Killing Bottlenecks

paul reinheimer

@preinheimer

Goals for Today

- Introduce Tools
- Help you recognize problems
- Hints on implementation

Horror Stories

Gearman

- Provides a simple queue service
- Workers connect and identify work they can complete, clients hand off jobs when available
- Created by Danga Interactive (livejournal)

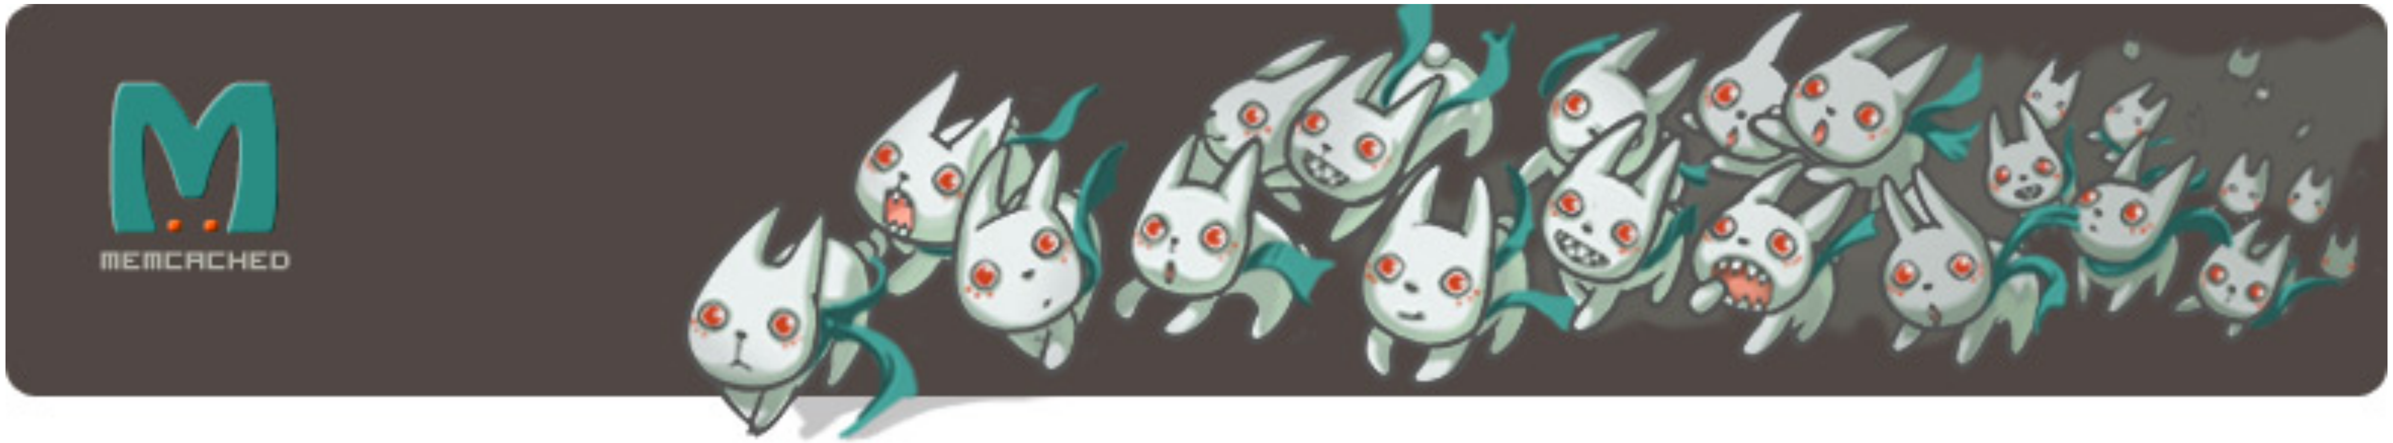

- Stores data in memory
  - really really fast!
- Many servers can be used at once to provide a single large datastore
- Created by Danga Interactive (livejournal)

Sphinx

- Excellent search tool
- Import data from a variety of sources (commonly a database), index it quickly, then run searches

Know Thyself

Gearman

- Workers sign up with work they can do, many workers can each sign up with a different set of jobs.
- Clients connect to Gearman and pass off jobs, this is done asynchronously, they get to keep working
- Failed jobs are re-submitted a specified number of times.

Gearman

- Workers sign up with work they can do, many workers can each sign up with a different set of jobs.
- Clients connect to Gearman and pass off jobs, this is done asynchronously, they get to keep working
- Failed jobs are re-submitted a specified number of times.

Clients

Workers

Managing Workers

Managing Workers

Managing Workers

Managing Workers

Managing Workers

Supervisor

Gearman Alternatives

- RabitMQ
- Amazon Simple Queue Service
- ZeroMQ

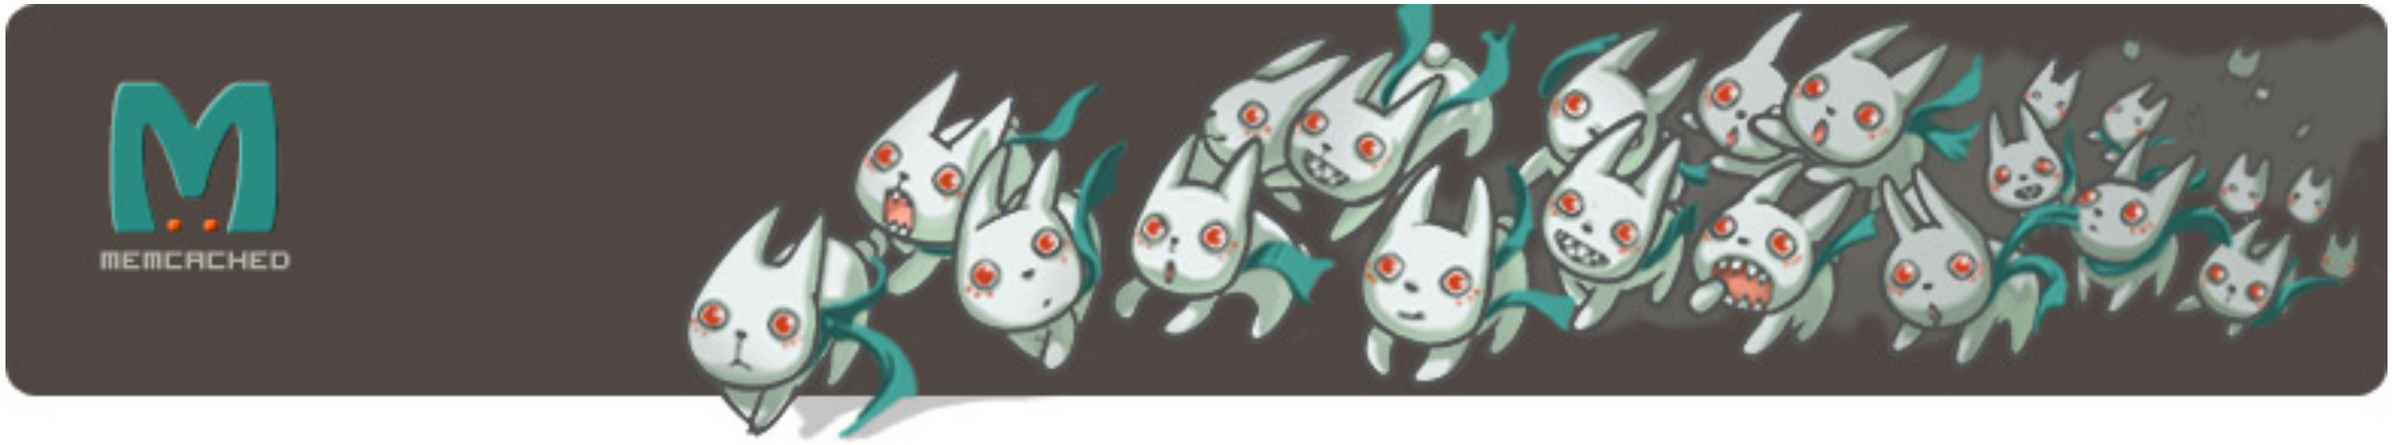

- Stores data in memory
  - really really fast!
- Many servers can be used at once to provide a single large datastore
- Created by Danga Interactive (livejournal)

Promises

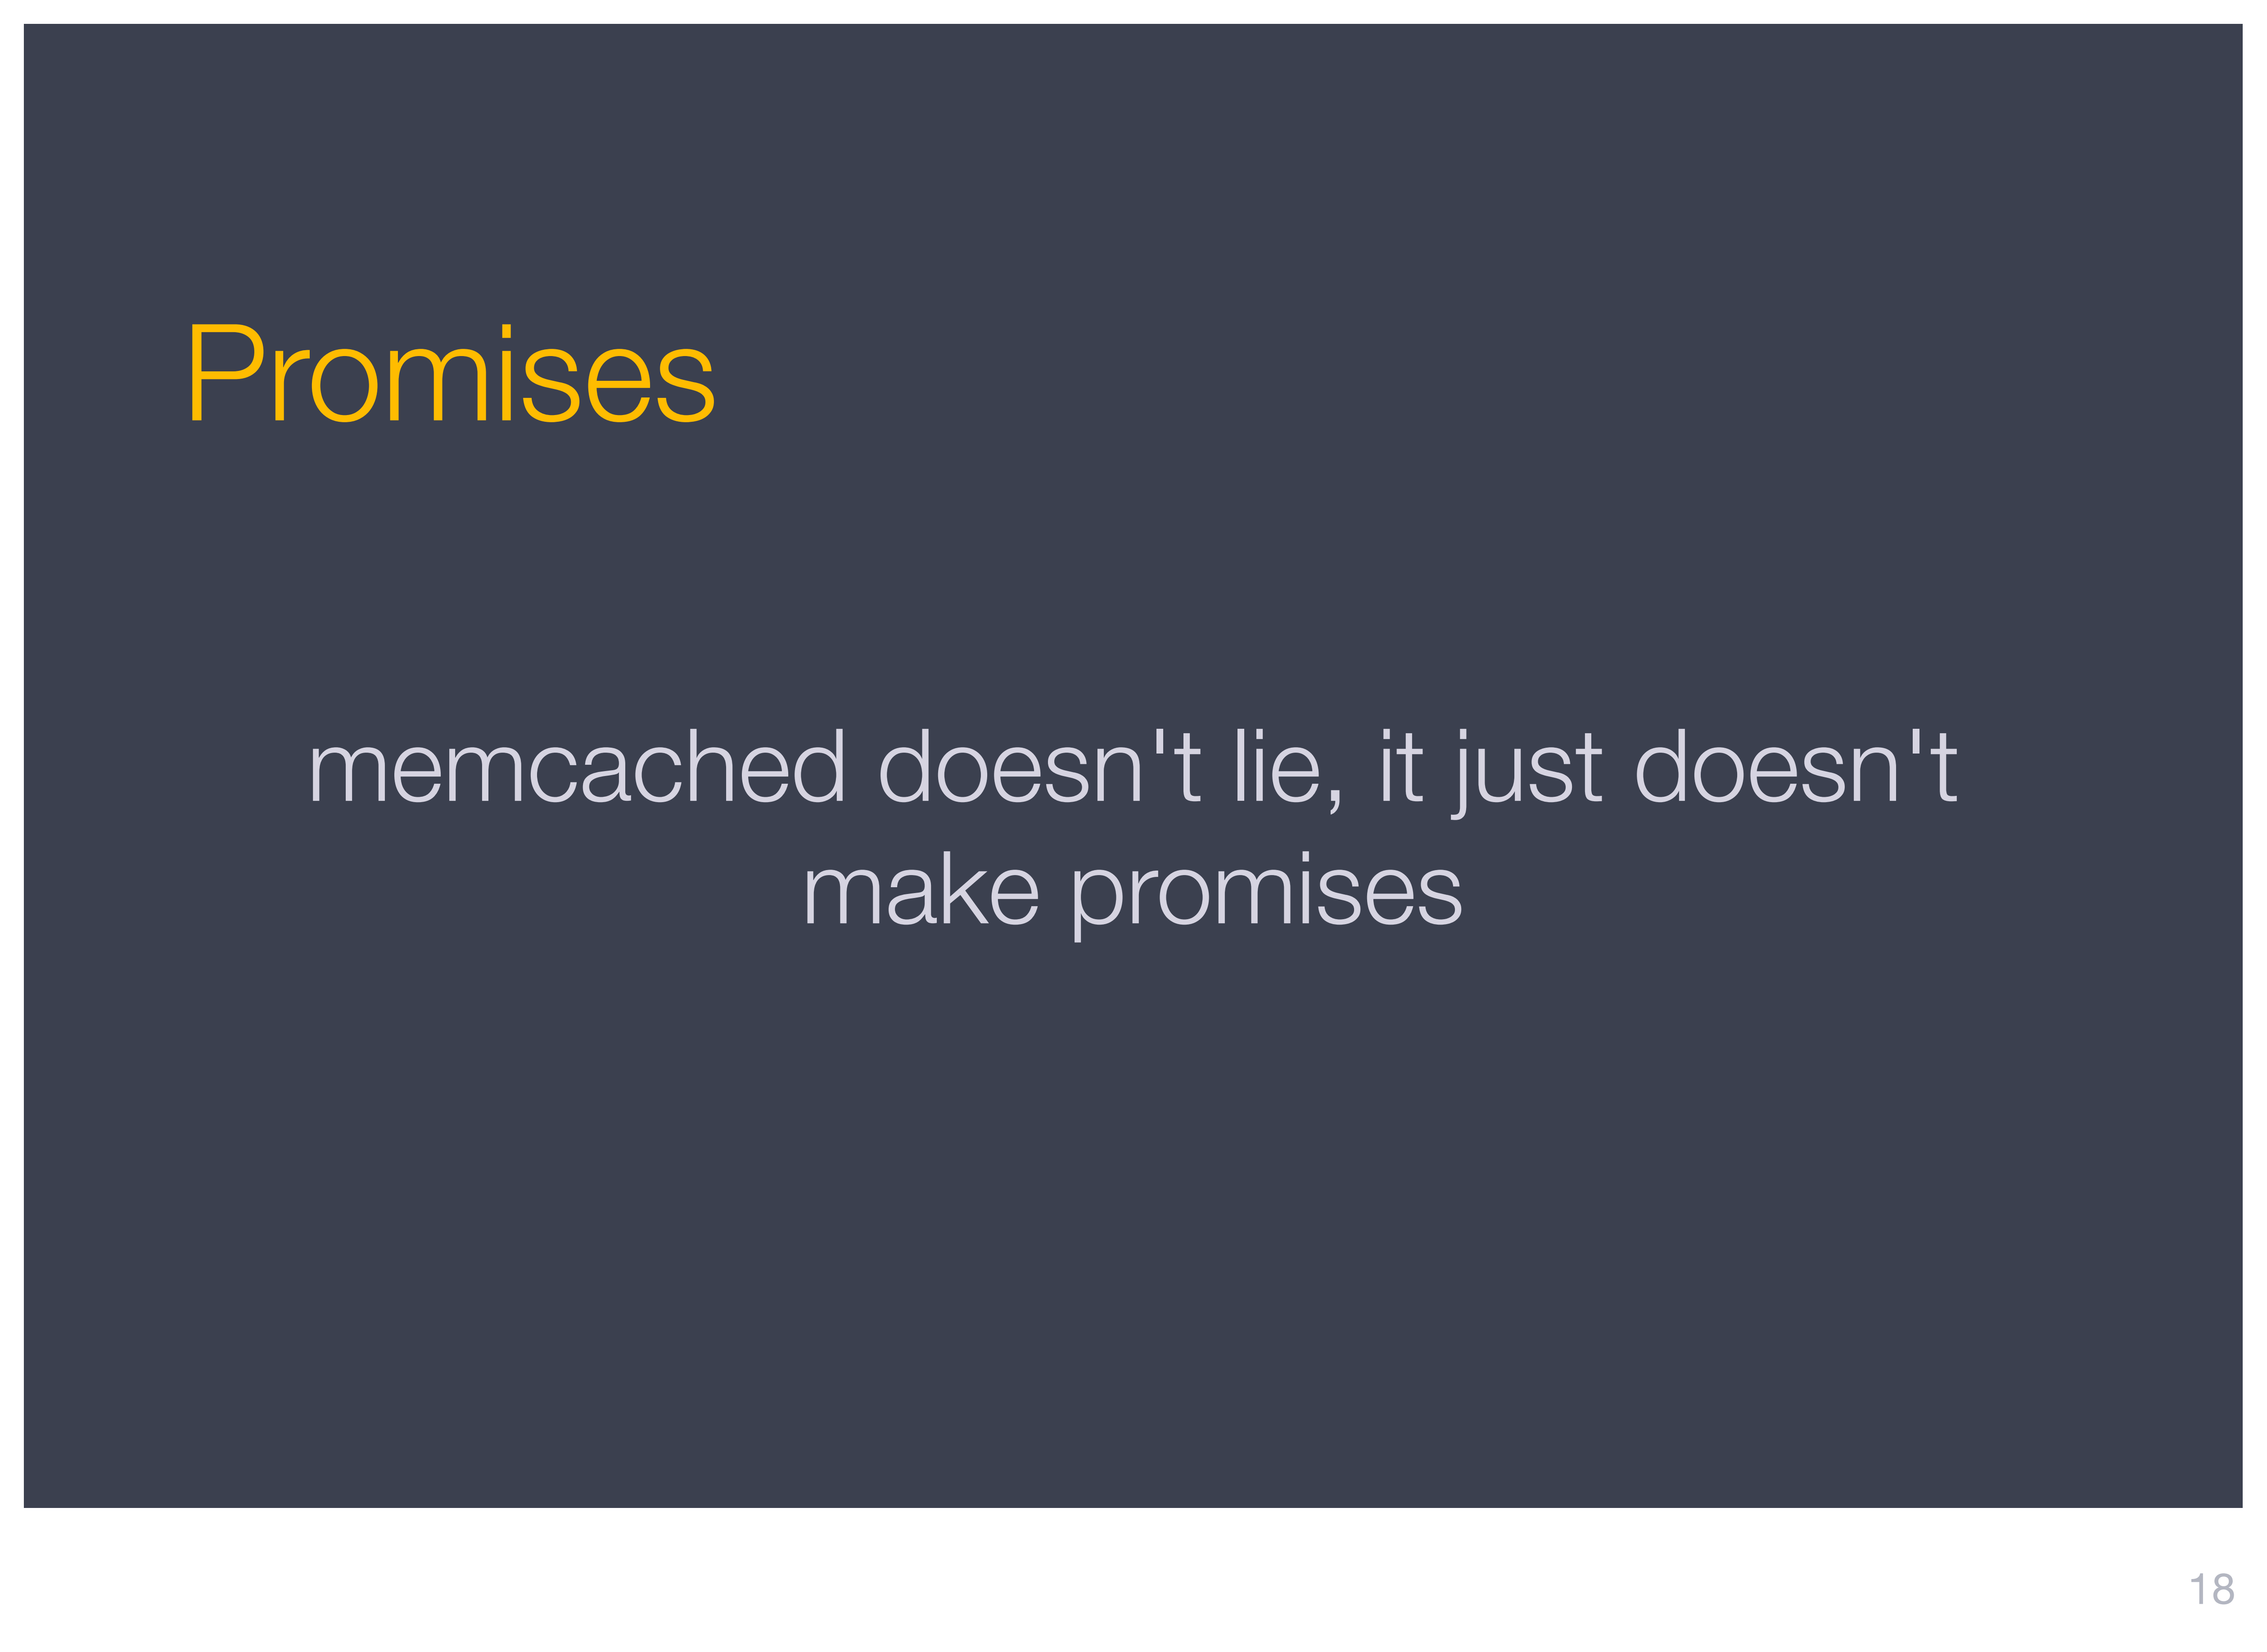

Promises

memcached doesn't lie, it just doesn't make promises

Promises

it makes **a** promise

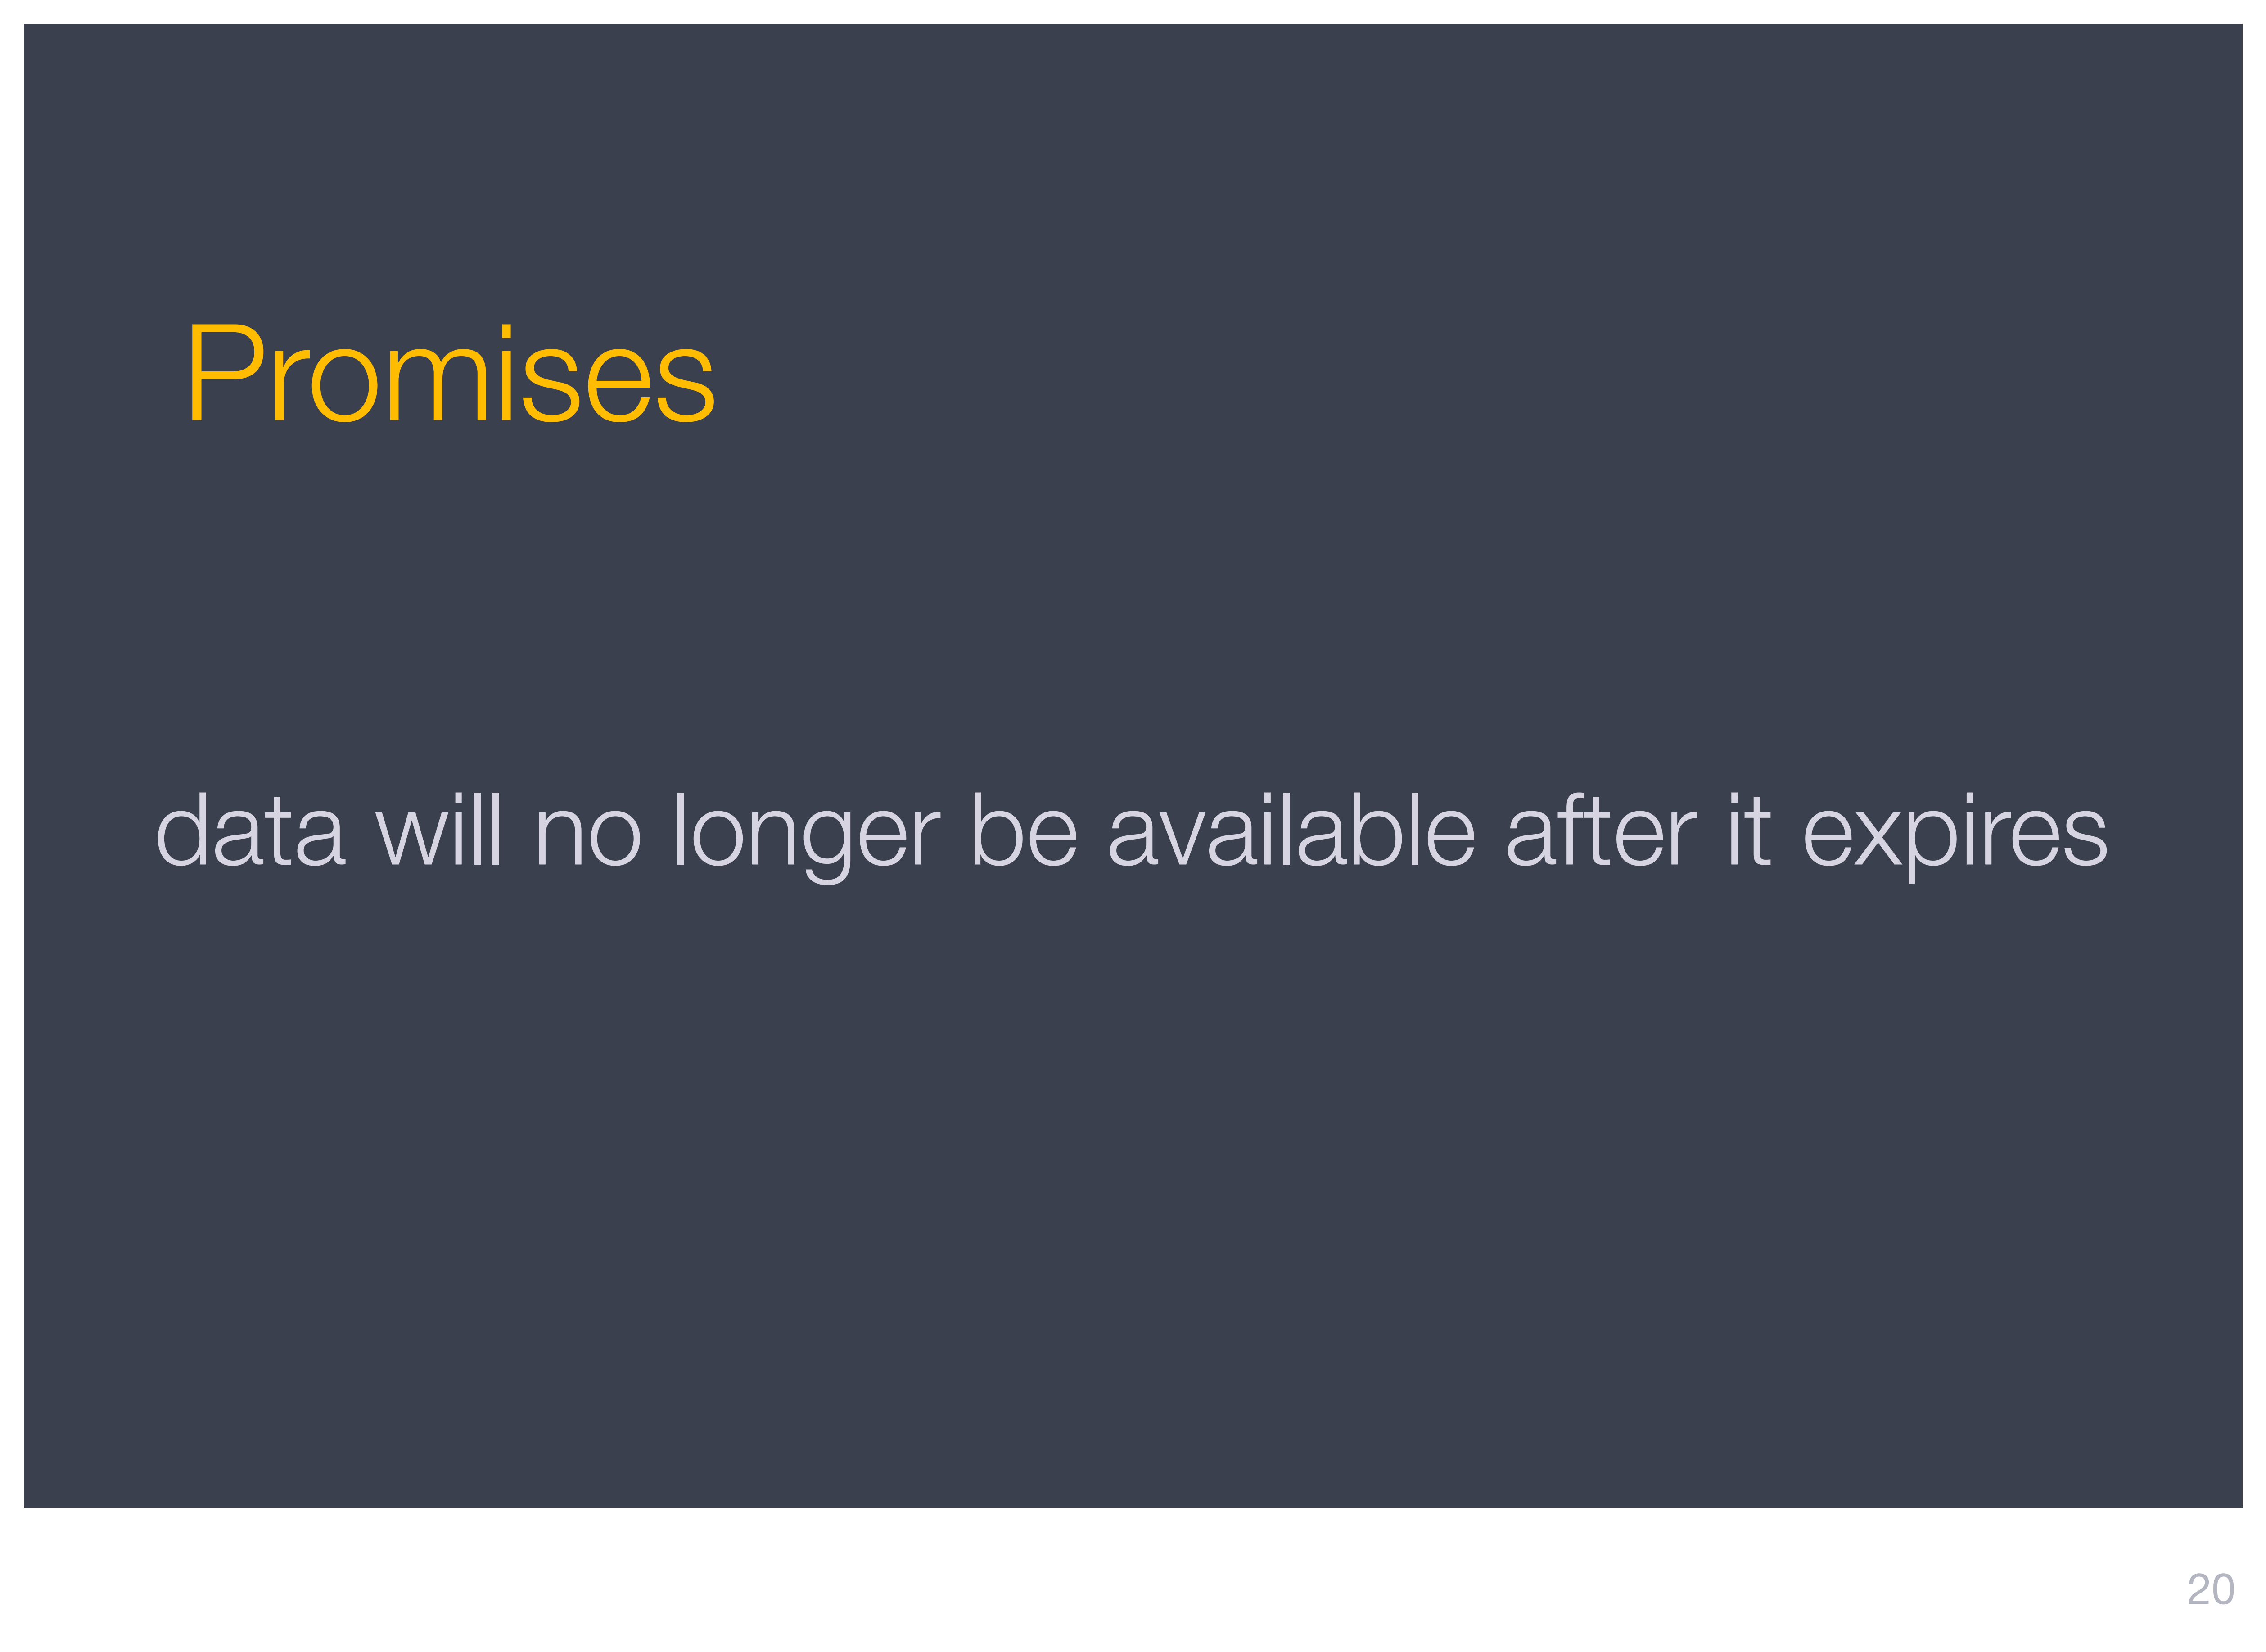

Promises

data will no longer be available after it expires

Promises

- When memcached doesn't have the information you need, get it elsewhere.
  - Use a read-through caching callback or put it in your caching object.
  - memcached is allowed to lose data!

Consistent Hashing

Hot Items = Hot Servers

Hot Items = Hot Servers

Memcached Alternatives

- redis
- Berkeley DB
- Tokyo Cabinet
- Tokyo Tyrant
- Voldemort
- Riak

Sphinx

- Excellent search tool
- Import data from a variety of sources (commonly a database), index it quickly, then run searches

id	title	year	rating
1	TMNT	1988	9.5
2	Goonies	1989	9.0
3	The Postman	1992	3.5
4	Alien	1994	8.5

id	title	year	rating	studio	category
1	TMNT	1988	9.5	x	action
2	Goonies	1989	9.0	y	adventure
3	The Postman	1992	3.5	z	boring
4	Alien	1994	8.5	a	awesome

id	title	year	rating	studio	category
1	TMNT	1988	9.5	x	action
2	Goonies	1989	9.0	y	adventure
3	The Postman	1992	3.5	z	boring
4	Alien	1994	8.5	a	awesome

film_id	tag_id
1	2
2	3
3	4
4	5

id	director
1	blue
2	red
3	green
4	orange

film_id	description
1	Four awesome turtles...
2	Friends, pirates, goonies.
3	Nap
4	They tried to turn around

id	title	year	rating	studio	category
1	TMNT	1988	9.5	x	action
2	Goonies	1989	9.0	y	adventure
3	The Postman	1992	3.5	z	boring
4	Alien	1994	8.5	a	awesome

id	actor
1	Mel Gibson
2	Shredder
3	Sigourney Weaver

id	tag
1	kids
2	awesome
3	sci-fi

id	producer
1	1
2	2
3	3
4	4

film_id	tag_id
1	2
2	3
3	4
4	5

id	director
1	blue
2	red
3	green
4	orange

film_id	description
1	Four awesome turtles...
2	Friends, pirates, goonies.
3	Nap
4	They tried to turn around

id	title	year	rating	studio	category
1	TMNT	1988	9.5	x	action
2	Goonies	1989	9.0	y	adventure
3	The Postman	1992	3.5	z	boring
4	Alien	1994	8.5	a	awesome

id	actor
1	Mel Gibson
2	Shredder
3	Sigourney Weaver

id	tag
1	kids
2	awesome
3	sci-fi

id	producer
1	1
2	2
3	3
4	4

film_id	tag_id
1	
2	
3	
4	

film_id	tag_id
1	1
2	2
3	3
4	4

id	director
1	
2	
3	
4	

film_id	director_id
1	1
2	2
3	3
4	4

d	description

film_id	desc_id
1	1
2	2
3	3
4	4

id	title	year	rating	studio	category
1	TMNT	1988	9.5	x	action
2	Goonies	1989	9.0	y	adventure
3	The Postman	1992	3.5	z	boring
4	Alien	1994	8.5	a	awesome

id	actor
1	
2	
3	

film_id	actor_id
1	1
2	2
3	3
4	4

id	tag
1	
2	

film_id	tag_id
1	1
2	2
3	3
4	4

id	producer
1	
2	
3	
4	

film_id	producer_id
1	1
2	2
3	3
4	4

```
SELECT
  `dvd_scenes`.`scene_id`,
  `dvds`.`dvd_id`,
  `dvds`.`name`,
  `dvds`.`name`                    AS sort_title,
  `dvds`.`featured`                AS premium,
  `dvds`.`description`,
  `dvds`.`date_released`,
  `dvds`.`channel_id`,
  `dvds`.`studio_id`,
  `dvd_scenes_dyn`.`times_viewed` AS sub_order,
  `dvd_scenes_dyn`.`rating`       AS rating,
  `channels`.`name`                AS channel_name,
  `studios`.`name`                 AS studio_name,
  `channels_keywords`.`keywords`   AS channel_keywords,
  GROUP_CONCAT(DISTINCT `categories`.`name` ORDER BY `categories`.`name` DESC SEPARATOR ', ') AS categories_list,
  GROUP_CONCAT(DISTINCT `models`.`name` ORDER BY `models`.`name` DESC SEPARATOR ', ') AS models_list
FROM `dvd_scenes`
  INNER JOIN dvds
    ON dvds.dvd_id = dvd_scenes.dvd_id
  INNER JOIN dvd_scenes_dyn
    ON dvd_scenes_dyn.id = dvd_scenes.scene_id
  LEFT JOIN scenes_categories
    ON dvd_scenes.scene_id = scenes_categories.scene_id
  LEFT JOIN categories
    ON `scenes_categories`.`category_id` = categories.category_id
  LEFT JOIN studios
    ON studios.studio_id = dvds.studio_id
  LEFT JOIN scenes_models
    ON scenes_models.scene_id = dvd_scenes.scene_id
  LEFT JOIN `channels`
    ON channels.channel_id = dvds.channel_id
  LEFT JOIN `channels_keywords`
    ON dvds.channel_id = channels_keywords.channel_id
  LEFT JOIN `models`
    ON scenes_models.model_id = models.model_id
WHERE dvds.is_active = 1
    AND dvd_scenes.is_active = 1
    AND dvds.date_released < UNIX_TIMESTAMP()
    AND dvds.date_released > 0
    AND dvds.channel_id NOT IN (21,26)
GROUP BY scene_id
```

sphinx

http://www.flickr.com/photos/waldenpond/4237758518/

Sphinx Alternatives

- Lucene
  - With Elastic Search
- solr

Assigning Problems

Appendix

- Paul Reinheimer - @preinheimer - http://blog.preinheimer.com
- Memcached - http://memcached.org/
  - supervisor - http://supervisord.org/
- Sphinx - http://sphinxsearch.com/
- Gearman - http://gearman.org/

5 Stars!

Feedback

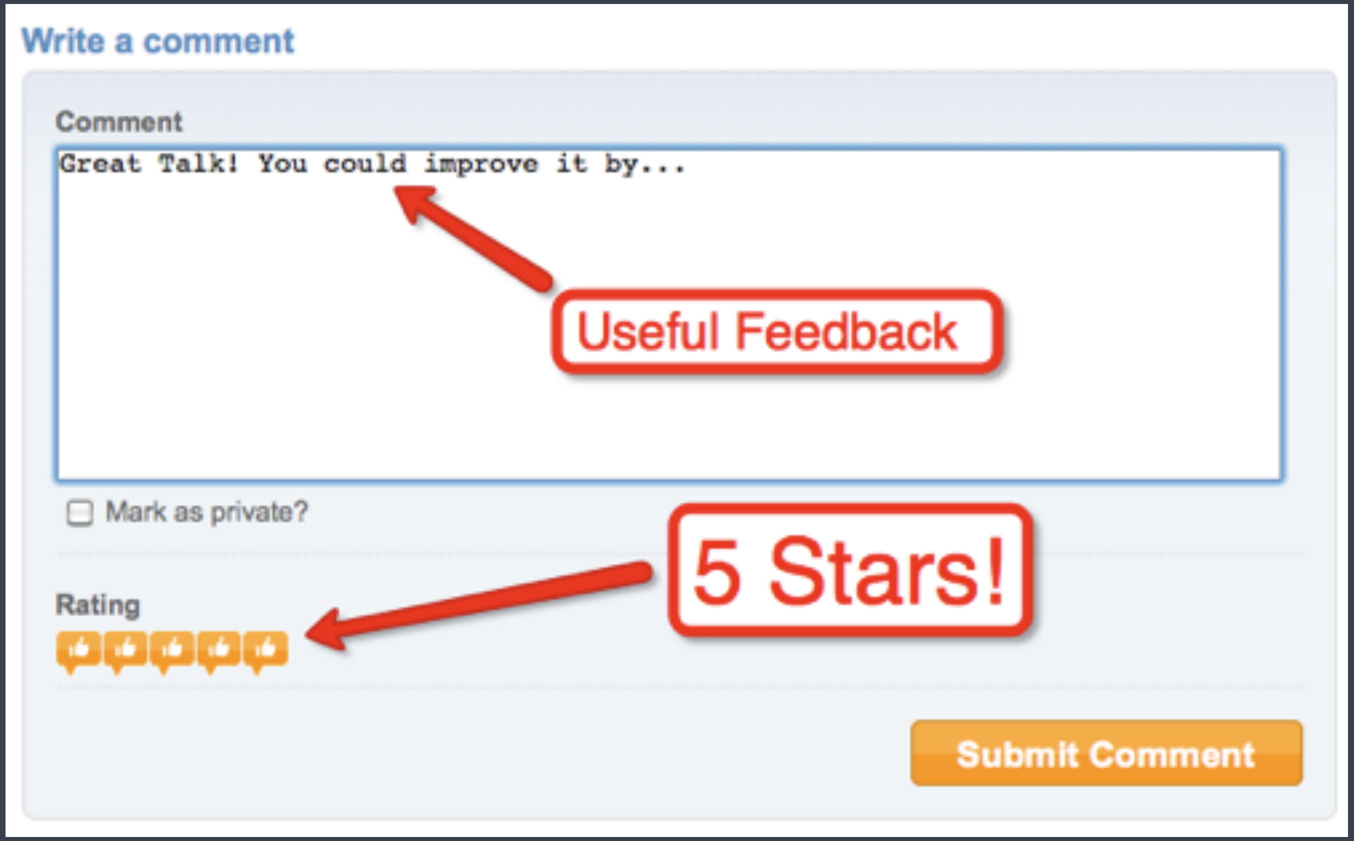

Write a comment

Comment

Great Talk! You could improve it by...

http://joind.in/2867