

# Building Killer Web Applications

Dave King

December 6th, 2011

## Once Upon a Time

- “I need a program that does...”

# Conditional Probability Calculator

Fork me on GitHub

Previous

Next

## Configure Model

Model Name:

Number of Variables:

## Variables

Name:

Range: 1 to

# My 'Education'

- Lots of 'server-side' work
- TAs compiled my code, seeing if it worked against a text file
- Native application development
  - Windows Sound API
  - Directing UNIX Process Output to Files
  - Compilers and Type-Checkers

# The New World

- Jeff Atwood (Coding Horror) in 2009: “All programming is web programming.”
- As a developer, the web is now the delivery mechanism for your code
- How are web applications built and improved?

# We Use These Every Day

- Web applications...
  - Web pages
  - Dynamic behavior
  - Persist information between visits

## Stop Dave, I'm Afraid

**RustyCT1** 17 videos



**466,423**

Uploaded by [RustyCT1](#) on Mar 4, 2008

Excerpt from Stanley Kubrick's phenomenal 1968 film featuring Arthur C. Clark's bestseller, "2001: A Space Odyssey," starring Keir Dullea.

1,233 likes, 24 dislikes

As Seen On:

[US Politics](#) | [AMERICA...](#)

Top Comments



**Hal 9000 sings Daisy**

by drinkmokolony  
84,160 views



**I'm sorry, Dave**

by jdoniach  
201,527 views



**Best Death Scenes: HAL 9000**

by BestDeathScenes  
195,596 views



**Hal wont Open the Pod Bay Door**

by blondegal1957  
131,055 views



**HAL takes out Frank**

by Excaveneau  
92,977 views



**Lustige Leute =)**

by happyFee1993  
[PLAYLIST](#) 50 videos



**Daisy...Daisy...**

by splosion7  
24,777 views

# Delicious

**delicious** featured recent search

hi, tildedave [+ link](#) [+ stack](#)

**my links** my stacks my inbox

20 links newest | oldest

tildedave tag bundle: **api** x - show all tags in bundle

select all create stack add to stack make public make private manage tags delete

29 dec 2010 50 saves  
**Testing REST APIs with Cucumber and Rack::Test « All Things (v8)**  
<http://www.anthonieden.com/2010/11/testing-rest-apis-with-cu...>  
ruby rails rest api

21 oct 2010 1953 saves  
**Twisted**  
<http://twistedmatrix.com/trac/>  
python concurrency api

19 oct 2010 55 saves  
**RESTx - The simplest way to create RESTful web services | RESTx**  
<http://restx.mulesoft.org/>  
api rest

18 oct 2010 56 saves  
**rest-api-code-gen - Project Hosting on Google Code**  
<http://code.google.com/p/rest-api-code-gen/>  
wadi rest api

18 oct 2010 149 saves

**tildedave**  
joined 06 mar 2007  
following 0 users  
followed by 0 users  
20 links  
rss: [public](#) | [private](#)  
**tag bundles** manage  
api  
exercise  
literature  
music  
programminglanguages

# Rackspace Webmail

The screenshot displays the Rackspace Webmail interface. At the top left is the Rackspace Hosting logo. The navigation bar includes 'Email', 'Contacts', 'Calendar', 'Tasks', and 'Notes'. On the right, there are links for 'Settings', 'Support', 'Feedback', and 'Log Out', along with a search box for 'OpenStack'. Below the navigation bar is a toolbar with actions: Compose, Check Email, Reply, Reply All, Forward, Delete, Move, and More. The left sidebar shows a tree of 'Email Folders' including Inbox, Drafts, Blacksburg (with sub-folders like Ackers, R&D, REM, Unit5), Builds (with sub-folders like Bamboo, Jenkins, Career), Dev (with sub-folders like Apps CP, Changelog, Cloud CP, devtests, Frontend-Eng, Reach), and Mailing Lists (with sub-folder Asrg). A storage indicator at the bottom left shows 1.2 GB used of a 10.0 GB total. The main area contains a table of emails with columns for From, Subject, Received, and Size. The email from Stefano Maffulli is selected. Below the table, a message pane displays 'No message selected'. At the bottom right, there are navigation icons for the interface.

**rackSPACE**  
HOSTING

Email | Contacts | Calendar | Tasks | Notes

Settings | Support | Feedback | Log Out

Search in OpenStack Search

Compose | Check Email | Reply | Reply All | Forward | Delete | Move | More

From	Subject	Received	Size
Lorin Hochstein	Re: [Op...	Today 7:10 pm	14 KB
Cole	Re: [Op...	Today 5:08 pm	10 KB
Tim Bell	Re: [Op...	Today 2:11 pm	5 KB
Muriel	Re: [Op...	Today 12:19 pm	5.7 KB
Leandro Reox	Re: [Ope...	Fri 12/2 10:04 pm	9.5 KB
Stefano Maffulli	Re: [Ope...	Fri 12/2 5:37 pm	5.7 KB
Stefano Maffulli	[Opensta...	Fri 12/2 5:01 pm	16 KB
Daryl Walleck	Re: [Ope...	Fri 12/2 2:34 pm	7.8 KB
Nachi Ueno	Re: [Ope...	Fri 12/2 2:17 pm	7.6 KB

1 - 100 of 5520

No message selected

1.2 GB 10.0 GB



# Rackspace Cloud Control Panel

The screenshot displays the Rackspace Cloud Control Panel interface. At the top left is the Rackspace Cloud logo. The top right contains support links (Cloud Status, Knowledge Base, Forums, Tickets, Live Chat), contact information (1-877-934-0407, 800-883-3012), and the user's login status (You are logged in as: mailtrustmercury).

The left sidebar contains navigation options: Home, Hosting, Cloud Sites, Cloud Files, Cloud Servers (highlighted), Load Balancers, Clients, Your Account, Support, and Logout.

The main content area is titled "hermes0" and includes tabs for Overview, DNS, Images, and Diagnostics. Under "Actions For This Server", there are buttons for Console, Reboot, Rescue, Rebuild, Resize, Reset Password, and Delete. Below this is the "Cloud Server Details" section, which includes a "Name & Status" table:

Name & Status	
<b>Name:</b>	hermes0 <a href="#">Change Name</a>
<b>Status:</b>	Active
<b>Current Action:</b>	None
<b>Age:</b>	449 Days

At the bottom of the main content area is a "Technical Details" section. The footer of the page contains helpful links and the version information: "Control Panel Version 2.4 - Release Notes".

# How Do You Build That?

- Three Main Parts

# The Client

- Display content with the browser
- Handle presentation and interaction
  - JavaScript and CSS

# The Server

- Handle requests from the user
- Serve content to the browser

# The Data Store

- Persist data
- Usually this is stored in relational databases
  - MySQL, PostgreSQL, Microsoft SQL, Oracle, SQLite
- Sometimes the data is stored beneath a service layer
  - External APIs

## Small is Easy

- Lots of good technology out there that makes building a small application pretty easy
- How do you grow?
- How do you improve the experience?
- What are you likely to see as you go from five users to five thousand?

# Rackspace

- Cloud: 110,000+ customers
- Email & Apps: 2.5 million mailboxes (and growing)



# Making The Client Awesome

- The client interacts with the website through the browser
- Client-side behavior: JavaScript
- Client-side presentation: CSS



# Minimize HTTP Requests

- Hit a web page, start grabbing content from the server
- HTTP requests are slow
- Ideally, browser should make one request per type of content

The screenshot shows the Rackspace Cloud dashboard. At the top, there's a navigation bar with the Rackspace Cloud logo, support links (Cloud Status, Knowledge Base, Forums, Tickets, Live Chat), and contact information (1-877-934-0497, 6090-983-3012). Below the navigation bar, there's a "Welcome" section with "Welcome" and "Getting Started" buttons. The main content area is divided into two columns. The left column is titled "Account Activity" and contains a table with columns "Cloud Sites", "Percent Used", "Used", and "Allocated". The right column is titled "Cloud News" and contains a message about the Knowledge Center. Below the main content area, there's a "Network" section with a table of network activity. The table has columns: Name Path, Method, Status, Type, Initiator, Size Content, Time Latency, and Timeline. The table shows two rows of data for "base.js" and "ga.js".

Support: [Cloud Status](#) [Knowledge Base](#) [Forums](#) [Tickets](#) [Live Chat](#) 1-877-934-0497 6090-983-3012  
You are logged in as: maltrustrmercury

Home  
Hosting  
Clients  
Your Account  
Support  
Logout

### Welcome

Welcome Getting Started

#### Account Activity

Amounts shown here are for the billing cycle of November 11, 2011 through December 10, 2011.

Cloud Sites	Percent Used	Used	Allocated
<a href="#">Bandwidth</a>	No Data To Display	-	500.0 GB
<a href="#">Compute Cycles</a>	No Data To Display	-	10,000.00 cc

#### Cloud News

##### Check out the new Knowledge Center!

We listened to your feedback and revised our knowledge bases. We now have the Knowledge Center that contains articles, videos and white papers to help answer questions or

For helpful apps and tools check out [Cloud Tools](#) and the [Rackspace Cloud iPhone page](#). Control Panel Version 2.4 - [Release Notes](#)


Elements Resources Network Scripts Timeline Profiles Audits Console Search Network

Name Path	Method	Status	Type	Initiator	Size Content	Time Latency	Timeline
base.js js_cscript_974079078/bundle	GET	200 OK	text/jav...	Home de:54	88.15KB 321.52KB	276ms 95ms	
ga.js ssl.google-analytics.com	GET	200 OK	text/jav...	Home de:186	0B 31.85KB	37ms 20ms	

2 / 17 requests | 88.15KB / 156.57KB transferred

# Sprites

- Merge images together into one, access through CSS offsets



# Minimize HTTP Requests

- At Rackspace: Cloud Control Panel
- On application startup: combine CSS/JavaScript/images together

## Combine and Minify CSS/JavaScript

- Strip whitespace and send content gzipped to client
- Less time to get content, less time to start rendering the web page

```
.mozilla input[type="button"]::-moz-focus-inner{border:0px;}.mozilla button{margin-right:6px;}.mozilla button span{margin-left:-3px;margin-right:-3px;margin-top:-1px;padding-left:11px;padding-bottom:1px;line-height:190%;}.mozilla button.primary.disabled span{padding-right:10px;padding-left:11px;padding-bottom:1px;}.mozilla button.secondary.disabled span{padding-right:10px;padding-left:11px;}.mozilla .list-search .clear-search{position:relative;top:2px;}.mozilla .scroll-table .scroll-table-title h2{line-height:1.5em;}.mozilla .live-chat-option .chat-title{left:75px;}Win32 .safari #pageTop{background-color:black;}msie #pageTop{background-color:black;}msie #enter-password{width:110px;}msie #content{overflow-x:hidden;}msie .container-box{border:1px solid #dddddd;}msie #current-page-number{margin-top:2px;}msie #breadCrumb #lastCrumb{border-left:none;padding-left:12px;}msie .list-search .clear-search{position:relative;float:right;right:10px;top:-10px;}msie .list-search{width:190px;}msie .list-search input.gwt-TextBox{position:relative;width:150px;top:6px;}msie .getting-started-header{top:3px;}msie #please-wait-modal-content{display:block;}msie #file-attachment{padding-left:10px;}msie .scroll-table .check{text-align:left;}msie .live-chat-option{height:auto;}msie .live-chat-option .chat-title{left:75px;}msie .getting-started-header{top:5px;}#version8 .getting-started-header{top:4px;}msie .action-bar .action-item a{background-color:#f1f1f2;}msie .action-bar .action-item a:hover{background-color:#fafafb;}msie .scroll-table .scroll-wrapper{overflow-x:hidden;}msie .apply-to-all .check{text-align:right;padding-right:2px;}msie .apply-to-all #port-column{padding-left:24px;}msie .apply-to-all #status-column{padding-left:24px;}button{padding:0px;font-size:12px;border-radius:3px;-moz-border-radius:3px;-
```

# JavaScript

- JavaScript is the language of the web
- But lots of things that make it hard to develop and maintain
  - DOM (interacting with the page) – every browser has its quirks
  - Has some poorly thought out language features



# Awesome JavaScript

- Use a DOM library that hides browser differences
  - JQuery
- Use style checker
  - jslint/jshint
- Write JavaScript like any other language
  - We've had luck using an object-oriented style
- Write JavaScript unit tests
  - QUnit
  - Jasmine
- Use a build system (Google Closure)
  - Closure compiler uses type annotations to guarantee strong typing

# Making the Server Awesome

- Web servers ('web heads') process incoming requests
- Lots of 'glue' at the web server level
- Translate from database model into presentation model
- Application Server level: Apache (mod\_php), Tomcat/Glassfish (Java), Mongrel (Rails)


## Minimize Backend Calls

- Be really careful when the server needs to call external information
- It's very easy to create performance problems by doing backend queries in for loops!
  - n+1 problems
  - “First you get all the customer ids, and then for each customer id, you make a query...”
  - SQL Subselects are bad




# Minimize Backend Calls

- At Rackspace: Log Search

Timestamp	Status	Sender	Subject
▲ 2011-11-30 00:13:39 EST	 sent	openstack-bounces+dave.king@mailtrust.com@lists.launchpad.net	[Openstack] How to using keystone with ldap
• Message Id: <a href="mailto:CAO4uuUsRHqOanW6S9j:Z3-fm_0K7mJvteTBioNTWTEt4Du8Tg@mail.gmail.com">CAO4uuUsRHqOanW6S9j:Z3-fm_0K7mJvteTBioNTWTEt4Du8Tg@mail.gmail.com</a>			

Recipient	Status	Detail Status	Client IP	Relay
▲ dave.king@mailtrust.com	 sent	delivered via deliver service	adelie.canonical.com[91.189.90.139]	deliver
• Speed: 0s • Relay: deliver				

Timestamp	Client IP	Hop Server	Relay	Details	Local Id	Next Id	Bytes
2011-11-30 00:13:39 EST	91.189.90.139	exp2.gate.dfw1a	ord1b[0.0.0.0]	250 2.0.0 Ok: queued as AD/12-06432-48BB5DE4	73/4D-03599-38BB5DE4	AD/12-06432-48BB5DE4	4744
2011-11-30 00:13:40 EST	98.129.184.12	smtp11.gate.ord1b	store15a.mail.ord1b.rsapps.net [10.130.64.15]	250 2.0.0 Ok: queued as 61C10140070	AD/12-06432-48BB5DE4	61C10140070	5445
2011-11-30 00:13:40 EST	10.130.68.11	store15a.mail.ord1b	deliver	delivered via deliver service	61C10140070		5582

# Minimize Backend Calls

- Incoming messages hit many servers, each server generates a log line
- Querying log lines for matching messages
  - ... and fetching all message information for each matching log lines (slow)
  - ... and displaying these as possibly incomplete messages (fast)

# Avoid Long Lived Connections

- Web servers have a number of request handlers (thread pool)
- If your server can't keep up with incoming requests, it will fall over
  - Nobody likes this :(

# Avoid Long Lived Connections

- File downloads are a common long-lived connection
- At Rackspace: streaming attachment downloads
  - User has a mailbox at one datacenter, web servers in another datacenter
  - Serve attachment downloads from the same datacenter as the user's mailbox
  - Web servers aren't tied up by long-lived download threads

# Making the Data Store Awesome

- The data store contains most of the value that customers are trying to get out of your application
- Easy for early database design decisions to become horrible

## Be Careful with ORMs

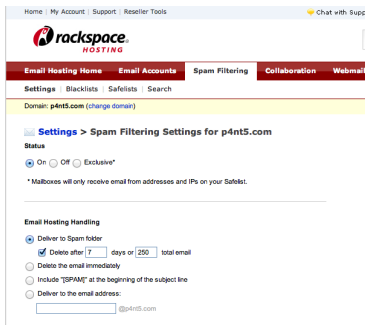
- ORMs (Object-Relation Mappers) map tables to objects
  - Hibernate (Java), NHibernate (C#)
  - Active Record (Ruby on Rails)
- Makes it easy to...
  - write database-interacting code
  - get to your value-add faster
  - introduce extremely non-performant behavior
- Eventually, you will have to understand what your ORM is doing under the hood

# Database Slaves

- Reads and writes have different usage patterns
- Most traditional applications are heavy on reads, light on writes
- Make writes to the master DB, reads go to the slaves

# Database Slaves

- At Rackspace: Spam Policy in the Control Panel
- Incoming mail needs to check a user's spam policy to route it correctly
  - Is it a whitelist (never check spam)?
  - Is it a blacklist (always mark as spam)?
- Heavy incoming volume, lots of queries



The screenshot shows the Rackspace Hosting Control Panel interface. At the top, there are navigation links: Home, My Account, Support, Reseller Tools, and a Chat with Support button. The Rackspace logo is prominently displayed. Below the logo is a navigation bar with tabs for Email Hosting Home, Email Accounts, Spam Filtering (which is selected), Collaboration, and Webmail. Underneath, there are links for Settings, Blacklists, Safelists, and Search. The main content area shows the domain p4nt5.com and the current page is 'Settings > Spam Filtering Settings for p4nt5.com'. The 'Status' section has three radio buttons: 'On' (selected), 'Off', and 'Exclusive\*'. A note states: '\* Mailboxes will only receive email from addresses and IPs on your Safelist.' The 'Email Hosting Handling' section includes several options: 'Deliver to Spam folder' (selected), 'Delete after 7 days or 250 total email' (checked), 'Delete the email immediately', 'Include "[SPAM]" at the beginning of the subject line', and 'Deliver to the email address:' with an input field containing '@p4nt5.com'.



## Database Slaves


- Master/Read Slaves per datacenter
- Control Panel writes to the master
- Servers handling incoming mail query the read slaves

# Non-Relational Data Stores

- 'NoSQL': non-relational data stores
- Key-value stores (redis, memcached)
- Document-based stores (mongoDB)
- Block storage (Hadoop)
- Direct relationship between code interacting with the data and how the data is represented

# Non-Relational Data Stores

- Rackspace Email & Apps Control Panel
- NHibernate ORM
- Use memcached as a cache layer for commonly-accessed data
- On memcache miss, fetch the data from the database, otherwise retrieve from memcached servers



# How Can You Start?

- Pick a language
  - Ruby on Rails
  - Django (Python)
- Learn JavaScript as necessary
  - Start with JQuery
- Put together something simple and iterate

# Resources

- Websites
  - Best Practices for Speeding Up Your Web Site
  - Frontend Development Guidelines
  - JavaScript Garden
- Books
  - Eloquent JavaScript
  - Agile Web Development with Rails
  - Ruby on Rails Tutorial Book
- Blogs
  - ORM is an anti-pattern

# My Website

- `http://www.davehking.com`
  - (slides posted tonight)