

DATASTAX:

Highly available, scalable and secure data wind Cassandra and DataStax Enterprise

GOTO Berlin 27<sup>th</sup> February 2014

### About Us

### DATASTAX

### Steve van den Berg

Regional Director Western Europe <u>www.datastax.com</u> @DataStaxEU



### **Johnny Miller** Solutions Architect



svandenberg@datastax.com

©2014 DataStax Confidential. Do not distribute without cons

@CyanMiller https://www.linkedin.com/in/johnnymiller jmiller@datastax.com

### DataStax

DATASTAX

- Founded in April 2010
- We drive Apache Cassandra™
- 400+ customers (24 of the Fortune 100)
- 220+ employees
- Contribute approximately 80% of the code to Cassandra
- Home to Apache Cassandra Chair & most committers
- Headquartered in San Francisco Bay area

©2014 DataStax Confidential. Do not distribute without consent.

• European headquarters established in London

### **Our Goal**

To be the first and best database choice for online applications

### оятязтях:

# Training

- Checkout the DataStax academy for free online virtual training!
  - <u>http://www.datastax.com/virtual-training</u>
- Public courses
  - <u>http://www.datastax.com/course-catalog</u>

Do not distribute without conser

• On-site training



http://www.datastax.com/training

# DataStax Enterprise for start-ups

- DataStax gives qualifying start-ups access to DataStax Enterprise for free!
- For more information:
  - http://www.datastax.com/startup

Do not distribute witho



### DataStax

• DataStax supports both the open source community and enterprises.

### Open Source/Community

- Apache Cassandra (employ Cassandra chair and 80+% of the committers)
- DataStax Community Edition
- DataStax OpsCenter
- DataStax DevCenter
- DataStax Drivers/Connectors
- Online Documentation
- Online Training
- Mailing lists and forums

### Enterprise Software

- DataStax Enterprise Edition
  - Certified Cassandra
  - Built-in Analytics
  - Built-in Enterprise Search
  - Enterprise Security
- DataStax OpsCenter
- Expert Support
- Consultative Help
- Professional Training

### DATASTAX

### DataStax Enterprise

- **DataStax Enterprise:** LOB\* applications, with analytics and search for online/real-time application data.
- Hadoop: data warehouse applications with analytics and search for data warehouse.



\*Line of business

©2014 DataStax Confidential. Do not distribute without consent

# Availability and Speed Matters for online apps!

- UK retailers **lost 8.5 billion last year to slow web sites**, which is 1 million for every 10 million in online sales
- Over half of all web users **expect a response time of 2 seconds or less**
- A 1 second delay causes a nearly 10% reduction in customer interactions
- A 1 second decrease in Amazon page load time costs the company \$1.6 billion in sales



### DATASTAX

### Apache Cassandra™

- Apache Cassandra™ is a massively scalable, open source, NoSQL, distributed database built for modern, mission-critical online applications.
- Written in Java and is a hybrid of Amazon Dynamo and Google BigTable
- Masterless with no single point of failure

Do not distribute without conse

- Distributed and data centre aware
- 100% uptime

©2014 DataStax Confidential

• Predictable scaling



### **Cassandra** – Core Values

- Fase of use
- Massive scalability •
- High performance •
- Always Available

Nathan Milford @NathanMilford

©2014 DataStax Confidential. Do not distribute without consent

Follow

Man, I just love Cassandra. Lost a data center Hurricane Sandy, nodes came up and started working with no pain.





Following

Follow

1-

1-

"Cassandra ... dealt with the loss of one third of its regional nodes without any loss of data or availability." techblog.netflix.com/2012/07/lesson... -

Nice!

**Jake Luciani** @tiake

1-Following DATASTAX

We've lost multiple SSDs in our #Cassandra cluster and the JBOD support in C\* 1.2 kept the nodes running while we swapped them out live!



Follow

Coming to the conclusion that #cassandra is kind of indestructible. "Robust" doesn't do it justice.

### Cassandra – Performance and Scale

"In terms of scalability, there is a clear winner throughout our experiments. Cassandra achieves the highest throughput for the maximum number of nodes in all experiments with a linear increasing throughput."

Solving Big Data Challenges for Enterprise Application Performance Management, Tilman Rable, et al., August 2012. Benchmark paper presented at the Very Large Database Conference, 2012. http://vldb.org/pvldb/vol5/p1724\_tilmannrabl\_vldb2012.pdf

### Solving Big Data Challenges for Enterprise Applicatio Performance Management





### Cassandra - Performance and Scale

- Cassandra works for small to huge deployments.
- Cassandra Footprint @ Netflix
  - 80+ Clusters
  - 2500+ nodes
  - 4 Data Centres (Amazon Regions)
  - > 1 Trillion transactions per day



DATAST

http://planetcassandra.org/functional-use-cases/

©2014 DataStax Confidential. Do not distribute without consent

### DATASTAX

### Cassandra – Overview

- Cassandra was designed with the understanding that system/hardware failures can and do occur
- Peer-to-peer, distributed system

©2014 DataStax Confidential. Do not distribute without consent

- All nodes the same
- Data partitioned among all nodes in the cluster
- Custom data replication to ensure fault tolerance
- Read/Write-anywhere and across data centres



# Cassandra – More Than One Server

- All nodes participate in a cluster
- Add or remove as needed
- All nodes the same masterless with no single point of failure

DATAST

Node 2

Node 3

Node 1

Node 4

- Each node communicates with each other through the Gossip protocol, which exchanges information across the cluster every second
- Data partitioned among all nodes in the cluster

Do not distribute without consen

- A commit log is used on each node to capture write activity. Node 5
  Data durability is assured
- Data also written to an in-memory structure (memtable) and then to disk once the memory structure is full (an SStable)
- More capacity? Add a server!
- More throughput? Add a server!

# Cassandra - Locally Distributed

- Client reads or writes to any node
- Node coordinates with others
- Data read or replicated in parallel
- Replication factor (RF): How many copies of your data?
- RF = 3 in this example
- Each node is storing 60% of the clusters total data i.e. 3/5

Handy Calculator: http://www.ecyrd.com/cassandracalculator/

©2014 DataStax Confidential. Do not distribute without consent



### Cassandra – Rack Aware

- Cassandra is aware of which rack (or availability zone) each node resides in.
- It will attempt to place each data copy in a different rack.
- RF = 3 in this example

©2014 DataStax Confidential. Do not distribute without consen



### Cassandra – Data Centre Aware



### **Cassandra** – Tunable Consistency Consistency Level (CL) • **Client specifies per read or write** Handles multi-data center operations • 5 µs ack Node 1 1<sup>st</sup> copy 12 µs ack ALL = All replicas ack • Write CL=QUORUM Parallel Write QUORUM = > 51% of replicas ack • Node 2 Node 5 2<sup>nd</sup> copy LOCAL\_QUORUM = > 51% in **local DC** ack • ONE = Only one replica acks ٠ 12 µs ack Plus more.... (see docs) ٠ Node 3 3<sup>rd</sup> copy 500 µs ack Blog: Eventual Consistency != Hopeful Consistency ٠ http://planetcassandra.org/blog/post/a-netflix-experiment-eventual-consistency-hopeful-consistency-by-christos-kalantzis/

©2014 DataStax Confidential. Do not distribute without cons

# Node Failure

- A single node failure shouldn't bring failure.
- Replication Factor + Consistency Level = Success
- This example:
  - RF = 3
  - CL = QUORUM

©2014 DataStax Confidential. Do not distribute without consent.





### Rack Failure

- Cassandra will place the data in as many different racks or availability zones as it can.
- This example:
  - RF = 3
  - CL = QUORUM
  - Rack 2 fails
- Data copies still available in Node 1 and Node 5
- Quorum can be honored i.e. > **51% ack**

©2014 DataStax Confidential. Do not distribute without consent



### оятязтях

# Don't be afraid of Weak Consistency

Do not distribute without conse

- More tolerant to failure
- Consistency Level of 1 is the most popular (I think)
- If you want stronger consistency go for LOCAL\_QUORUM i.e. quorum is honored in the local data centre.
- If you go stronger than LOCAL\_QUORUM understand what this means and why you are doing it.
- Remember you can have different consistency levels for reads and writes e.g. write with CL:1, read with CL:LOCAL\_QUORUM











# **Cassandra Clients** - Native Driver

- Clients that use the native driver also have access to various policies that enable the client to intelligently route requests as required.
- This includes:
  - Load Balancing
    - Data Centre Aware
    - Latency Aware
    - Token Aware
  - Reconnection policies
  - Retry policies
    - Downgrading Consistency
    - Plus others..

©2014 DataStax Confidential. Do not distribute without consent

<u>http://www.datastax.com/download/clientdrivers</u>



### Quotes

©2014 DataStax Confidential. Do not distribute without consen

 "Cassandra, our distributed cloud persistence store which is distributed across all zones and regions, dealt with the loss of one third of its regional nodes without any loss of data or availability". http://techblog.netflix.com/2012/07/lessons-netflix-learned-from-aws-storm.html

 "During Hurricane Sandy, we lost an entire data center. Completely. Lost. It. Our application fail-over resulted in us losing just a few moments of serving requests for a particular region of the country, but our data in Cassandra never went offline." <u>http://planetcassandra.org/blog/post/outbrain-touches-over-80-of-all-us-online-users-with-help-from-cassandra/</u>



NETFLIX





# **Ring-fenced resources**

- If you need to isolate resources for different uses, Cassandra is a great fit.
- You can create separate virtual data centres optimised as required different workloads, hardware, availability etc..
- Cassandra will replicate the data for you no ETL is necessary



### DATASTAX:

# Hybrid Cloud

- DataStax Enterprise and Cassandra are multi-data centre and cloud capable
- Data written to any node is automatically and transparently written to all other nodes in multiple data centres i.e. **no etl**



# Security in Cassandra

### DATASTAX



**Internal Authentication** Manages login IDs and passwords inside the database

- + Ensures only authorized users can access a database system using internal validation
- + Simple to implement and easy to understand
- + No learning curve from the relational world



**Object Permission Management** controls who has access to what and who can do what in the database

- + Provides granular based control over who can add/change/ delete/read data
- + Uses familiar GRANT/REVOKE from relational systems
- + No learning curve



**Client to Node Encryption** protects data in flight to and from a database cluster

- + Ensures data cannot be captured/stolen in route to a server
- + Data is safe both in flight from/ to a database and on the database; complete coverage is ensured

FEATURES

BENEFITS

# Advanced Security in DataStax Enterprise

### онтнатнх:



**External Authentication** uses external security software packages to control security

- + Only authorized users have access to a database system using external validation
- + Uses most trusted external security packages (Kerberos), mainstays in government and finance
- + Single sign on to all data domains



Transparent Data Encryption encrypts data at rest

- + Protects sensitive data at rest from theft and from being read at the file system level
- + No changes needed at application level



Data Auditing provides trail of who did and looked at what/when

- + Supplies admins with an audit trail of all accesses and changes
- + Granular control to audit only what's needed
- + Uses log4j interface to ensure performance and efficient audit operations

FEATURES

BENEFITS

# Data Replication Security in Cassandra

• A popular feature from a data security perspective is the ability to control at a keyspace/schema level which data centres data should be replicated to.

- What this means is that in a multi-data centre (both physical and virtual) cluster you can ensure that data is not shipped anywhere it shouldn't be and access to that data can be controlled.
- This is very simple to set-up and is extremely useful when you need to share some of your data, but not all of you data or if you have requirements around where your data is permitted to reside.



### оятязтях.

# DataStax Enterprise 4.0, OpsCentre 4.1

- DataStax Enterprise 4.0
  - **New in-memory option**. Brings all of the goodness of Cassandra to an in-memory database
  - Production-certified version of Apache Cassandra (2.0)
  - Enterprise search enhancements
- OpsCenter 4.1
  - Capacity planning updates

©2014 DataStax Confidential. Do not distribute without consent

- Better insight into node performance
- More information:

http://www.datastax.com/wp-content/uploads/2014/02/WP-WhatsNewDSE40.pdf

### http://www.datastax.com/download

DATASTA

# How does in-memory work?

- Developers can create new tables to be in-memory or alter existing tables to be in-memory
- Writes are **durable**
- 10-100x improvement

©2014 DataStax Confidential. Do not distribute without consen



### DATASTAX:

# **Find Out More**

•

•

•

٠

•

- DataStax: <u>http://www.datastax.com</u>
- Getting Started: <a href="http://www.datastax.com/documentation/gettingstarted/index.html">http://www.datastax.com/documentation/gettingstarted/index.html</a>
  - Training: <u>http://www.datatstax.com/training</u>
  - Downloads: <u>http://www.datastax.com/download</u>
    - Documentation: <u>http://www.datastax.com/docs</u>
    - Developer Blog: <u>http://www.datastax.com/dev/blog</u>
    - Community Site: <u>http://planetcassandra.org</u>
      - http://planetcassandra.org/Learn/CassandraCommunityWebinars
- Summit Talks: <u>http</u>

©2014 DataStax Confidential. Do not distribute without consen

Webinars:

http://planetcassandra.org/Learn/CassandraSummit



# Thank You

©2014 DataStax Confidential. Do not distribute without conser



# We power the big data apps that transform business.