

CASSANDRA JP 2015

PROBLEM FREE

CASSANDRA UPGRADES

Aaron Morton

@aaronmorton

Co-Founder & Team Member

**THE LAST PICKLE**

Licensed under a Creative Commons Attribution-NonCommercial 3.0 New Zealand License

About The Last Pickle.

**Work with clients to deliver and  
improve Apache Cassandra based  
solutions.**

**Based in New Zealand, USA,  
Norway.**

**Why Upgrade?**

When to Upgrade

Run Books

Upgrading from 1.2 to 2.0

Why Upgrade?

**Bug Fixes**

**New Features**

**Performance Improvements**

Why Upgrade?

**Because it is easy.**

Why Upgrade?

**Upgrading is a series of short  
term node failures.**

Why Upgrade?

Cassandra is designed to  
maintain availability when  
nodes fail.

Why Upgrade?  
**When to Upgrade**  
Run Books

Upgrading from 1.2 to 2.0



When To Upgrade

Upgrade before you need to  
upgrade.

When To Upgrade

Upgrade during business  
hours when people are at  
their best.

When To Upgrade

**If you are a startup upgrade to whatever release gives you features you need.**

(Make sure you test.)

When To Upgrade

**If you are a mature company  
generally avoid the first 6  
revisions of a major release.**

(Still do testing.)

When To Upgrade

**Ideally stay 1 or 2 revisions  
behind on the major version.**

Why Upgrade?  
When to Upgrade  
**Run Books**

Upgrading from 1.2 to 2.0

Run Books

Why are we doing this?

What are we doing?

How will we do it?

Run Books

Phase changes into chunks.

Describe the state of the system at the end of the Phase



Phases

Tasks that have checks and  
rollback.

cut-and-paste actions for  
operators.

Run Book Sections

**Overview**

**News**

**Changes**

**Tickets**

**Configuration**

**Schema**

Run Book Sections

**Current Situation**

**Assumptions**

**Risks**

**Phases**

Why Upgrade?  
When to Upgrade  
Run Books

**Upgrading from 1.2 to 2.0**

TLP Run Book

A detailed Run Book for this  
upgrade is available at  
<http://thelastpickle.com/blog/>

## Overview

- Prepare the cluster.
- To each node in turn:
  - Shutdown then node.
  - Upgrade the installation.
  - Restart the node.
- Upgrade on disk files.
- Cleanup.

# cassandra.yaml

Setting	Change
reduce_cache_sizes_at	Setting has been removed.
reduce_cache_capacity_to	Setting has been removed.
flush_largest_memtables_at	Setting has been removed.
native_transport_min_threads	Setting has been removed.

(There are more changes.)

# Schema Change

Setting	Change
read_repair_chance	The default has changed from 0.1 to 0.0.
local_read_repair_chance	The default has changed from 0.0 to 0.1.



News

Java 7 Required  
Cassandra 1.2.9 Required  
Blob Formats  
New LCS File Size  
Speculative Retry Added  
Configurable Metrics

Current Situation

```
nodetool status
```

```
nodetool version
```

```
nodetool gossipinfo
```

```
&& date
```

# Assumptions

Disk Space

No ERRORs log messages

Gossip Stable

No Dropped Messages

Backups Disabled

Repair Disabled

Assumptions

**On Disk File Format**  
**Minimum Version**  
**Monitoring**  
**Blob Constants**  
**Availability**

Risks

Latency  
Availability  
Incompatibility  
Disk Space

Phase 0. Check Assumptions

**Ensure the cluster works the  
way we think it does.**

Phase I. Prepare the cluster

# Snapshot

Update package repositories

Disable Hinted Handoff

Phase 2. Upgrade First Node to 2.0.14

**All nodes UP**

**Graceful shutdown**

**Clear CommitLog**

**Upgrade Package**



Phase 2. Upgrade First Node to 2.0.14

Update `cassandra.yaml`

Start Node

Tail `system.log`

watch `nodetool tpstats`

Phase 2. Upgrade First Node to 2.0.14

**Rollback (if needed)**

**Backup 2.0 SSTables.**

**Delete all data.**

**Restore from snapshot.**

**Repair.**

Phase 3. Upgrade Remaining Nodes to 2.0.14

**Ensure all nodes are UP**

**Rolling upgrade**

Phase 4. Cleanup

**Enabled Hinted Handoff**

Phase 5. Upgrade SSTables

```
nodetool clearsnapshot -t pre-2.0
```

```
nodetool upgradesstables
```

Phase 6. Restore Support Services

**Enabled Backups**

**Enabled Repair**

Thanks.

Aaron Morton  
@aaronmorton

Co-Founder & Team Member  
[www.thelastpickle.com](http://www.thelastpickle.com)

**THE LAST PICKLE**

Licensed under a Creative Commons Attribution-NonCommercial 3.0 New Zealand License