backport! ☐ ☐

```

@@ -82,7 +82,7 @@
82 82 // (don't need to record Token here since it's still part of tokenToEndpointMap until it's done leaving)
83 83 private final Set<InetAddress> leavingEndpoints = new HashSet<>();
84 84 // this is a cache of the calculation from {tokenToEndpointMap, bootstrapTokens, leavingEndpoints}
85 - private final ConcurrentMap<String, Multimap<Range<Token>, InetAddress>> pendingRanges = new ConcurrentHashMap<>();
85 + private final ConcurrentMap<String, PendingRangeMaps> pendingRanges = new ConcurrentHashMap<String, PendingRangeMap
86 86
87 87 // nodes which are migrating to the new tokens in the ring
88 88 private final Set<Pair<Token, InetAddress>> movingEndpoints = new HashSet<>();

@@ -673,23 +673,30 @@ public InetAddress getEndpoint(Token token)
673 673     return sortedTokens;
674 674 }
675 675
676 - private Multimap<Range<Token>, InetAddress> getPendingRangesMM(String keyspaceName)
676 + public Multimap<Range<Token>, InetAddress> getPendingRangesMM(String keyspaceName)
677 677 {
678 -     Multimap<Range<Token>, InetAddress> map = pendingRanges.get(keyspaceName);
679 -     if (map == null)

```

# Theory

---

Step 1: `git-format-patch 5f3b35981d281`

Step 2: `git am < lol.patch`

Step 3: ??????????????????????!

Step 4: PROFIT

@MissAmyTobey



# Practice

1. Create backport jira
2. checkout cassandra
3. Maven (ivy) download internet
4. Export Eclipse data
5. Import Eclipse data into IntelliJ
6. Find code
7. Step through patch and manually apply
8. Run dtests
9. Prepare patch
10. Mail to mailing list
11. Wait
12. Get yelled at
13. Fix complaints
14. Send another patch
15. Wait
16. Patch accepted



# Pickle to the Rescue! ☐

---

~~Step 1: Ask the consultants to do it!~~

~~Step 1: Tell the consultants to do it!~~

Step 1: Pay the consultants to do it!

Step 2: PROFIT



# Pickled Cassandra

---

- Started with a patched cassandra.jar
- Deployed to **new** clusters
- More waiting



# CHANGES.txt



ASF Git Repos - ca x

Secure | <https://git-wip-us.apache.org/repos/asf?p=cassandra.git;a=blob;f=CHANGES.txt;hb=8a5e88f635...>

[projects](#) / [cassandra.git](#) / blob

[summary](#) | [shortlog](#) | [log](#) | [commit](#) | [commitdiff](#) | [tree](#)  ? search:   re

[history](#) | [raw](#) | [HEAD](#)

**Use new token allocation for non bootstrap case as well.**

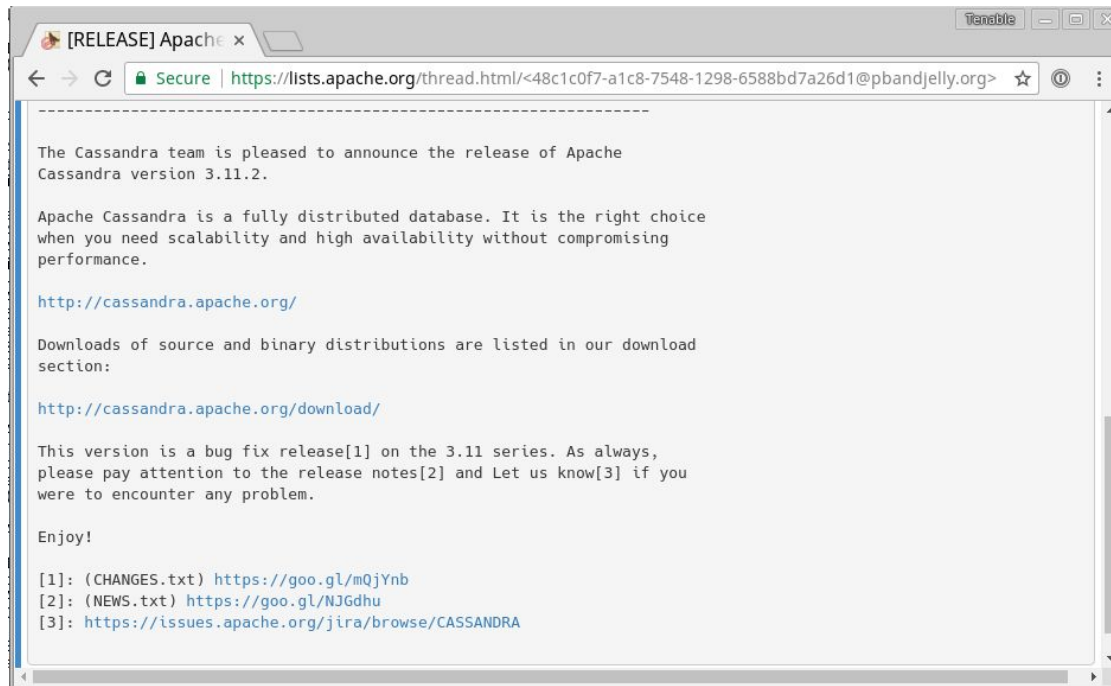
[\[cassandra.git\]](#) / [CHANGES.txt](#)

```
1 3.11.2
2 * Backport CASSANDRA-13080: Use new token allocation for non bootstrap case as well (CASSANDRA-14212)
3 * Remove dependencies on JVM internal classes from snitchServerUtils (CASSANDRA-14175)
4 * Add DEFAULT, UNSET, MBEAN and MBEANS to `ReservedKeywords` (CASSANDRA-14205)
5 * Add Unittest for schema migration fix (CASSANDRA-14140)
6 * Print correct snitch info from nodetool describecluster (CASSANDRA-13528)
7 * Close socket on error during connect on OutboundTcpConnection (CASSANDRA-9630)
8 * Enable CDC unittest (CASSANDRA-14141)
9 * Acquire read lock before accessing CompactionStrategyManager fields (CASSANDRA-14139)
10 * Split CommitLogStressTest to avoid timeout (CASSANDRA-14143)
11 * Avoid invalidating disk boundaries unnecessarily (CASSANDRA-14083)
12 * Avoid exposing compaction strategy index externally (CASSANDRA-14082)
13 * Prevent continuous schema exchange between 3.0 and 3.11 nodes (CASSANDRA-14109)
14 * Fix imbalanced disks when replacing node with same address with JBOD (CASSANDRA-14084)
15 * Reload compaction strategies when disk boundaries are invalidated (CASSANDRA-13948)
16 * Remove OpenJDK log warning (CASSANDRA-13916)
17 * Prevent compaction strategies from looping indefinitely (CASSANDRA-14079)
18 * Cache disk boundaries (CASSANDRA-13215)
19 * Add asm jar to build.xml for maven builds (CASSANDRA-11193)
```

# More Waiting 🕒



# RELEASE!





# Packages!

---

- Packages built & uploaded
- Update Ansible playbooks
- Run Ansible Playbooks
- Over and over and over and over.....
- ~~—PROFIT!~~
- Things suck less!

# TL;DR

---

- Local patches are awful for everyone
- Work with shops that value OSS (Like TLP!)

**@MissAmyTobey**

# TL;DR Part Deux ☐

---

- TNSTAAFL
- Pay your partners to push patches upstream!
- Tip your wait staff!

**@MissAmyTobey**

**Bye bye! 🙋 (We're hiring SREs & DBREs)**

---

**@MissAmyTobey**

**atobey@tenable.com**