Distributed Computing

at rackspace

Hai.Thai@rackspace.com

About: Me

ME

About: Me

- 09 Tech grad
- B.S. Computer Engineering
- 4 years at rackspace

About: Rackspace

About: Rackspace

- Managed + Cloud hosting
- Cloud Applications:
  - Email

About: Rackspace

- Office in Blacksburg
- 100 best companies to work for
- We’re hiring!

The Big Picture

Data is VALUABLE

Data is growing

- More sources + more data per source
- Faster than individual devices
- Years of information

The Big Picture: Rackspace

At Rackspace e-mail

- 2.5 Million mailboxes
- 50-100 Million messages / day
- 300-400 GB raw log data / day
- Hundreds of servers
- TBs of stored log data

The Big Picture: Rackspace

How do we…

- Aggregate
- Store
- Analyze
- Access

The Big Picture: Rackspace

How do we…

Get Value?

The Problem

With mail logs, we can:

- Help customers
- Diagnose the system
- Understand and plan

Aggregation

- Multi-Source Single-Sink
- Realworld network
- Hardware Failure
- 

Storage

- Distributed
- Fault tolerant
- Horizontally scalable
- Easy

Serving Logs

Make logs accessible for:

- **Support to help customers**
- **Operations to diagnose errors**

Serving Logs

The challenge: Volume

- **400+ GB / day = 300 MB / min**
- **Must be timely**
- **Related log data may be disjoint**

Serving Logs

- Index data with Hadoop MapReduce
- Serve indexes in Solr

Serving Logs: Indexing

Map Reduce:

- History on distributed systems:
  - Google
- Easily distributed
- Map step: key->value pair
- Reduce step: All values for a key

Serving Logs: Indexing

Map Reduce for mail logs:

- **Map step:**
  - Parse raw log
- **Reduce step:**
  - Aggregate related log lines
  - Generate relevant structure for queries
  - Output as Solr index

Serving Logs: Indexing

```
Nov 12 17:36:54 gate8.gate.sat.mlsrvr.com postfix/smtpd[2552]: connect from hostname

Nov 12 17:36:54 relay2.relay.sat.mlsrvr.com postfix/qmgr[9489]: 1DBD21B48AE:
from=<mapreduce@mailtrust.com>, size=5950, nrcpt=1 (queue active)

Nov 12 17:36:54 relay2.relay.sat.mlsrvr.com postfix/smtpd[28085]: disconnect from
hostname

Nov 12 17:36:54 gate5.gate.sat.mlsrvr.com postfix/smtpd[22593]: too many errors after
DATA from hostname

Nov 12 17:36:54 gate2.gate.sat.mlsrvr.com postfix/smtp[15928]: 732196384ED:
to=<mapreduce@mailtrust.com>, relay=hostname[ip], conn_use=2, delay=0.69,
delays=0.04/0.44/0.04/0.17, dsn=2.0.0, status=sent (250 2.0.0 Ok: queued as

02E1544C005)

Nov 12 17:36:54 gate5.gate.sat.mlsrvr.com postfix/smtpd[22593]: disconnect from
hostnameNov 12 17:36:54 gate10.gate.sat.mlsrvr.com postfix/smtpd[10311]: connect from
hostname

Nov 12 17:36:54 relay2.relay.sat.mlsrvr.com postfix/smtp[28107]: D42001B48B5:
to=<mapreduce@mailtrust.com>, relay=hostname[ip], delay=0.32, delays=0.28/0/0/0.04,
dsn=2.0.0, status=sent (250 2.0.0 Ok: queued as 1DBD21B48AE)

Nov 12 17:36:54 gate20.gate.sat.mlsrvr.com postfix/smtpd[27168]: disconnect from hostname

Nov 12 17:36:54 gate5.gate.sat.mlsrvr.com postfix/qmgr[1209]: 645965A0224: removed

Nov 12 17:36:54 gate2.gate.sat.mlsrvr.com postfix/qmgr[13764]: 732196384ED: removed

Nov 12 17:36:54 gate1.gate.sat.mlsrvr.com postfix/smtpd[26394]: NOQUEUE: reject: RCPT
from hostname 554 5.7.1 <mapreduce@mailtrust.com>: Client host rejected: The sender's
mail server is blocked; from=<mapreduce@mailtrust.com>  to=<mapreduce@mailtrust.com>
proto=ESMTP helo=<mapreduce@mailtrust.com>
```

Serving Logs: Indexing

Nov 12 17:36:54 gate8.gate.sat.mlsrvr.com postfix/smtpd[2552]: connect from hostname

Nov 12 17:36:54 relay2.relay.sat.mlsrvr.com postfix/qmgr[9489]: 1DBD21B48AE: from=<mapreduce@mailtrust.com>, size=5950, nrcpt=1 (queue active)

Nov 12 17:36:54 relay2.relay.sat.mlsrvr.com postfix/smtpd[28085]: disconnect from hostname

Nov 12 17:36:54 gate5.gate.sat.mlsrvr.com postfix/smtpd[22593]: too many errors after DATA from hostname

Nov 12 17:36:54 gate2.gate.sat.mlsrvr.com postfix/smtp[15928]: 732196384ED: to=<mapreduce@mailtrust.com>, relay=hostname[ip], conn_use=2, delay=0.69, delays=0.04/0.44/0.04/0.17, dsn=2.0.0, status=sent (250 2.0.0 Ok: queued as

02E1544C005)

Nov 12 17:36:54 gate5.gate.sat.mlsrvr.com postfix/smtpd[22593]: disconnect from hostnameNov 12 17:36:54 gate10.gate.sat.mlsrvr.com postfix/smtpd[10311]: connect from hostname

Nov 12 17:36:54 relay2.relay.sat.mlsrvr.com postfix/smtp[28107]: D42001B48B5: to=<mapreduce@mailtrust.com>, relay=hostname[ip], delay=0.32, delays=0.28/0/0/0.04, dsn=2.0.0, status=sent (250 2.0.0 Ok: queued as 1DBD21B48AE)

Nov 12 17:36:54 gate20.gate.sat.mlsrvr.com postfix/smtpd[27168]: disconnect from hostname

Nov 12 17:36:54 gate5.gate.sat.mlsrvr.com postfix/qmgr[1209]: 645965A0224: removed

Nov 12 17:36:54 gate2.gate.sat.mlsrvr.com postfix/qmgr[13764]: 732196384ED: removed

Nov 12 17:36:54 gate1.gate.sat.mlsrvr.com postfix/smtpd[26394]: NOQUEUE: reject: RCPT from hostname 554 5.7.1 <mapreduce@mailtrust.com>: Client host rejected: The sender's mail server is blocked; from=<mapreduce@mailtrust.com> to=<mapreduce@mailtrust.com> proto=ESMTP helo=<mapreduce@mailtrust.com>

Serving Logs: Searching

- Full text search + advanced search features
- Supports distributed operation
- Horizontally scalable

Serving Logs: Searching

Our Solr cluster:

- Separate from hadoop
  - Pulls indexed data and merges into memory
- Subset of logs searchable
- Shard data based on time

Hadoop Map Reduce

- Large sets of data
  - 100s of GBs per job; potentially TBs
- Full power of Map Reduce
- Hadoop Streaming

Challenges

Building on top of HDFS

- Easy, but simple
- Custom organization on top of filesystem

In Flight Refactor

- Original design assumed perfect information
- Redesign around delayed logs/events

Challenges

- **Parsing Application Logs Requires Domain Knowledge**

- **Develop services based on distributed systems for solutions to use rather than solutions build around technology**

The Future

- Streaming vs Batching
- Solr Cloud
- New Logging solution

Takeaway

- Use of Hadoop + Map Reduce to solve our data problem
- Solutions must be created to extract value from growing data
- Example of a realworld distributed system

Distributed Systems

Big Data is only one of the areas of growth in distributed systems

We need YOU

RackerTalent.com

Resources

- lucene.apache.org/solr
- hadoop.apache.org
- Hadoop: The Definitive Guide