



# Firehose Engineering

*designing  
high-volume  
data collection  
systems*

# Firehose Database Applications (FDA)

- (1) very high volume of data input from many automated producers
- (2) continuous processing of incoming data


# mozilla crash reports

Find Crash ID or Signature

Product: **Firefox** Current Versions Report: **Overview** [Advanced Search](#)

Firefox Crash Data **3 days** **7 days** 14 days

## Crashes per 100 Active Daily Users



## Crash Reports

**Firefox 9.0a1**

Top Crashers

Top Changers

Top Plugin Crashers

# Mozilla Socorro

**Firefox 7.0.2**

Top Crashers

Top Changers

Top Plugin Crashers

**Firefox 7.0.1**

Top Crashers

Top Changers

Top Plugin Crashers


**Firefox 6.0.2**

Top Crashers


Top Changers

Top Plugin Crashers

# Upwind



# Fraud Detection System



# Firehose Challenges



# 1. Volume



- 100's to 1000's facts/second
- GB/hour



# 1. Volume



- spikes in volume
- multiple uncoordinated sources





# 1. Volume



*volume always grows over time*



## 2. Constant flow

since data arrives 24/7 ...

*while the user interface can be down, data collection can never be down*

## 2. Constant flow



- can't stop receiving to process
- data can arrive out of order

# 3. Database size

- terabytes to petabytes
  - lots of hardware
  - single-node DBMSes aren't enough
  - difficult backups, redundancy, migration
  - analytics are resource-consumptive

# 3. Database size

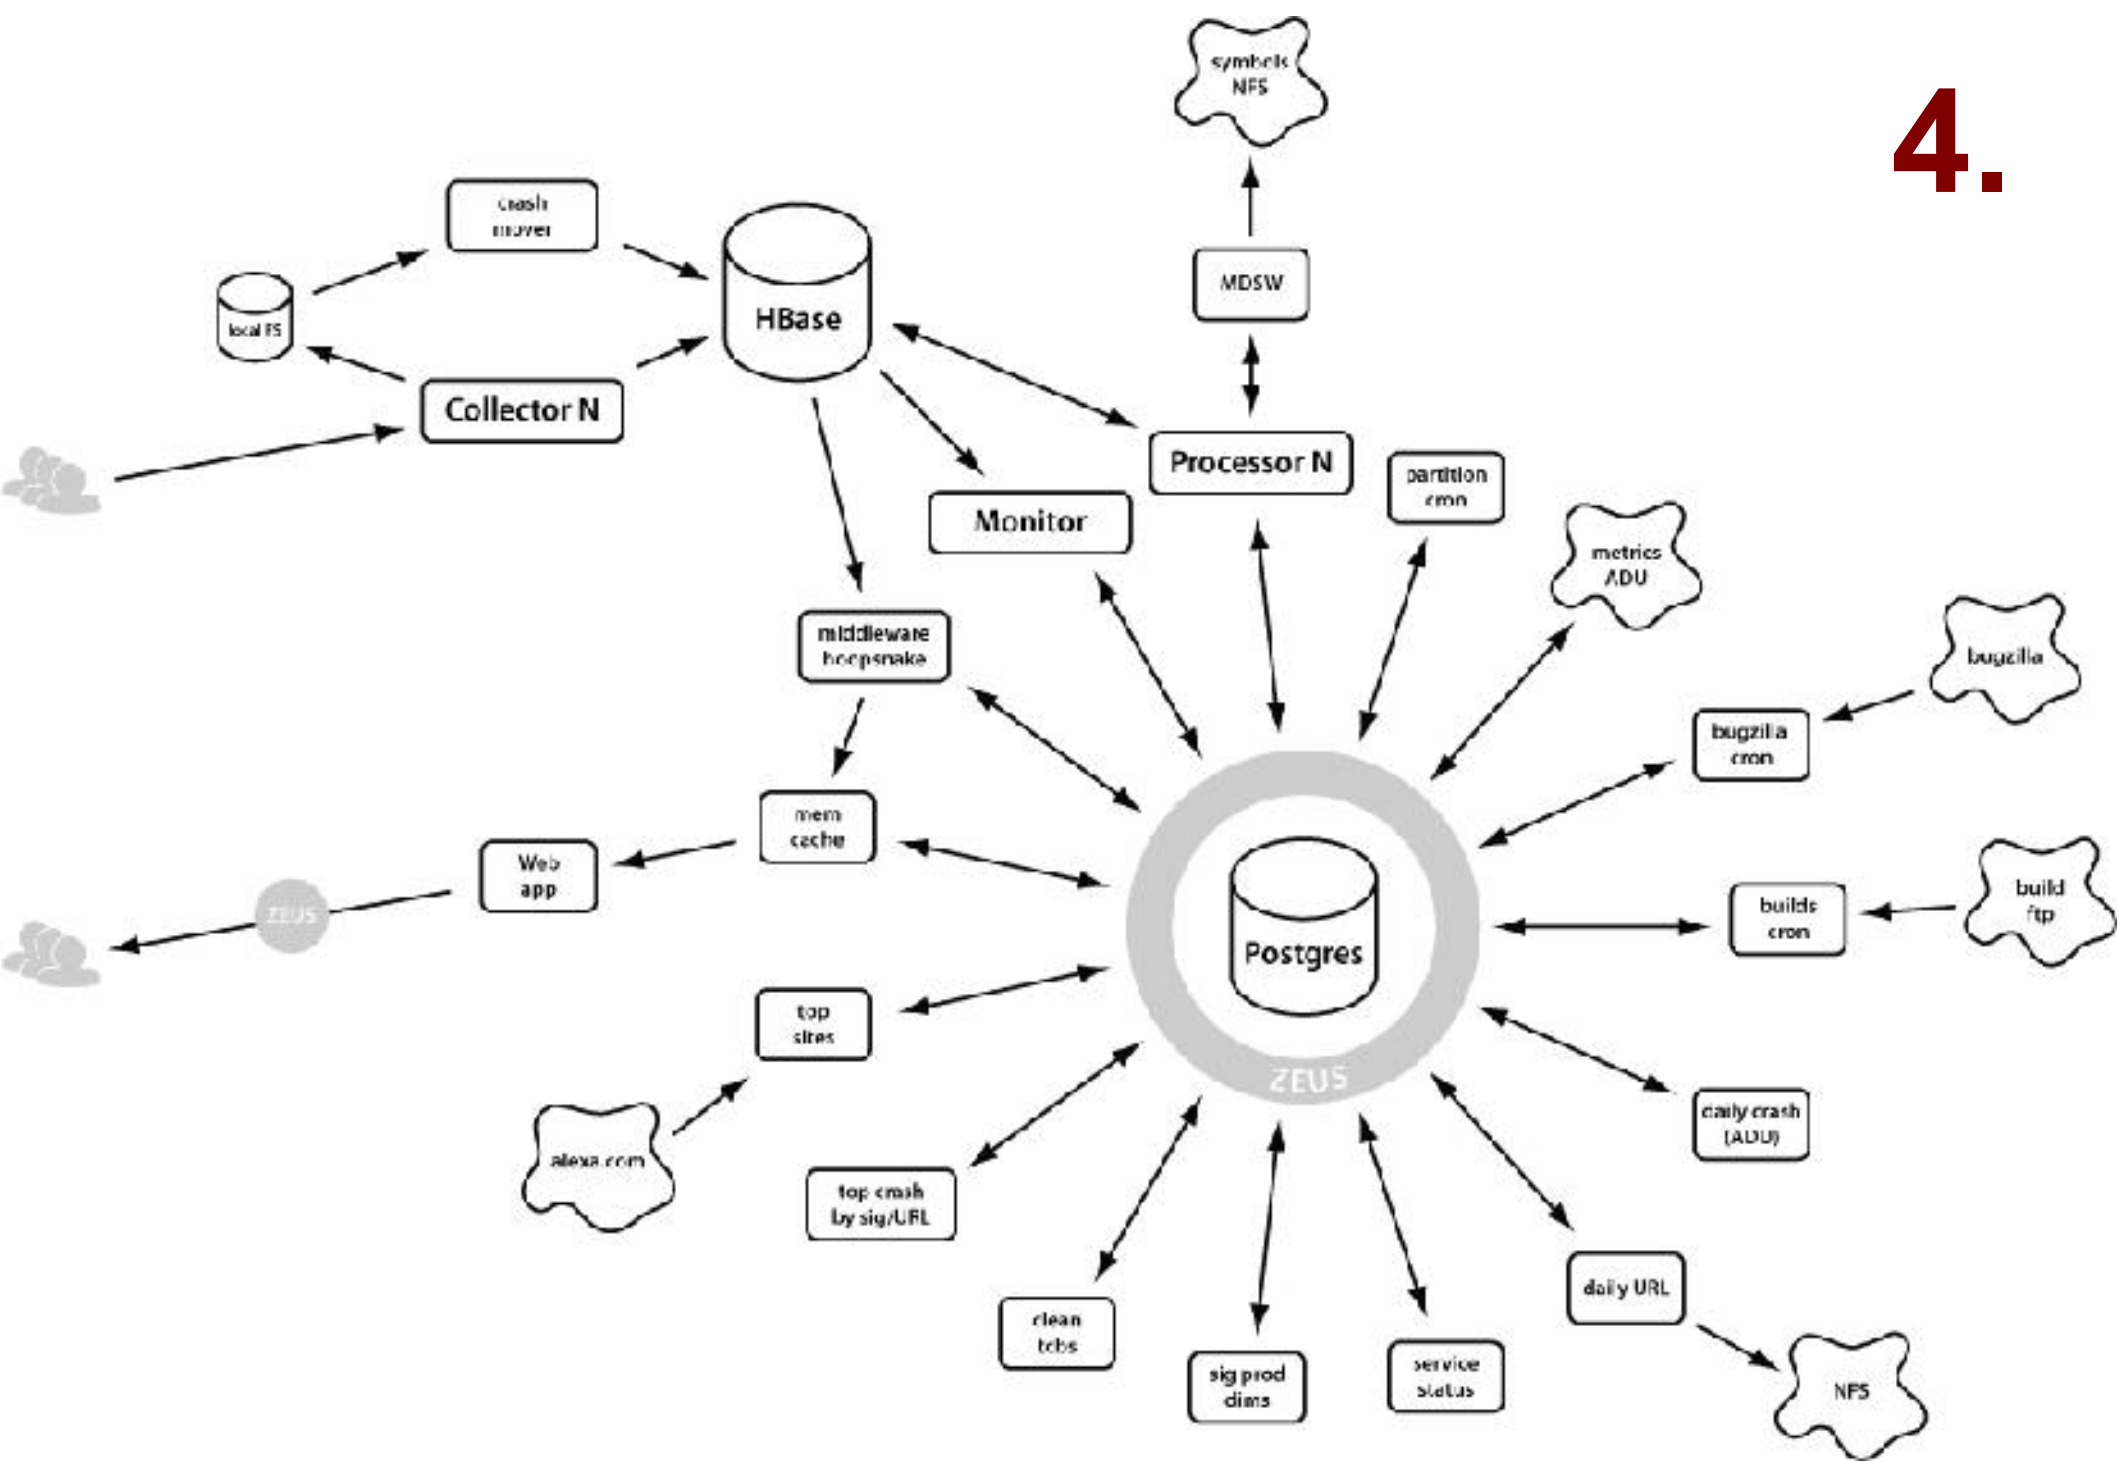
- database growth
  - size grows quickly
  - need to expand storage
  - estimate target data size
  - create data ageing policies

# 3. Database size

*“We will decide on a data retention policy when we run out of disk space.”*

– every business user everywhere

4.





**many components  
= many failures**



# 4. Component failure

- all components fail
  - or need scheduled downtime
  - including the network
- collection must continue
- collection & processing must recover



**solving firehose problems**

# socorro project



Mozilla Crash Reporter

**We're Sorry**

Firefox had a problem and crashed. We'll try to restore your tabs and windows when it restarts.

To help us diagnose and fix the problem, you can send us a crash report.

Tell Mozilla about this crash so they can fix it

[Details...](#)

Include the address of the page I was on

Email me when more information is available


Your crash report will be submitted before you quit or restart.

[Quit Firefox](#) [Restart Firefox](#)

# http://crash-stats.mozilla.com

Firefox Crash Data 3 days **7 days** 14 days

### Crashes per 100 Active Daily Users



### Crash Reports

**Firefox 9.0a1**

[Top Crashers](#)

[Top Changers](#)

[Top Plugin Crashers](#)

**Firefox 8.0a2**

[Top Crashers](#)

[Top Changers](#)

[Top Plugin Crashers](#)

**Firefox 7.0**

[Top Crashers](#)

[Top Changers](#)

[Top Plugin Crashers](#)

**Firefox 6.0.2**

[Top Crashers](#)

[Top Changers](#)

[Top Plugin Crashers](#)



# mozilla crash reports

Find Crash ID or Signature

Product: **Firefox** 6.0.2 Report: **Top Crashers**

[Advanced Search](#)

## Top Crashers for Firefox 6.0.2

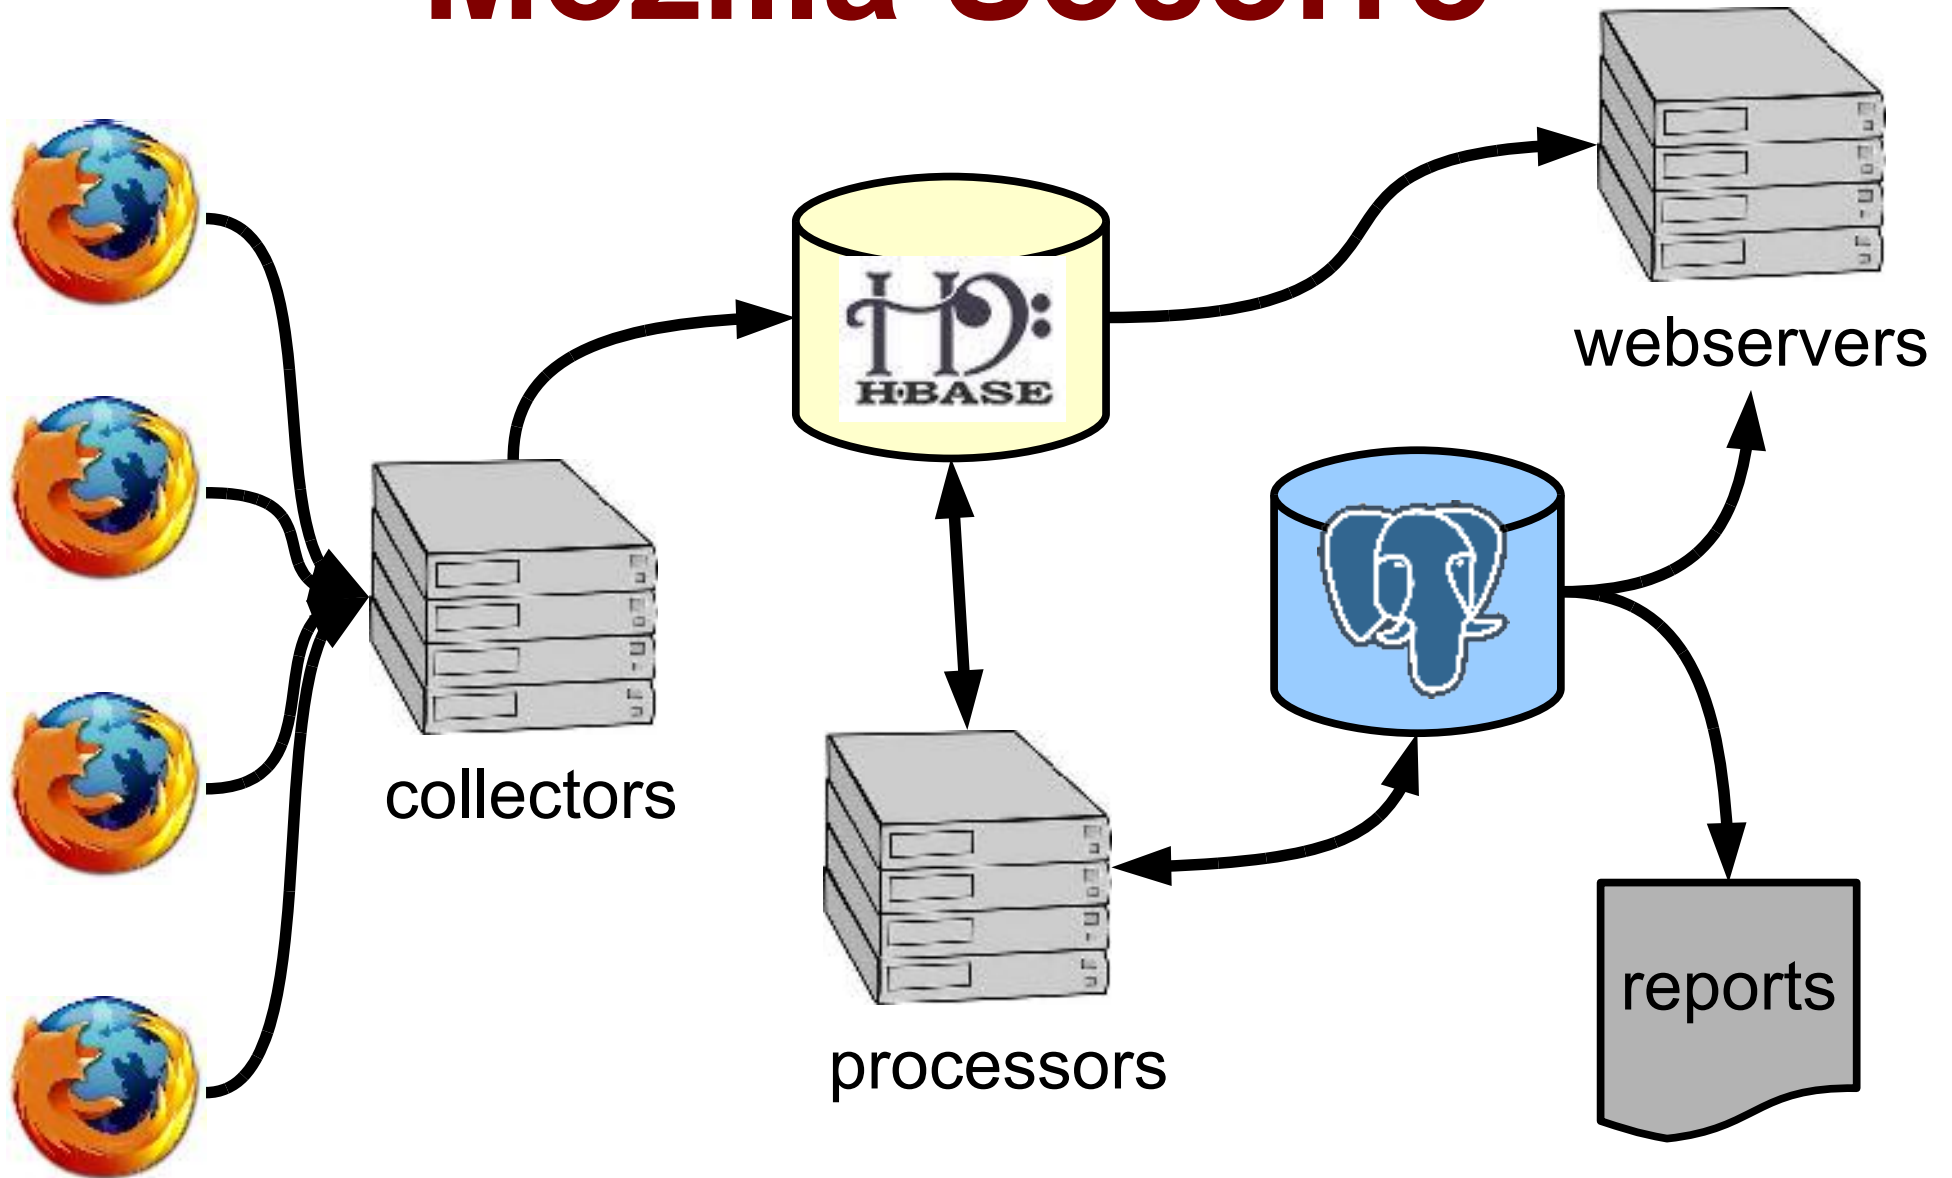
By Signature By URL By Domain By Topsite

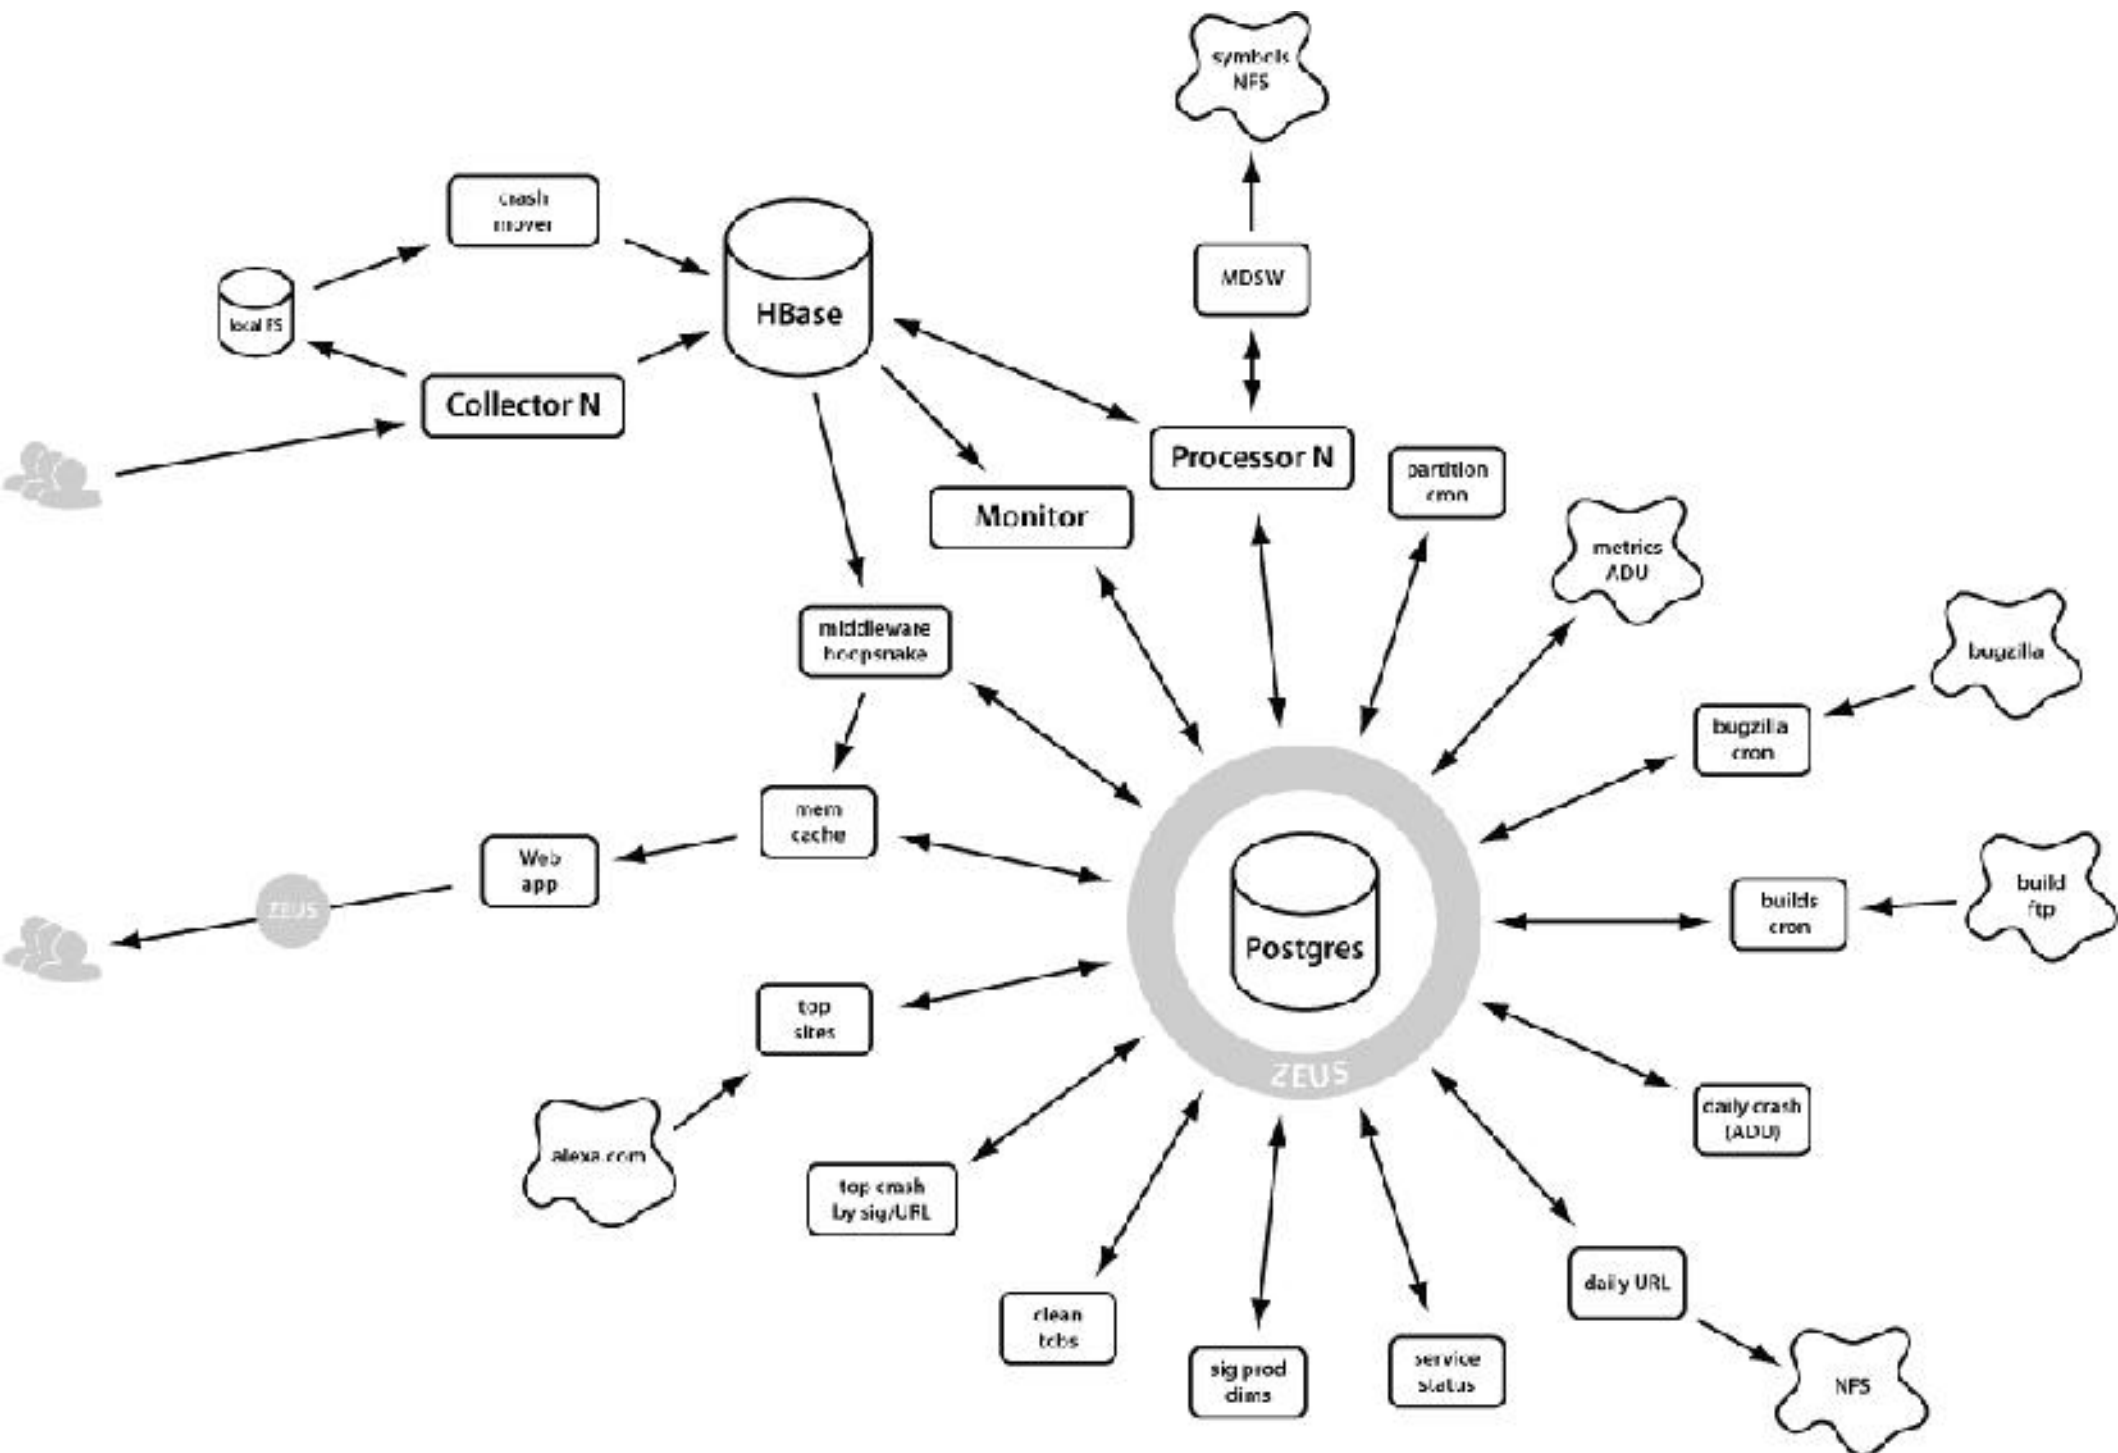
Top 300 Crashing Signatures. 2011-09-21 through 2011-09-28. The report covers 89.60% of all 775208 crashes during this period. Graphs below are dual-axis, having Count (Number of Crashes) on the left X axis and Percent of total of Crashes on the right X axis.

Type: Days: All **Browser** Plugin 1 3 7 14 28

Rank	Trend	%	Diff	Signature	Count	Win	Mac	Lin	Ver	First Appearance	Bugzilla IDs
1	0	10.80%	0.27%	<a href="#">[empty signature]</a> <a href="#">Learn More</a>	82606	73	2	0	-		
2	1	3.49%	0.04%	<a href="#">hang   mozilla::plugins::PPluginInstanceParent::CallPBrowserStreamConstructor(mozilla::</a>	27084	27084	0	0	63	2011-01-01	<a href="#">672111</a> , <a href="#">574084</a> , <a href="#">566062</a> , <a href="#">More</a>
3	1	2.01%	-0.04%	<a href="#">hang   mozilla::plugins::PPluginInstanceParent::CallNP_HandleEvent(mozilla::plugins::N</a>	22573	22573	0	0	62	2011-01-01	<a href="#">688092</a> , <a href="#">647400</a> , <a href="#">667036</a> , <a href="#">More</a>
4	1	2.50%	0.00%	<a href="#">hang   mozilla::plugins::PPluginScriptableObjectParent::CallHasProperty(mozilla::plugins::</a>	19394	19394	0	0	60	2011-01-01	<a href="#">678046</a> , <a href="#">669000</a> , <a href="#">664826</a> , <a href="#">More</a>
5	-3	2.09%	-0.00%	<a href="#">PInvokeCallWorker</a>	16195	16195	0	0	3	2011-02-04	<a href="#">664740</a> , <a href="#">More</a>

# Mozilla Socorro



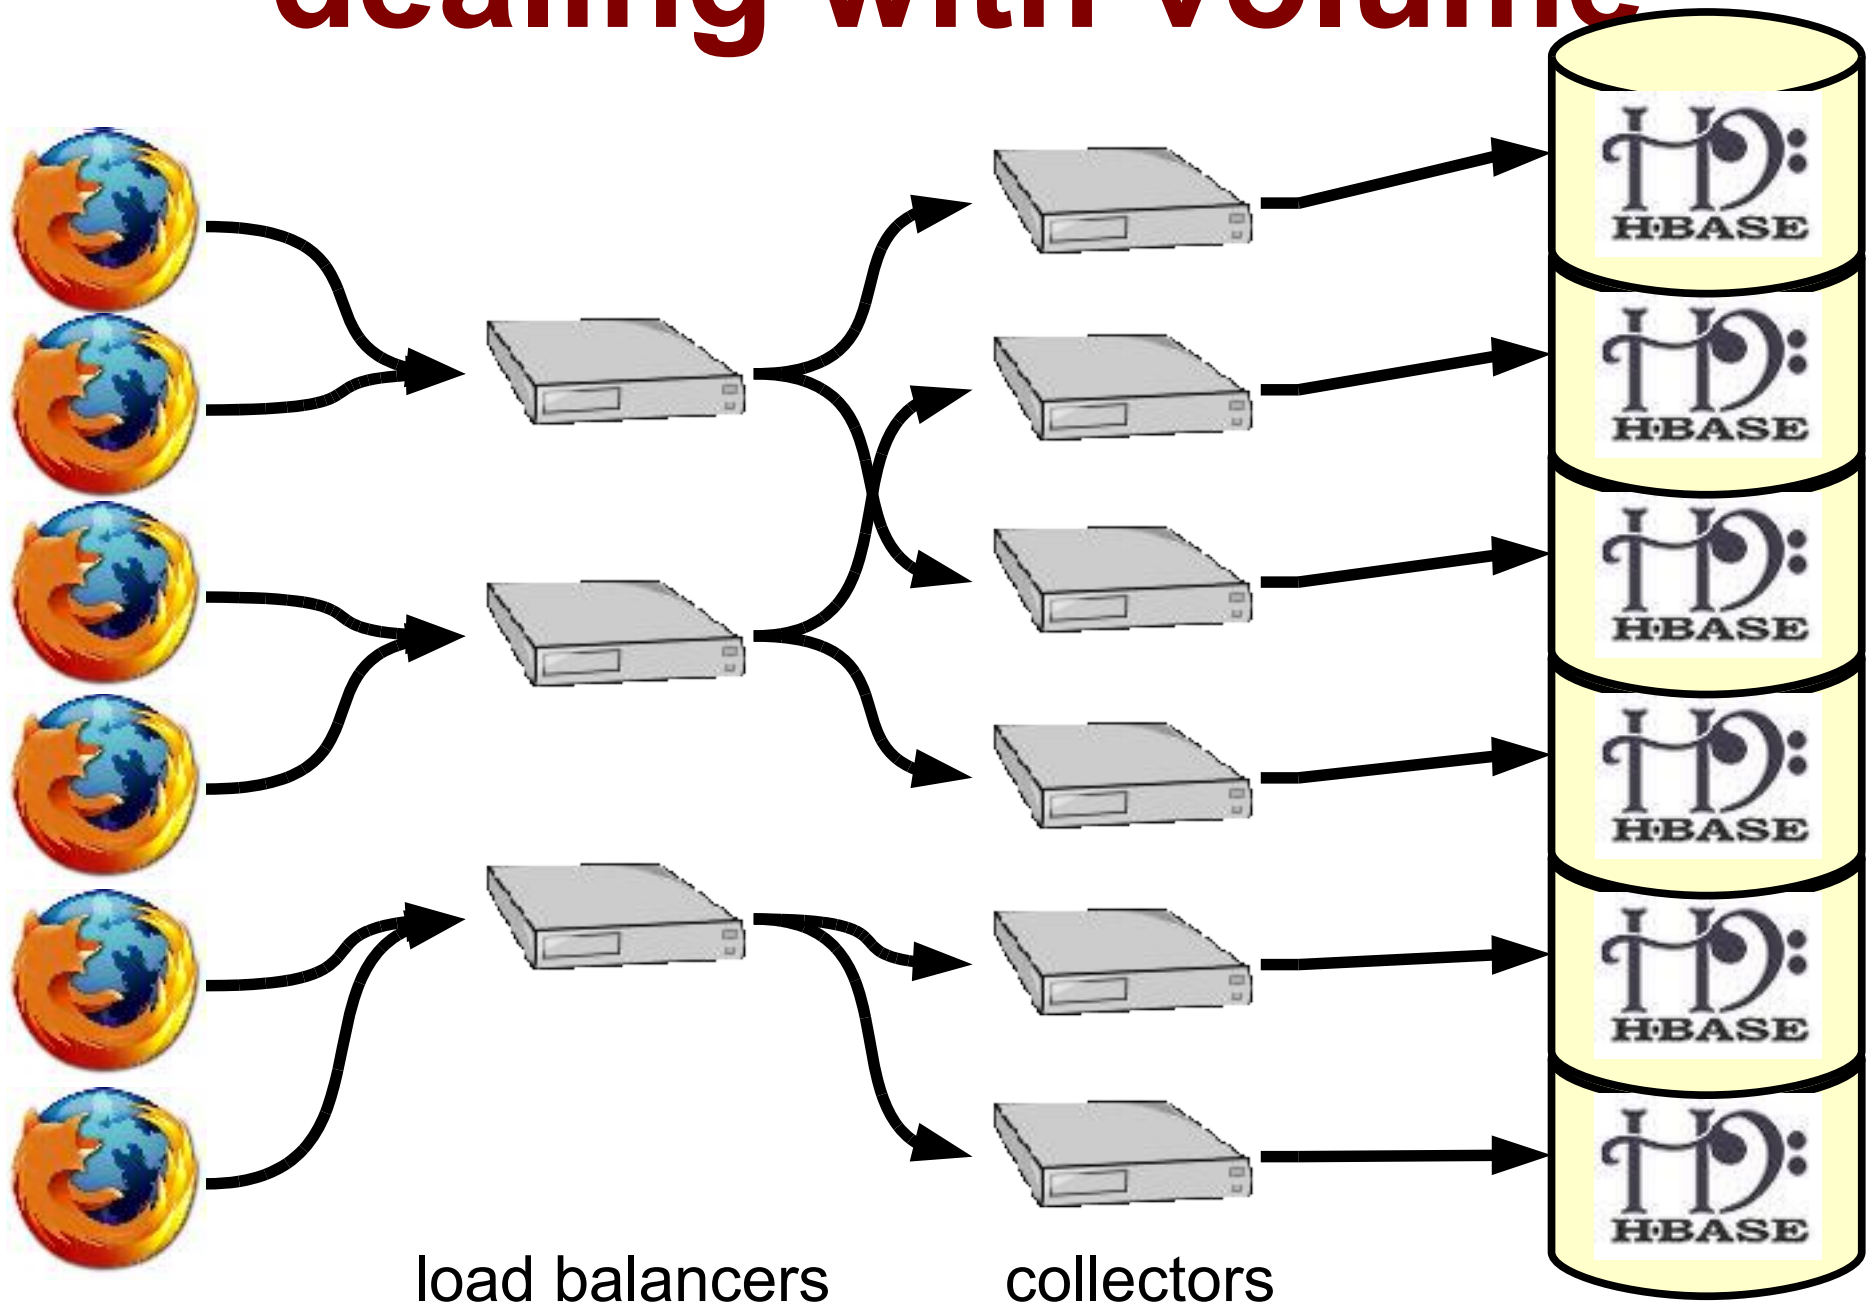


# socorro data volume

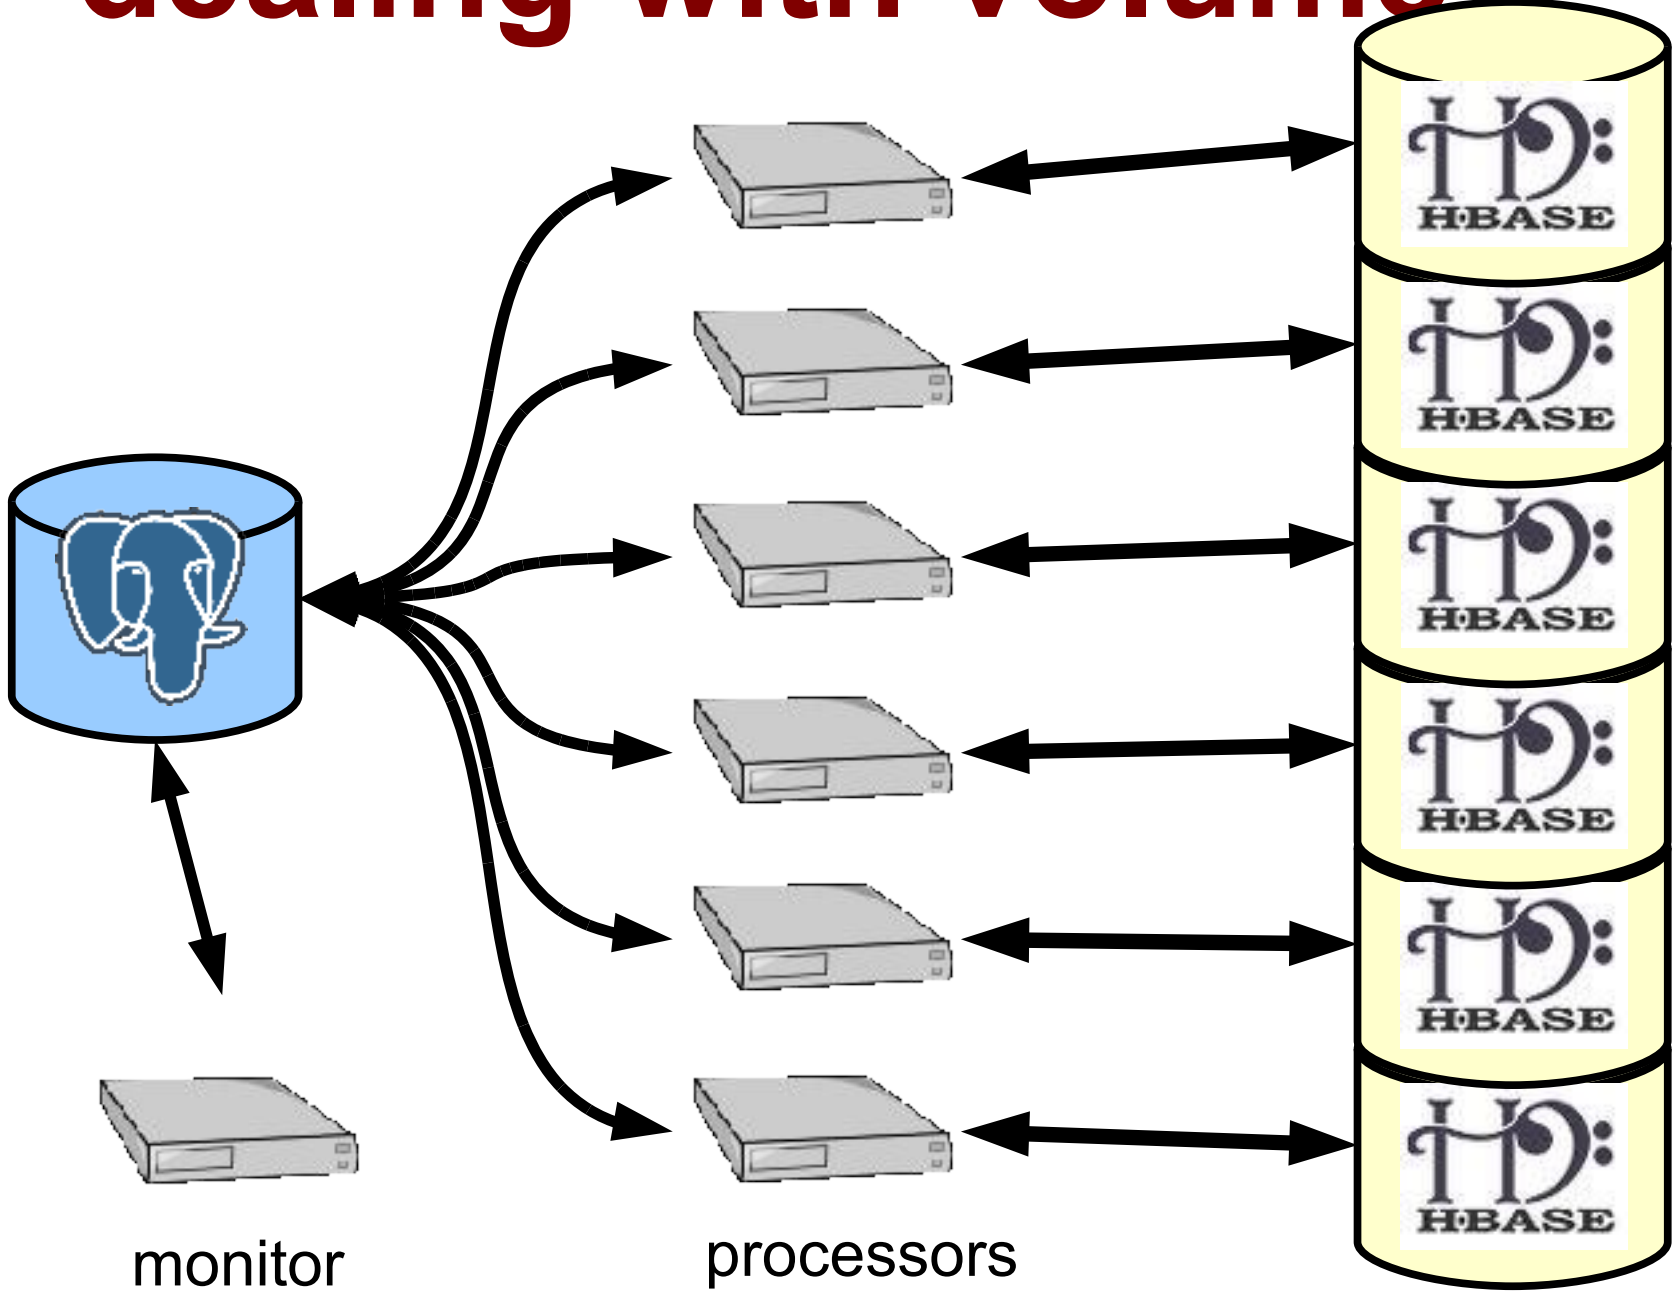
- 3000 crashes/minute
  - avg. size 150K
- 40TB accumulated raw data
  - 500GB accumulated metadata / reports



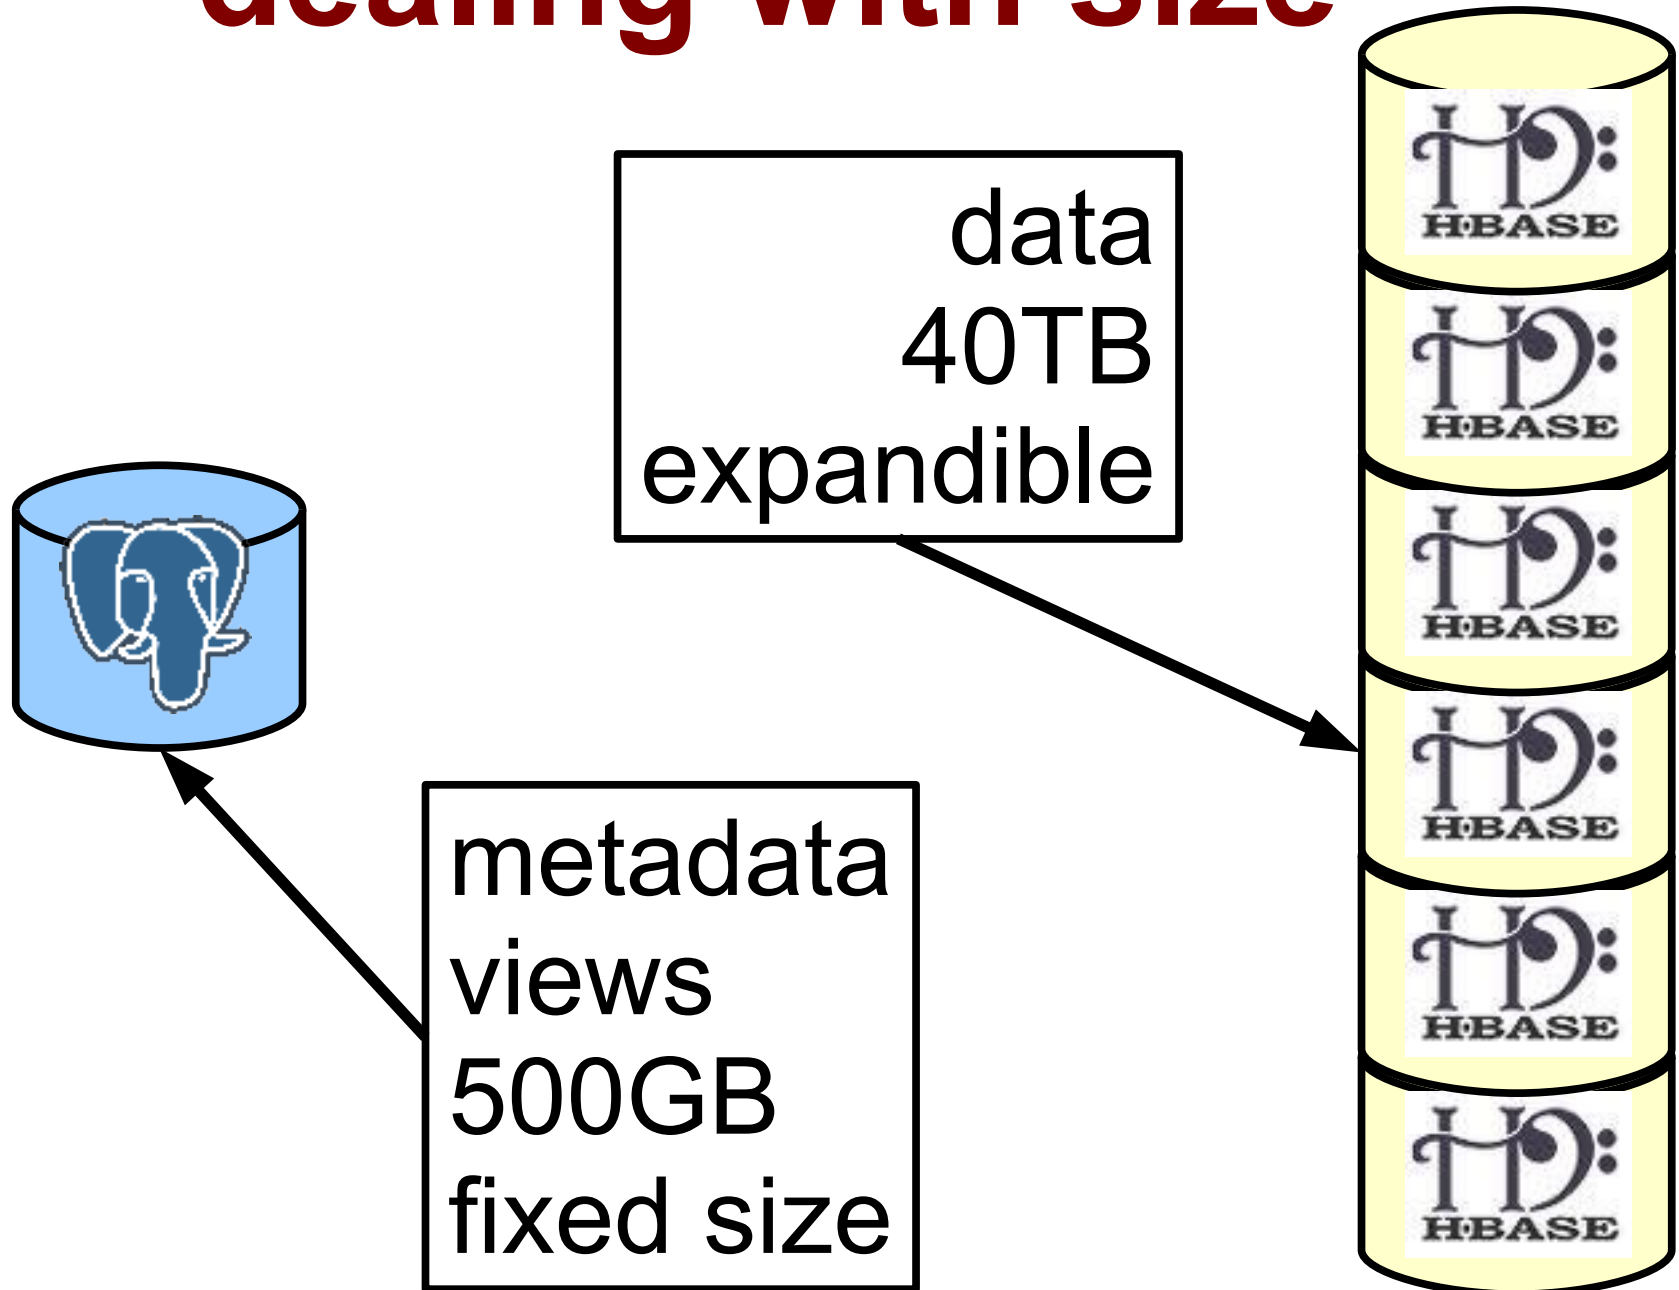
# dealing with volume



# dealing with volume



# dealing with size



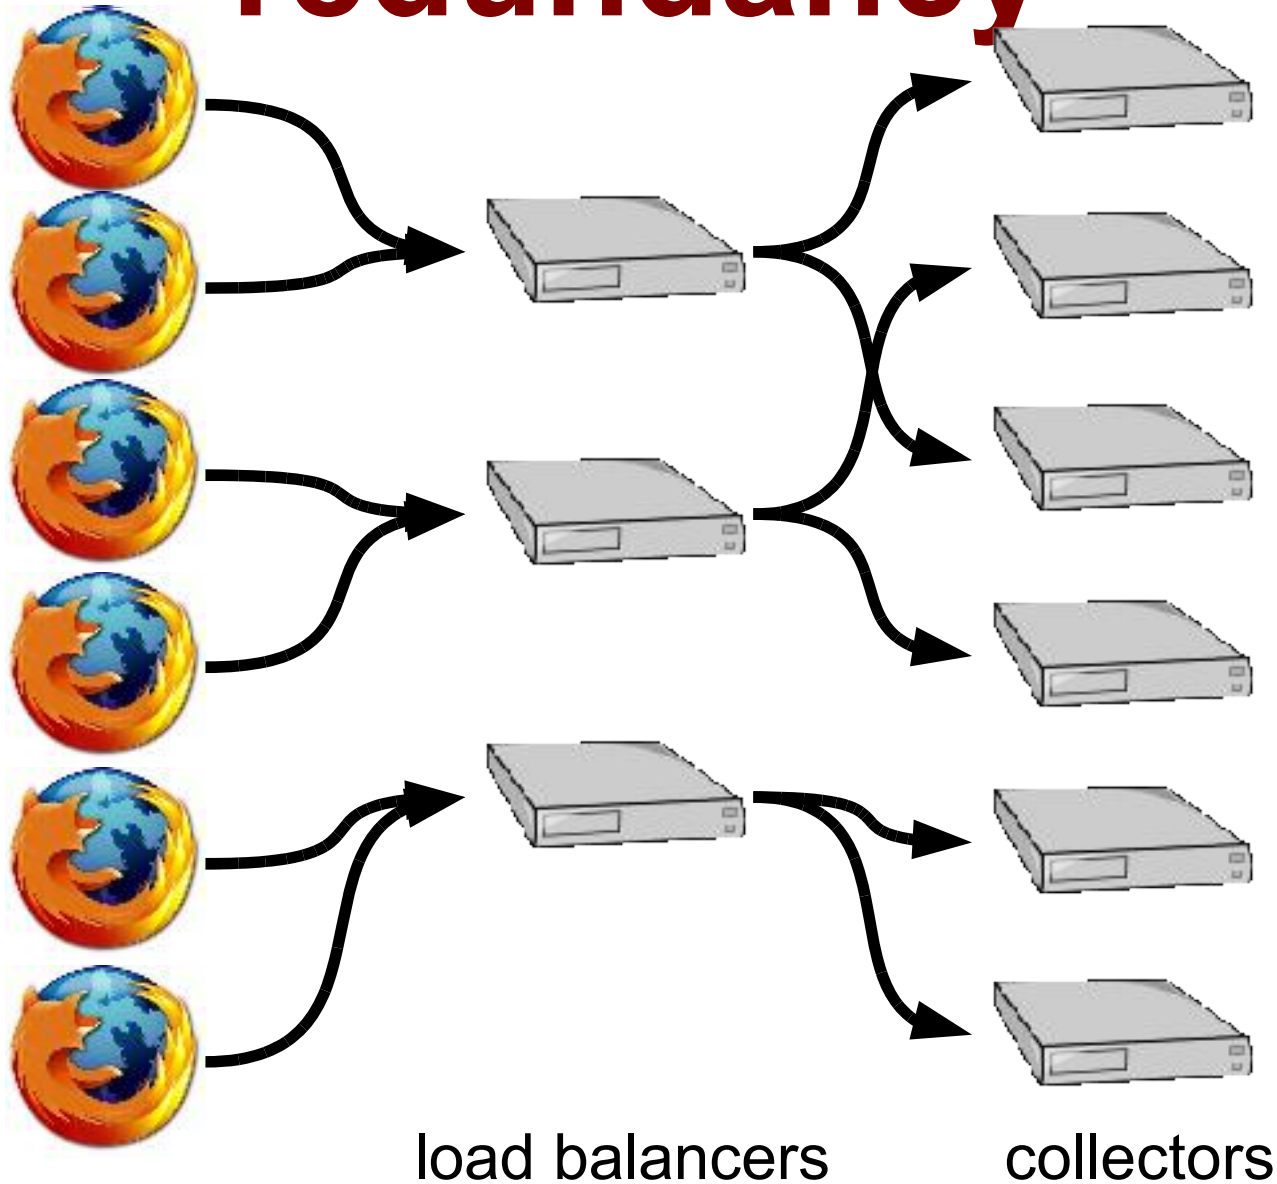
# dealing with component failure

Lots of hardware ...

- 30 Hbase nodes
- 2 PostgreSQL servers
- 6 load balancers
- 3 ES servers
- 6 collectors
- 12 processors
- 8 middleware & web servers

... lots of failures

# load balancing & redundancy

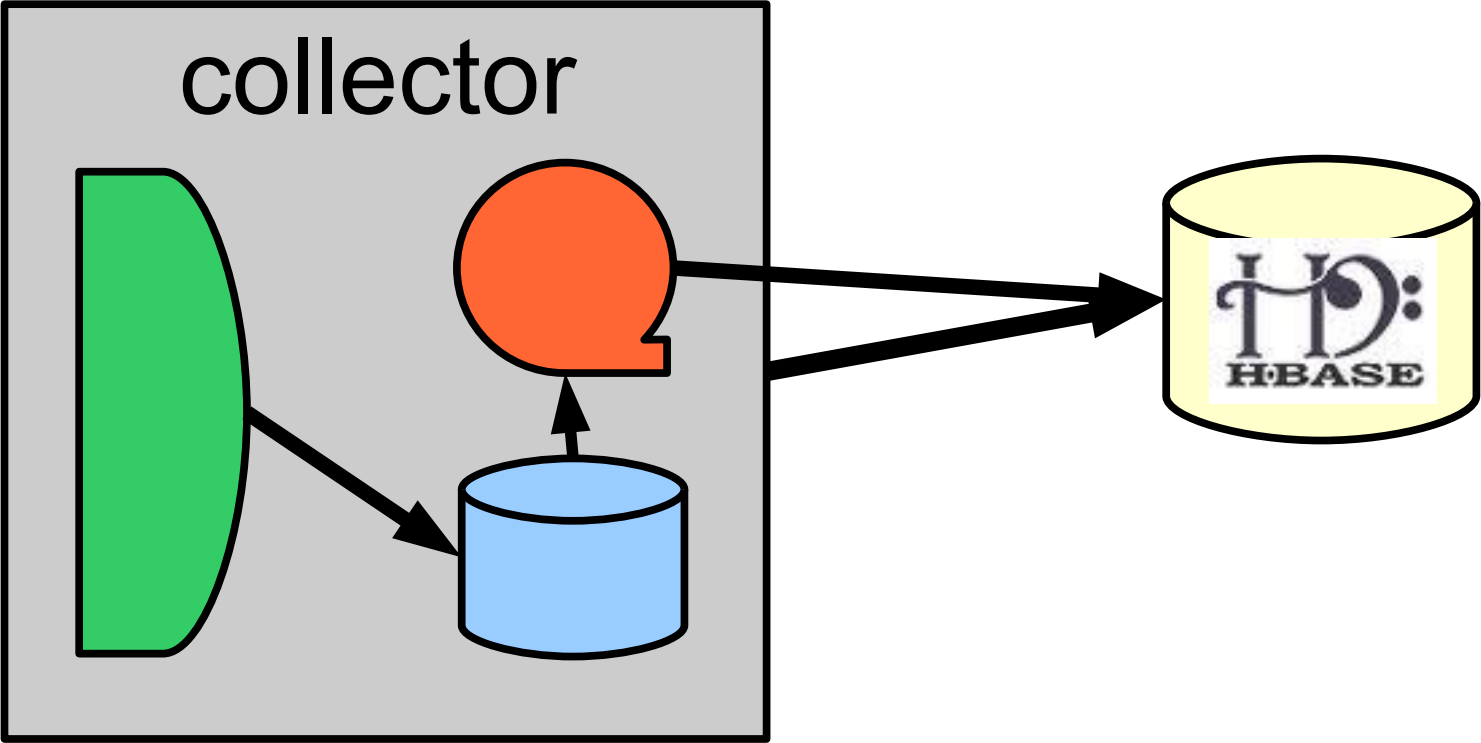


# elastic connections

- components queue their data
  - retain it if other nodes are down
- components resume work automatically
  - when other nodes come back up

# elastic connections

crash mover



reciever


local file queue

# server management

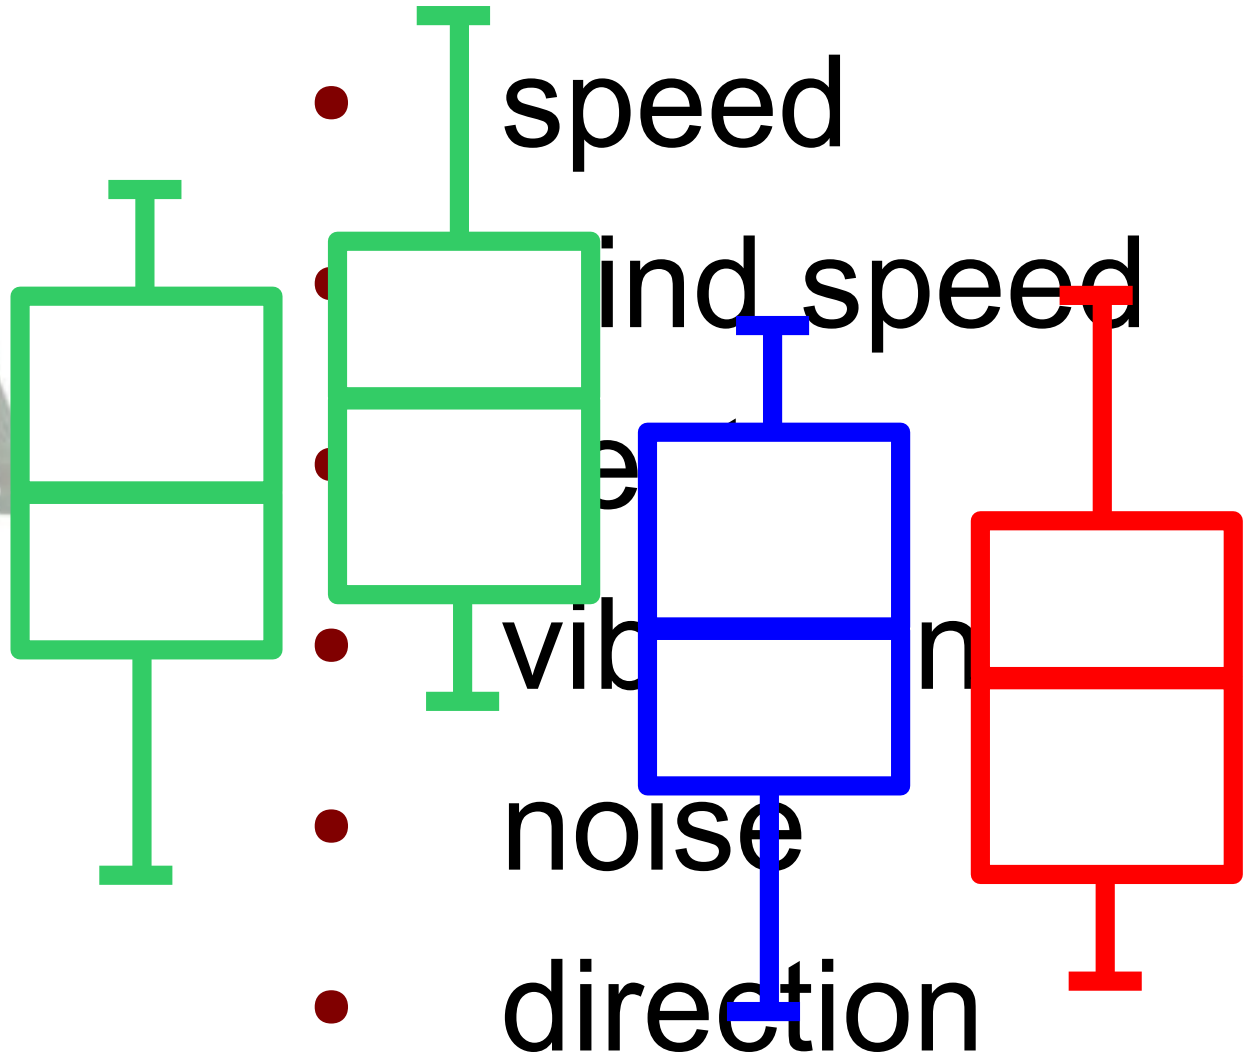
- puppet
  - controls configuration of all servers
  - makes sure servers recover
  - allows rapid deployment of replacement nodes



# Upwind



# Upwind



# Upwind



1. maximize power generation
2. make sure turbine isn't damaged

# dealing with volume

each turbine:

90 to 700 facts/second



windmills per farm: up to 100

number of farms: 40+

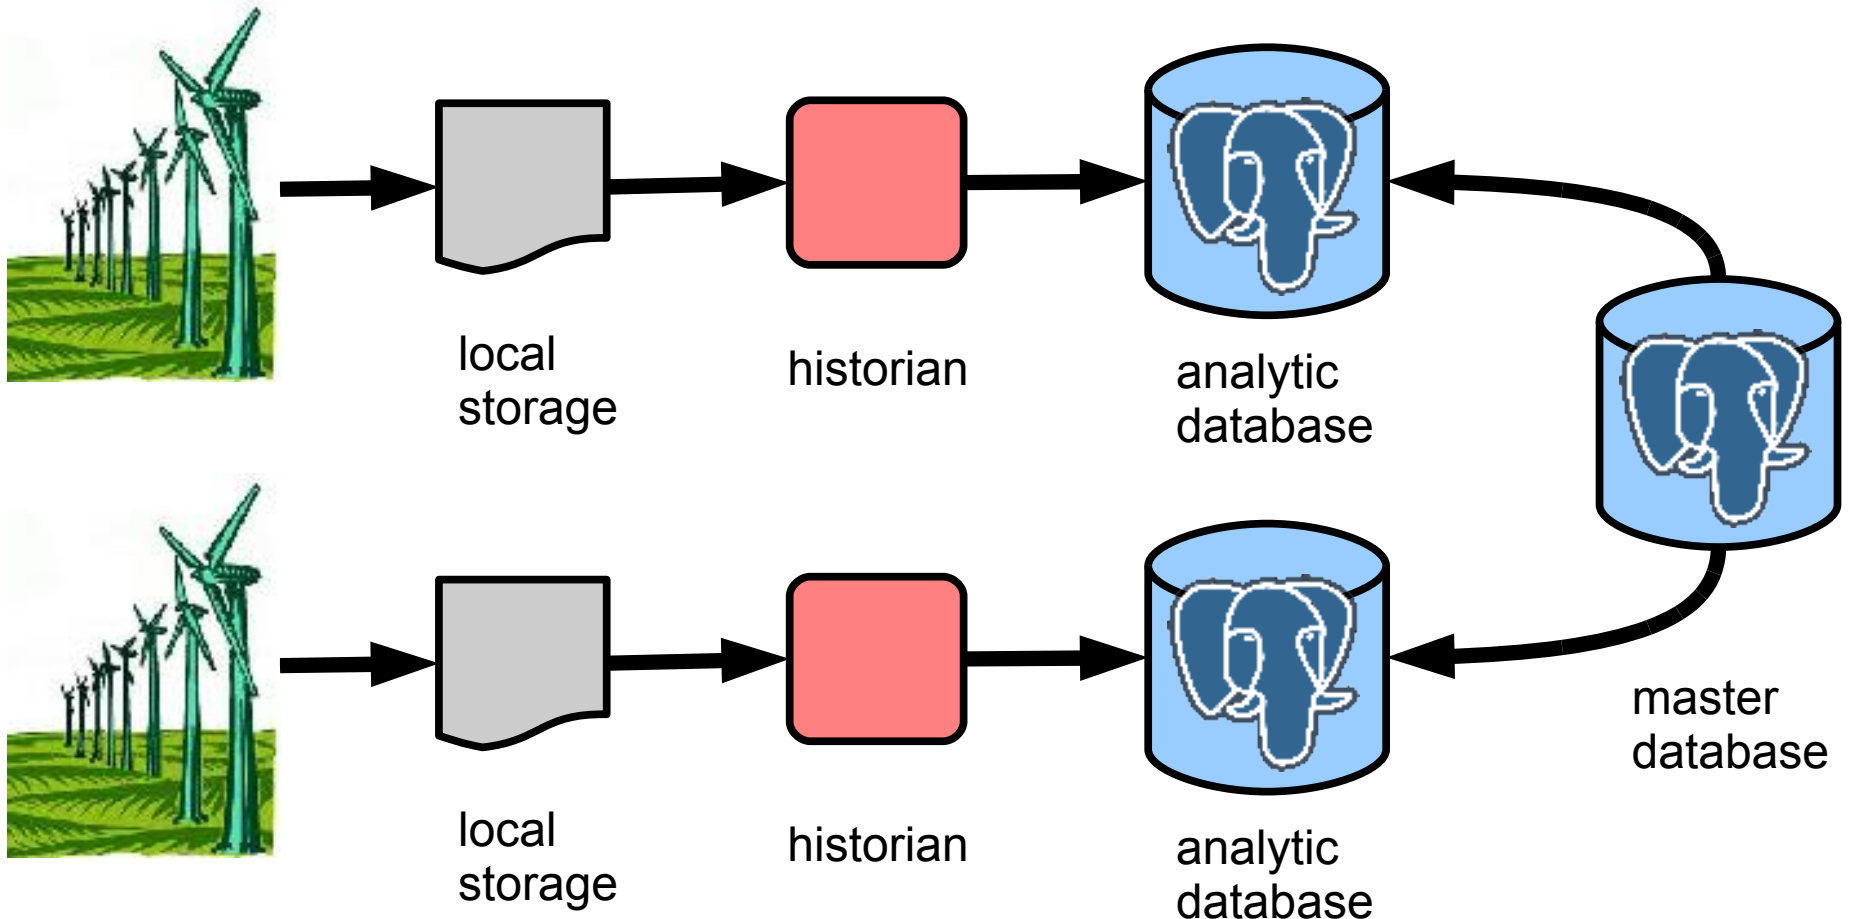
est. total: 300,000 facts/second

*(will grow)*

# dealing with volume



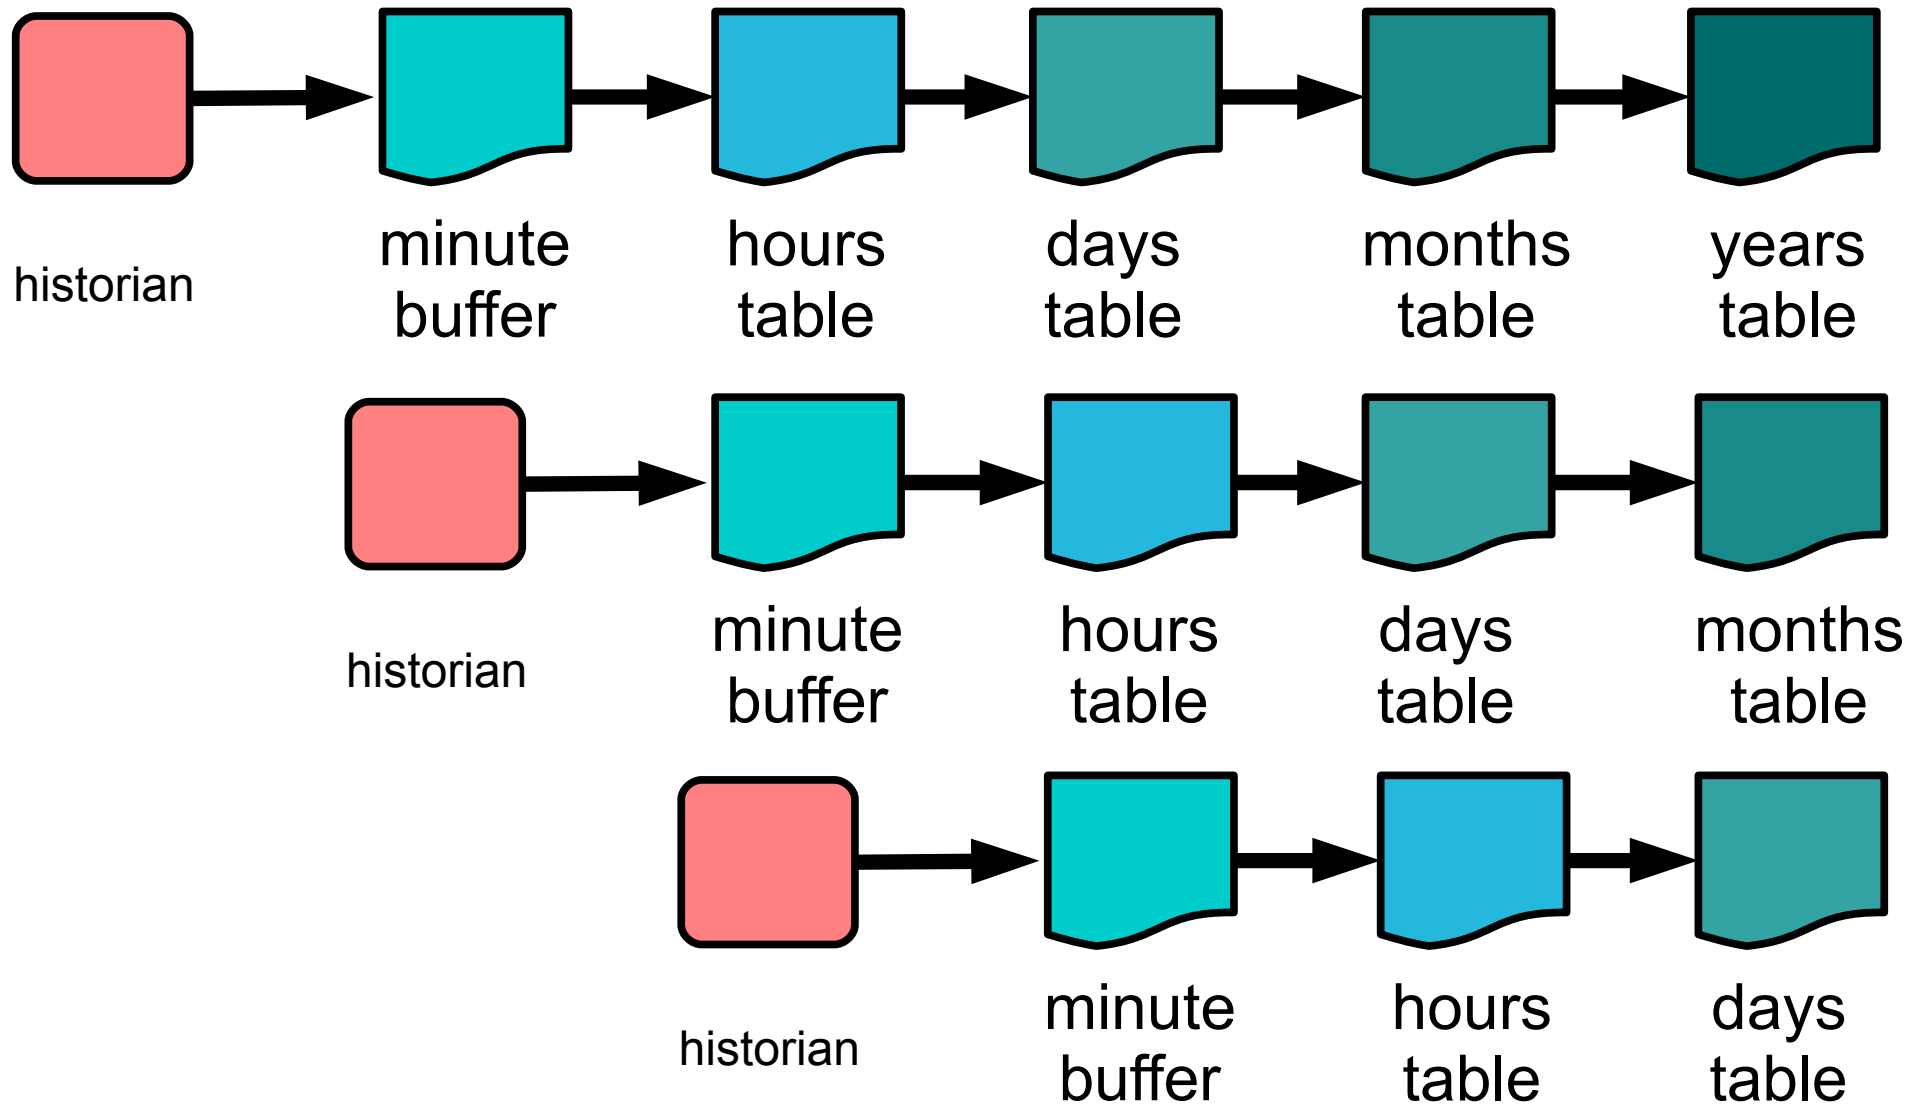
# dealing with volume



# multi-tenant partitioning

- partition the whole application
  - each customer gets their own toolchain
- allows scaling with the number of customers
  - lowers efficiency
  - more efficient with virtualization

# dealing with: constant flow and size





# time-based rollups

- continuously accumulate levels of rollup
  - each is based on the level below it
  - data is always appended, never updated
  - small windows == small resources

# time-based rollups

- allows:
  - very rapid summary reports for different windows
  - retaining different summaries for different levels of time
  - batch/out-of-order processing
  - summarization in parallel



**firehose tips**

# **data collection must be:**

- continuous
- parallel
- fault-tolerant

# data processing must be:

- continuous
- parallel
- fault-tolerant

# **every component must be able to fail**

- including the network
- without too much data loss
- other components must continue

# 5 tools to use

1. queueing software
2. buffering techniques
3. materialized views
4. configuration management
5. comprehensive monitoring

# 4 don'ts

1. use cutting-edge technology
2. use untested hardware
3. run components to capacity
4. do hot patching



firehose mastered?



# Contact

- Josh Berkus: [josh@pgexperts.com](mailto:josh@pgexperts.com)
- blog: [blogs.ittoolbox.com/database/soup](http://blogs.ittoolbox.com/database/soup)
- PostgreSQL: [www.postgresql.org](http://www.postgresql.org)
- pgexperts: [www.pgexperts.com](http://www.pgexperts.com)
- Upcoming Events
  - PostgreSQL Europe: <http://2011.pgconf.eu/>
  - PostgreSQL Italy: <http://2011.pgday.it/>



The text and diagrams in this talk is copyright 2011 Josh Berkus and is licensed under the creative commons attribution license. Title slide image is licensed from iStockPhoto and may not be reproduced or redistributed. Socorro images are copyright 2011 Mozilla Inc.

