



SUPPORT APACHE CASSANDRA IN PRODUCTION

Anuj Wadehra
Architect & Cassandra SME
Ericsson R & D

AGENDA

- › Building In-House Support
- › Case Study
- › Challenges
- › Best Practices
- › How We Fixed It





BUILDING IN-HOUSE SUPPORT

BUILDING IN-HOUSE SUPPORT



Identify Vision

- Faster & Predictable Turnaround Time
- Cost-Effectiveness
- Open Source Contributions



Define Scope

- 24*7 On-Call Support
- Hot Fixes
- Consultancy
- Community Contributions



Build Competence

– What to learn ?

- > Cassandra
- > Programming
- > Troubleshooting
- > Product

– How to learn?

- > Documentation
- > Trainings
- > Community
- > Product issues
- > PoCs



Setup Infrastructure

- Continuous Integration setup
- Test Lab for reproducing defects
- R&D Lab for PoCs and new Feature Evaluation
- Collaboration tools
 - > Wiki
 - > Mailing Lists
 - > Forums etc.

CASE STUDY



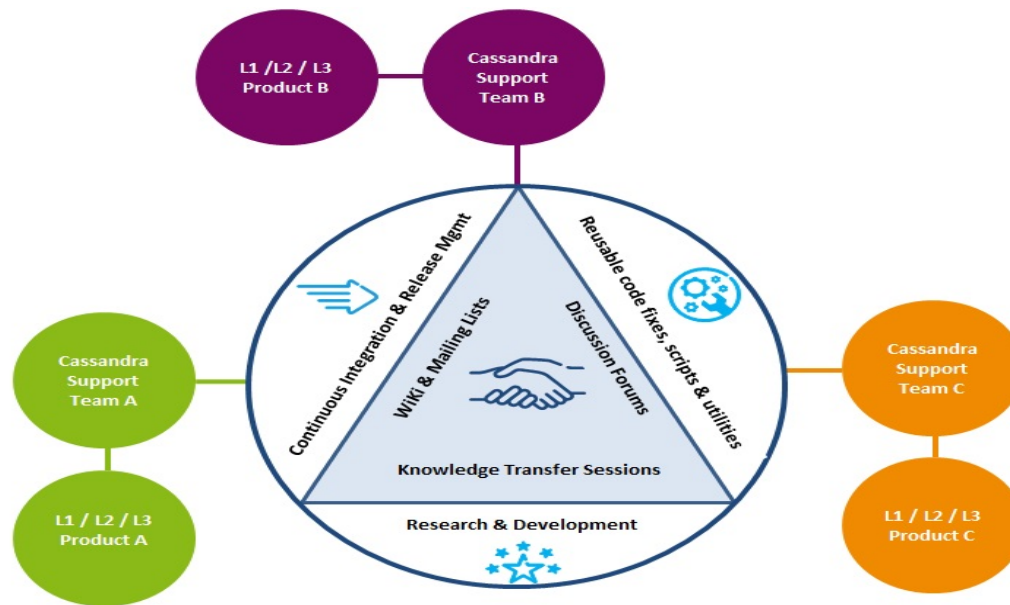
Cassandra In-House Support

X Technologies is a product based organization which has decided to replace RDBMS with C* in 3 product offerings. The decision is expected to positively impact scalability, fault-tolerance & availability of these products. Moreover, the organization aims to cut its expenditure on licensing and expensive RDBMS support.

CASE STUDY CONTD..



Proposed solution:



CHALLENGES



CHALLENGES



Our Biggest Challenges

- Keeping pace with Releases
 - Support for EOL versions
 - Insufficient/outdated docs
 - Lengthy upgrades
 - Compaction strategies
- Selling NoSQL concepts
 - > Why Repair ?
 - > How deleting data can increase disk space ?
 - > How RF determines fault-tolerance ?

CHALLENGES CONTD..



Addressing Challenges

- Stability over fancy features
- Upgrades aligned with EOL
- Well tested Compaction Strategies
- CI Infrastructure for Patches
- Well documented
 - > NoSql concepts
 - > Operational procedures
- Be Updated
 - > Mailing Lists
 - > JIRAs



BEST PRACTICES

C* Operations



BEST PRACTICES



› Daily Maintenance

- Backup
- Clearing Snapshots

› Repair

- Be Selective
- Repair once:
 - › Primary Range
 - › Incremental

› Monitoring

- Repair
- Dropped Mutations
- All time blocked
- Read/Write Latencies
- GC Pauses
- Disk Usage
- Upgrade sstables

BEST PRACTICES



› Automate

- Scaling/Down Sizing
- Multi-DC setup
- Upgrades
- Backup
- Restoration
- Clearing Snapshots
- Repairs

› Security

- SSL for JMX
- Encryption



HOW WE FIXED IT

C* Issues & Solutions



HOW WE FIXED IT



› **GC Pauses**

- Faster Memtable flushing
- Aggressive compactions
- JVM Tuning

› **Slow Table Scans**

- Spark

HOW WE FIXED IT



› Hung Repair

- Firewall timeout
- TCP Keep alive

› Disk Crunch

- Aggressive compaction
- LCS

HOW WE FIXED IT




› Performance & OOM

- Disable THP/Zone Reclaim
- Max map count

› Wide Rows

- Divide into buckets
- Add buckets when needed

QUESTIONS ?





ERICSSON