



CoolingTower and CirrAS

Bob McWhirter
JBoss Fellow



CoolingTower and CirrAS

Bob McWhirter
JBoss Fellow

Who is Bob?

- JBoss Fellow
- Chief Architect of Middleware Cloud Computing
- Founder of...
 - The Codehaus
 - Drools
 - TorqueBox

Who is Marek?

- Marek Goldmann
- Polish!
- Electronic musician
- **CirrAS** Lead

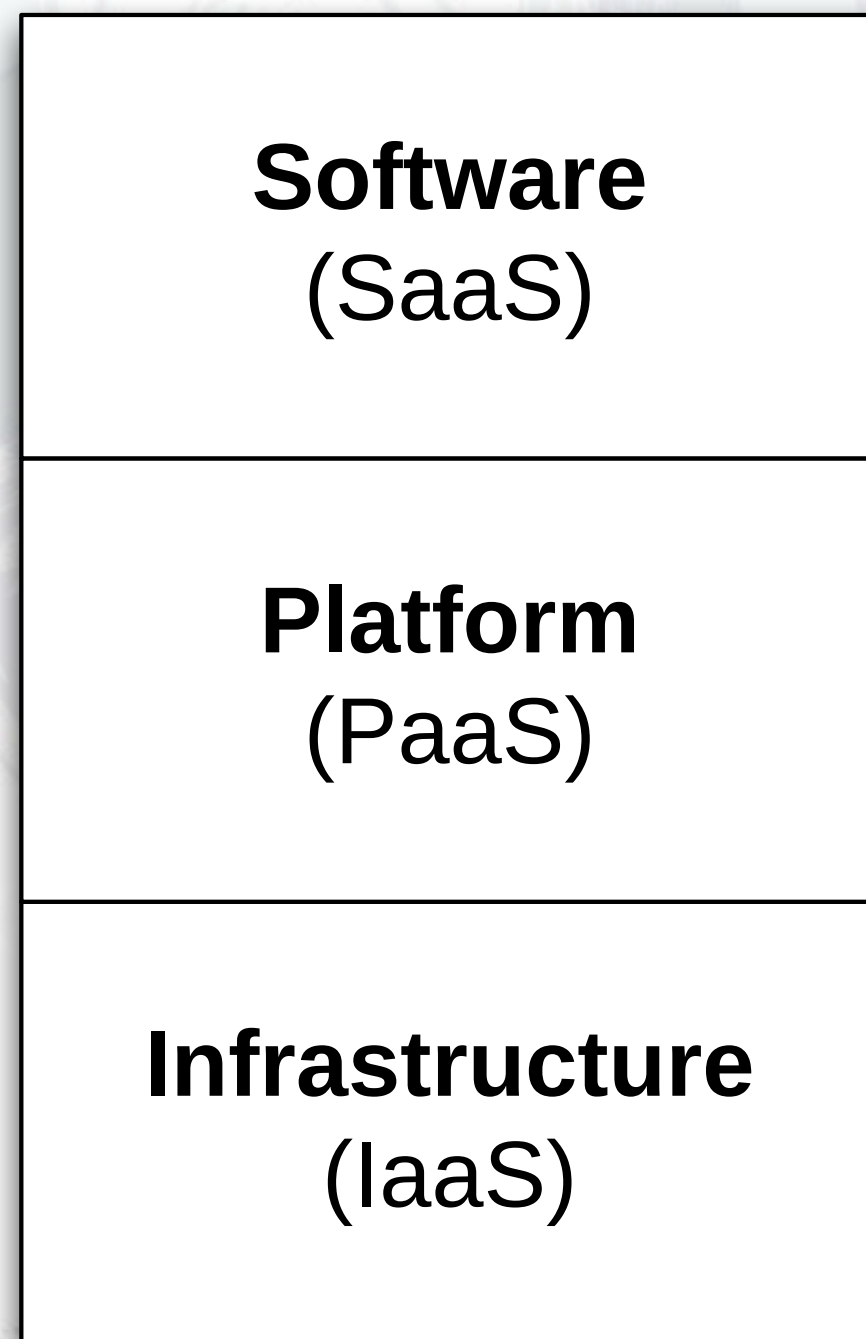


Who is Jim?

- Jim Crossley
- Loves grits!
- Acoustic musician
- **CoolingTower** Lead

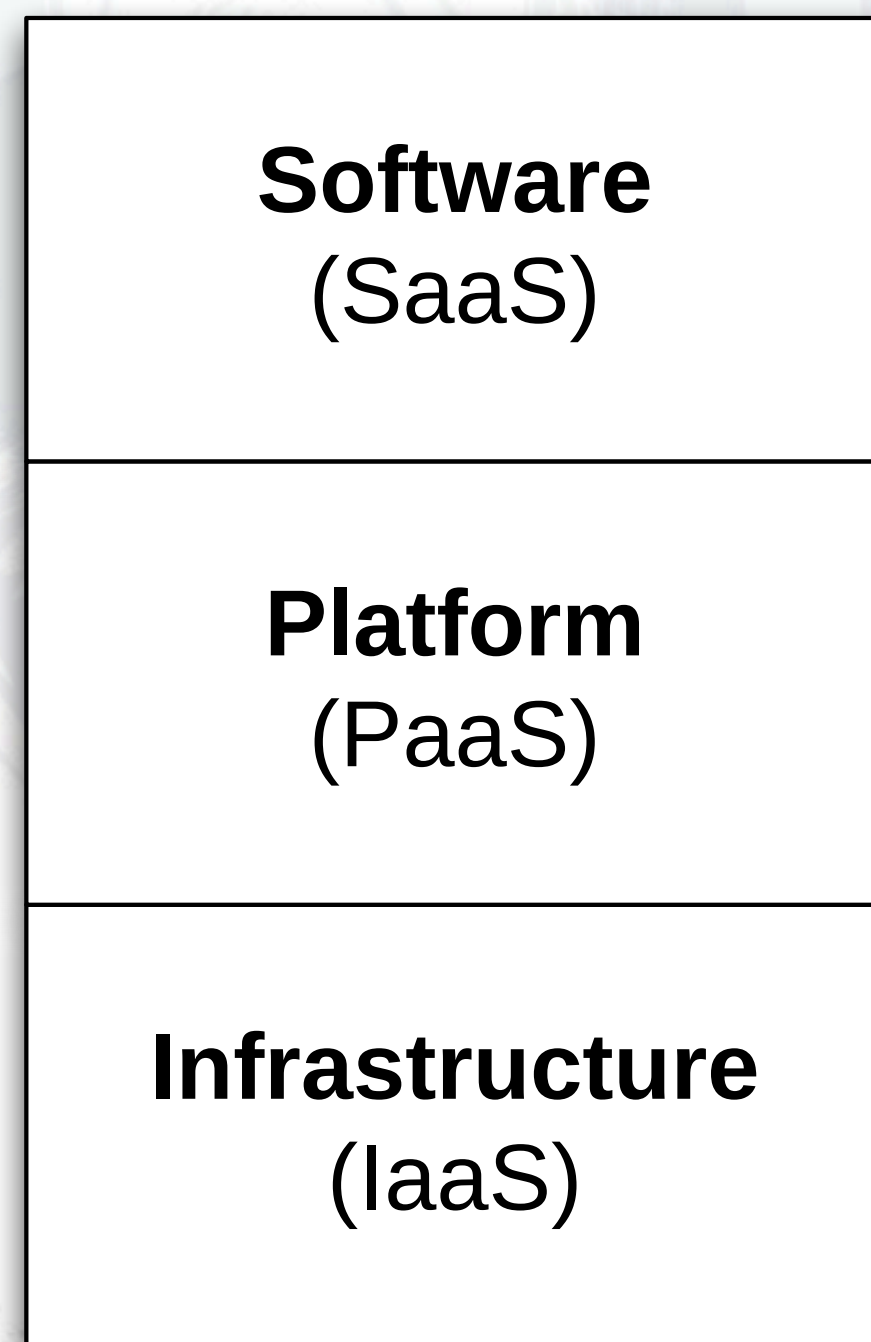


Background



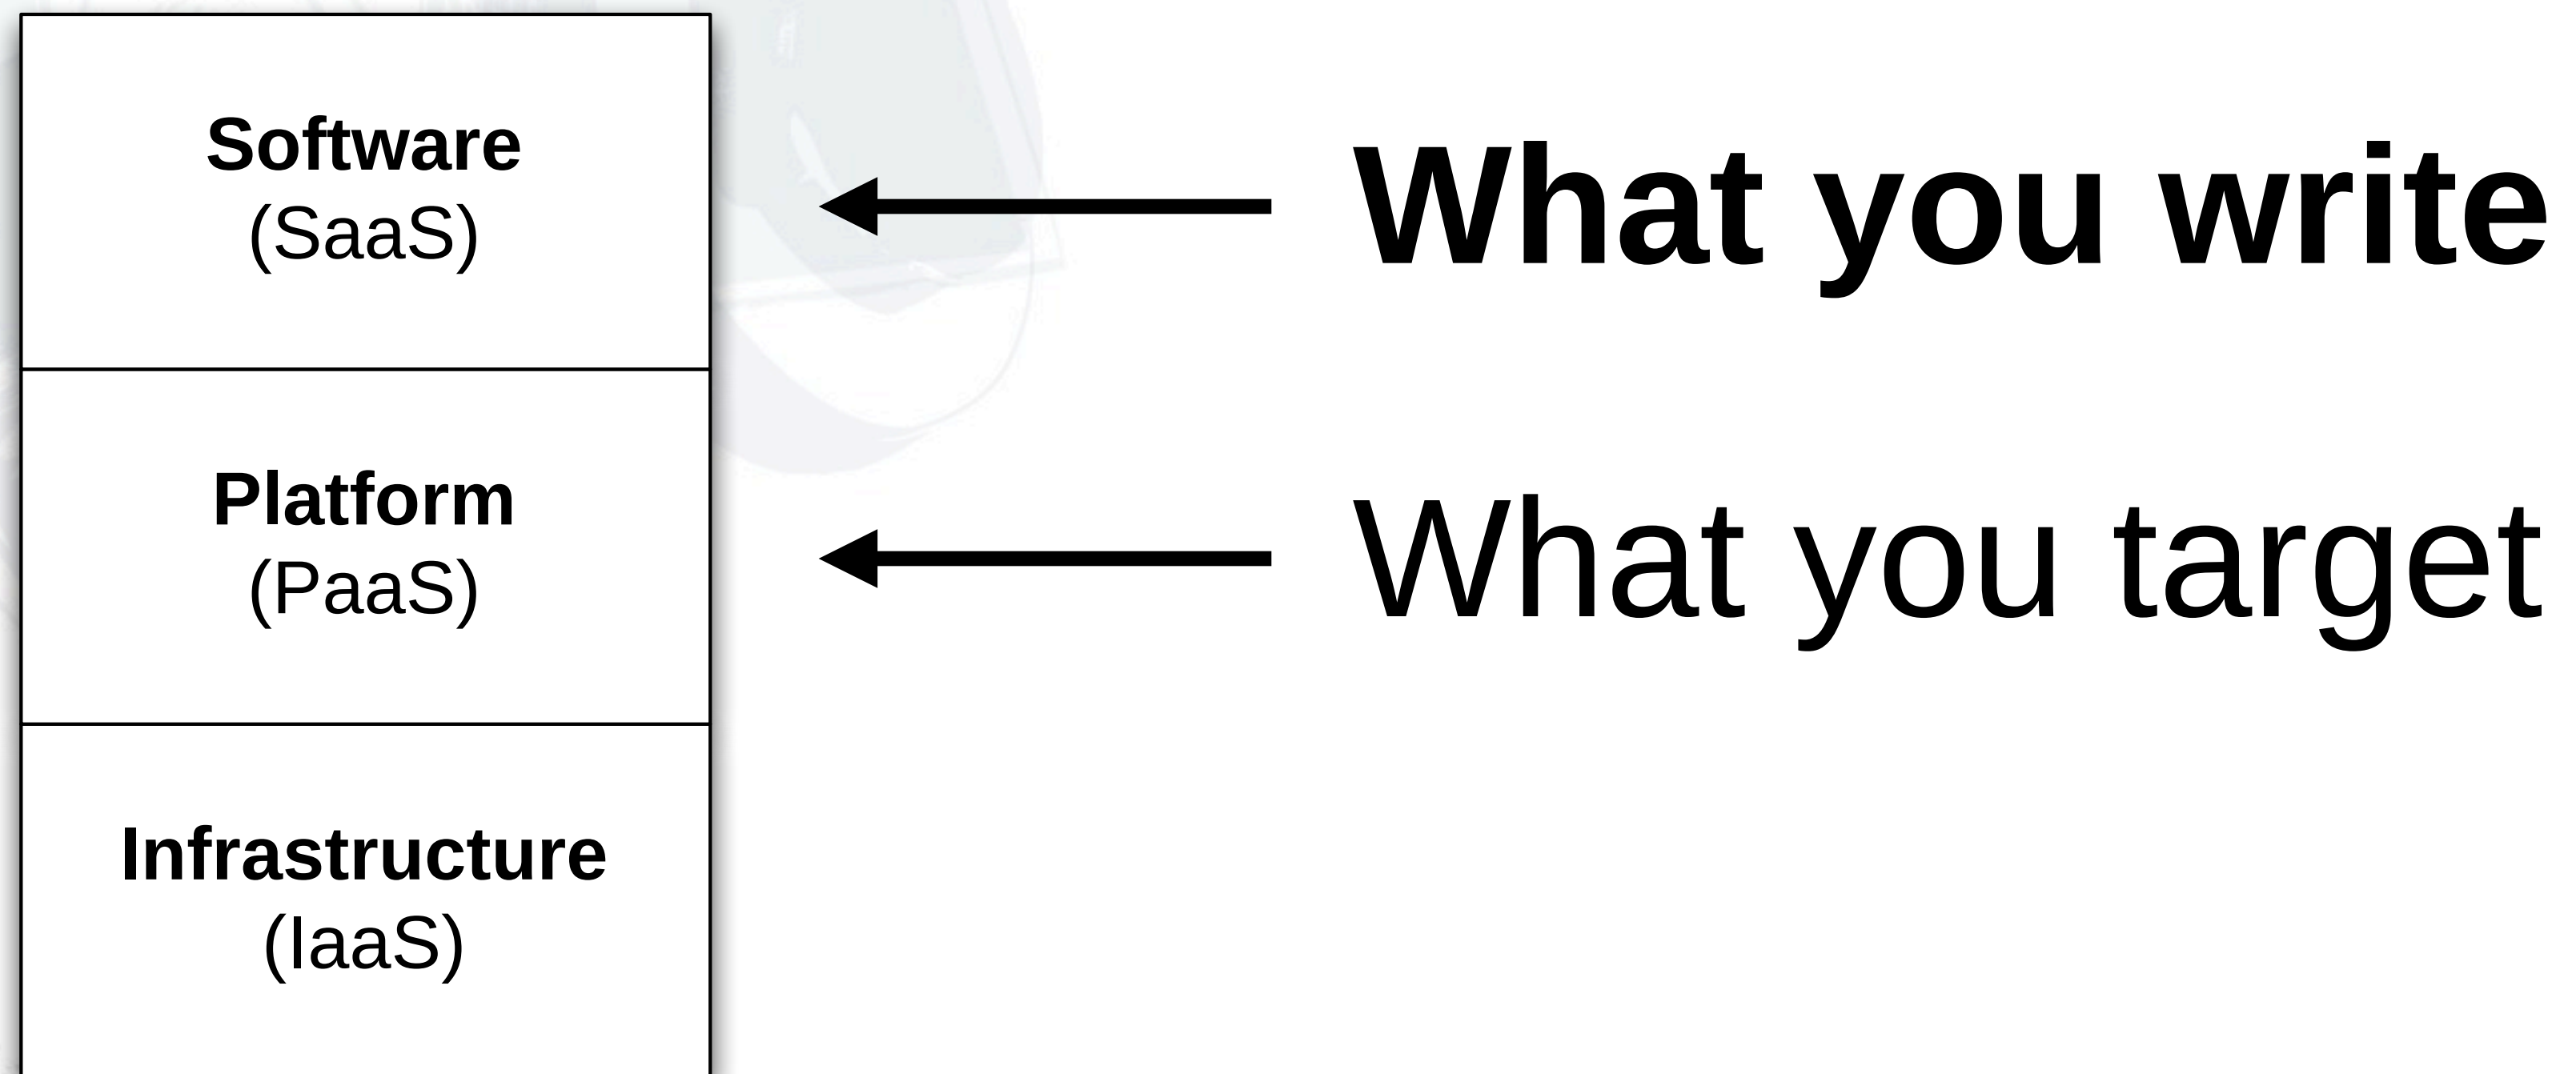
**Simplified
“cloud” stack**

Background

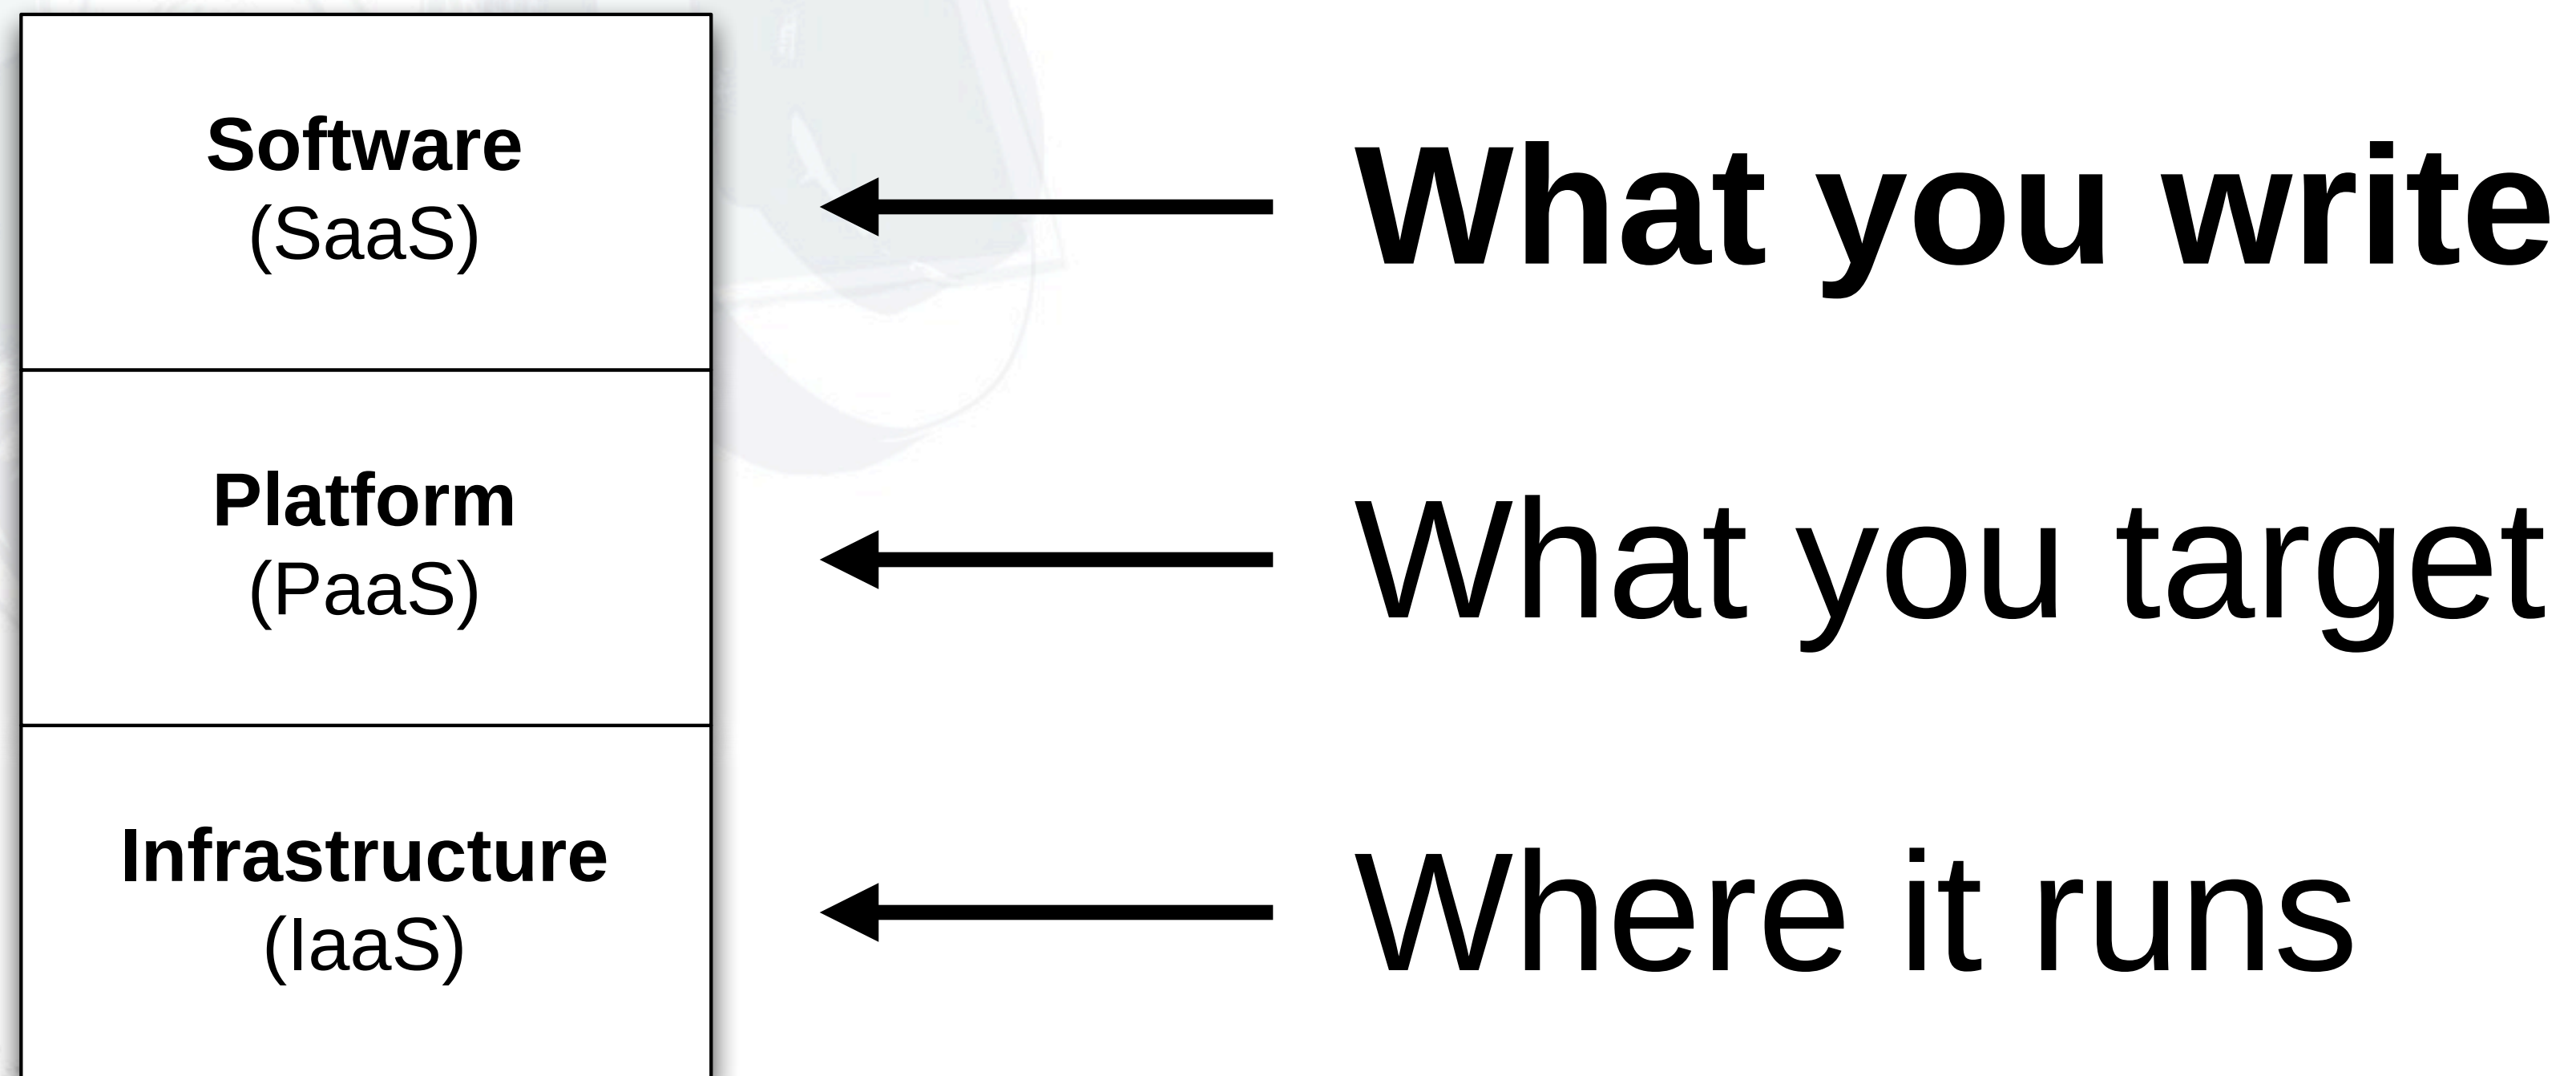


← **What you write**

Background



Background



Background

Software
(SaaS)



What you write

Platform
(PaaS)



What you target

You don't want to care

Infrastructure
(IaaS)



Where it runs

Background

Software
(SaaS)



What you write

Platform
(PaaS)

But you want JavaEE

Infrastructure
(IaaS)

and easy



**What you target
where it runs**


CoolingTower

A simple-to-deploy (and manage)
cloud catalyst.

Deltacloud

To be cloud-agnostic, **CoolingTower** uses **Deltacloud** to interface with providers.

Deltacloud



Deltacloud

```
$ gem install deltacloud-core
```

```
$ deltacloud -i ec2
```

```
Starting Deltacloud API :: ec2 ::  
http://localhost:3001/api
```

```
>> Thin web server (v1.2.7 codename No Hup)
```

```
>> Debugging ON
```

```
>> Maximum connections set to 1024
```

```
>> Listening on localhost:3001, CTRL+C to stop
```


Bootstrap your PaaS

CoolingTower is a simple **Ruby-on-Rails** application you deploy in **JBoss**.

*See our **TorqueBox** talk at 5pm!

CoolingTower

```
$ cd $JBOSS_HOME
```

```
$ ./bin/run.sh
```

```
...
```

```
JBossAS [6.0.0.20100429-M3  
"Neo"] Started in 43s:728ms
```


CoolingTower

Deltacloud

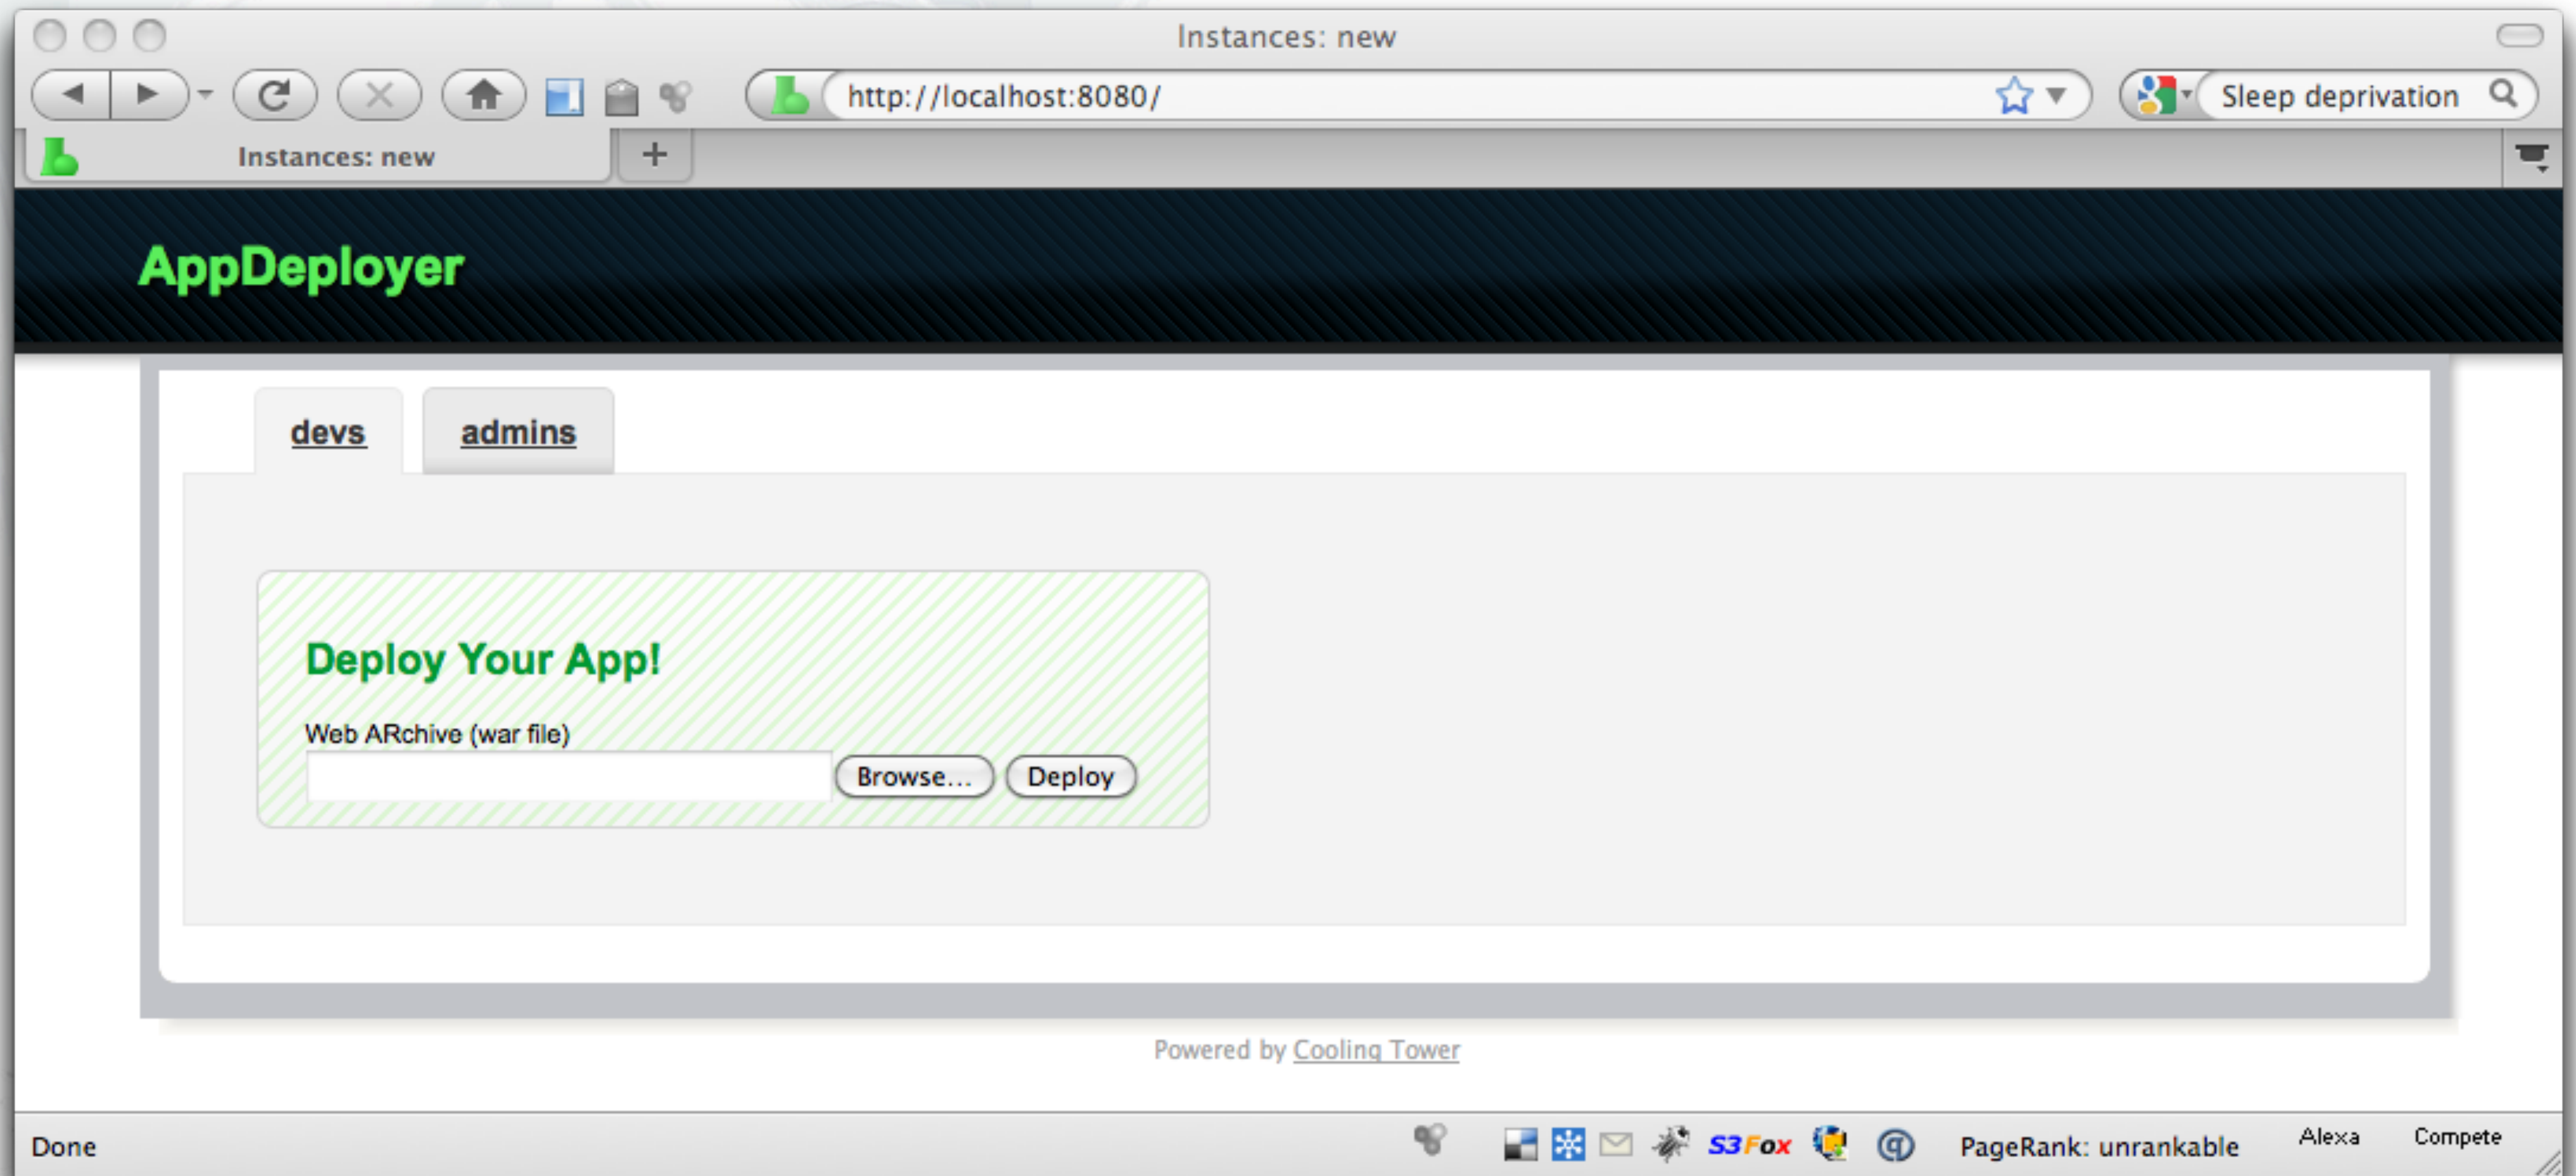
localhost



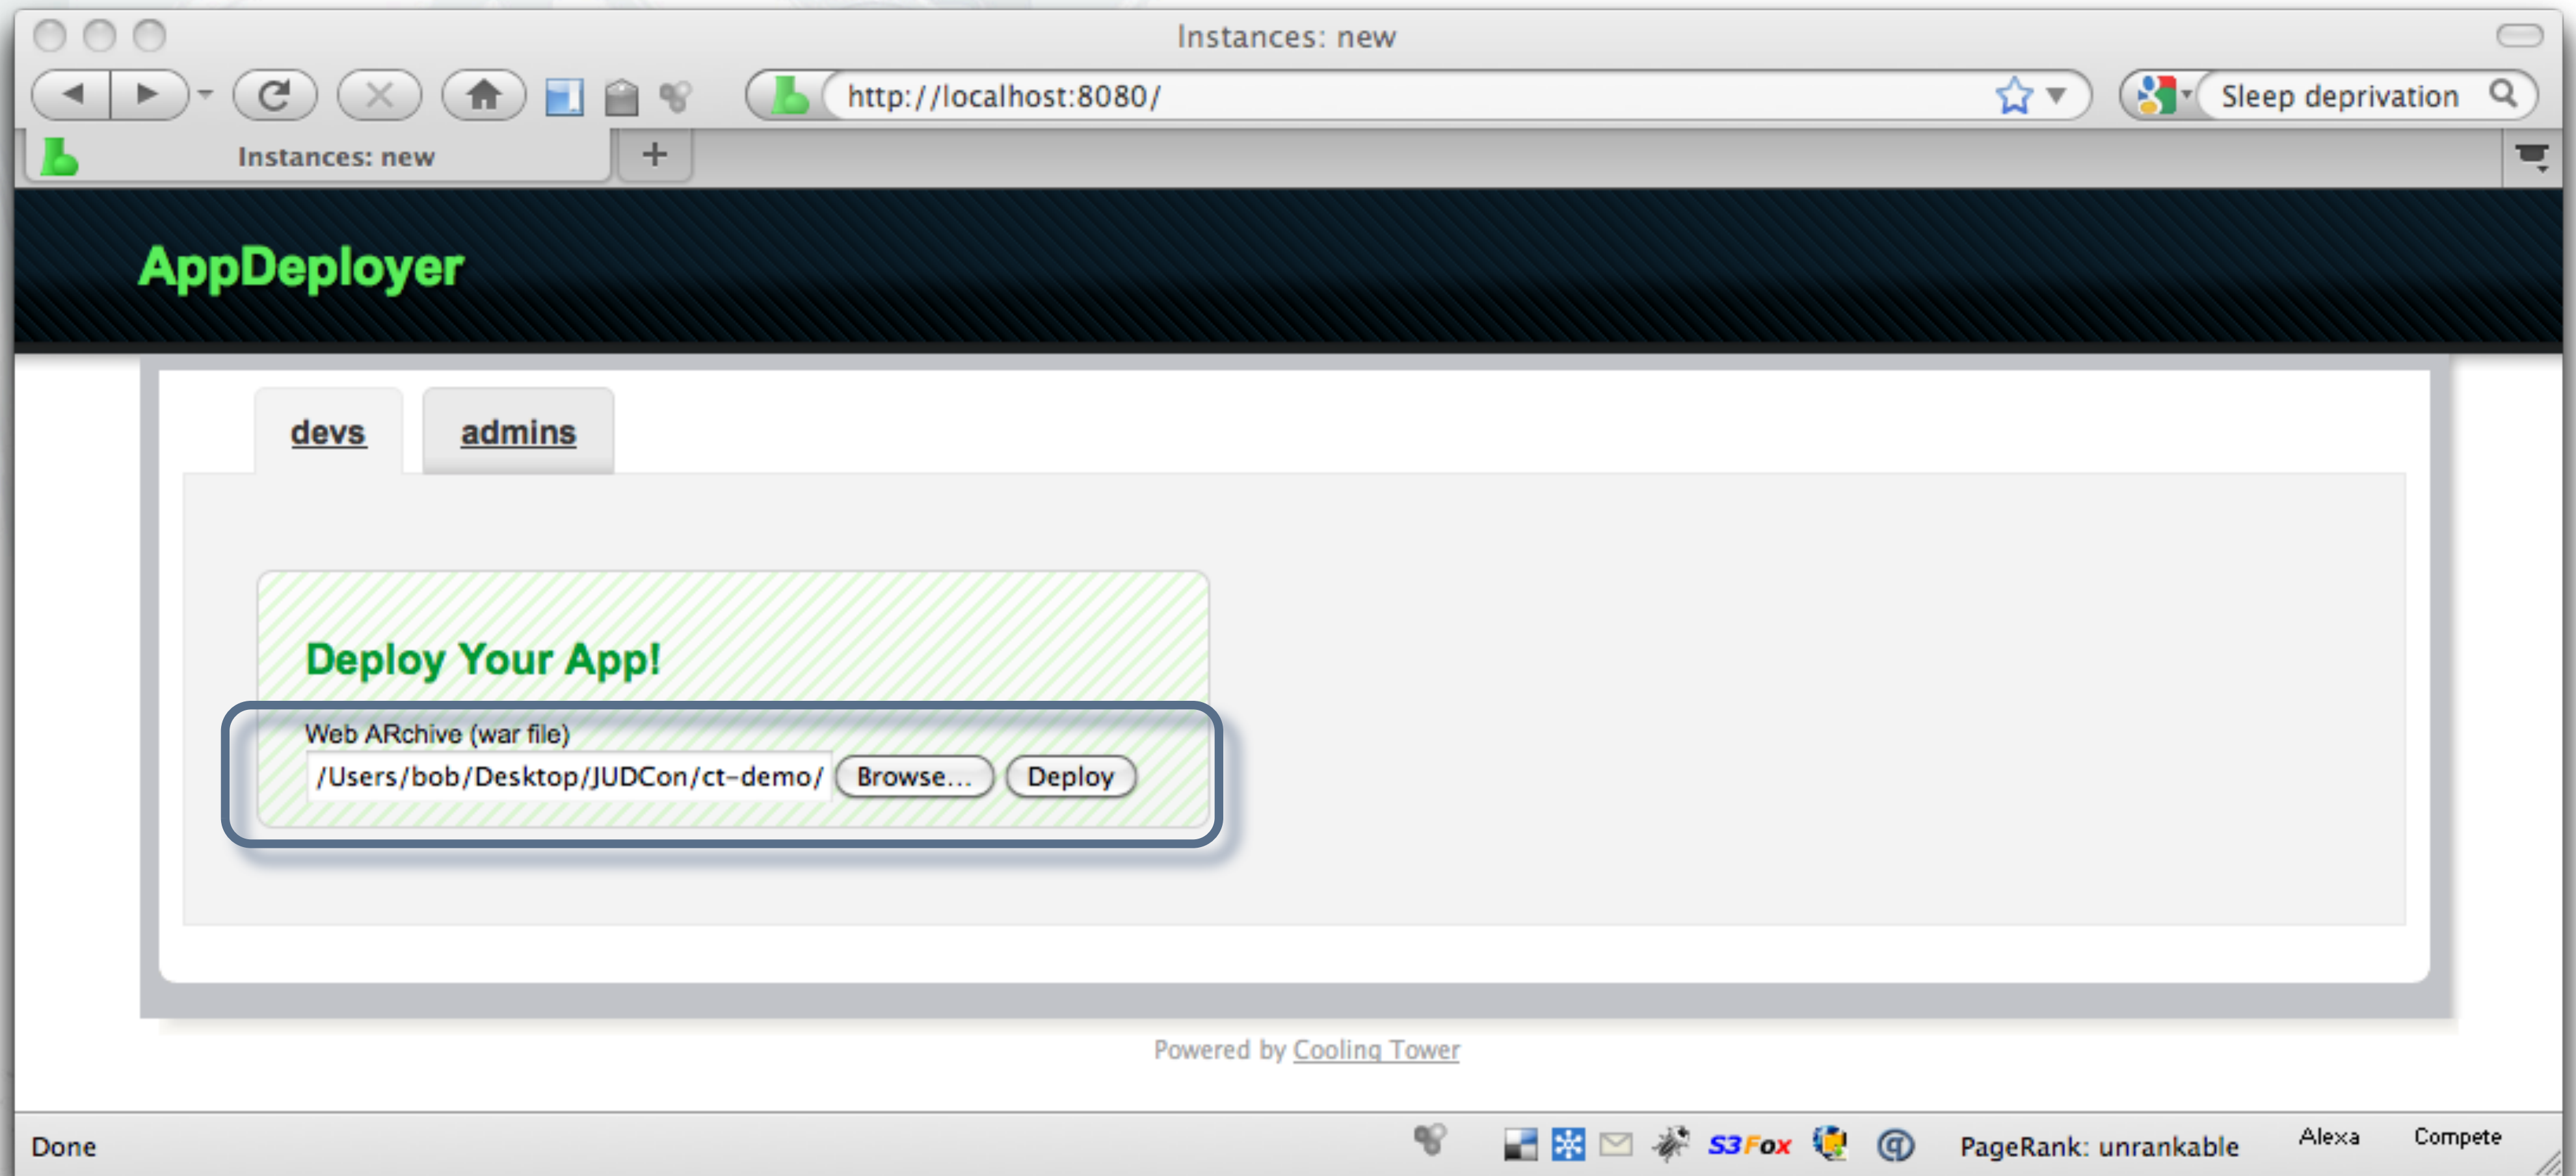
Role #1

“I want my app running”

Developer Portal

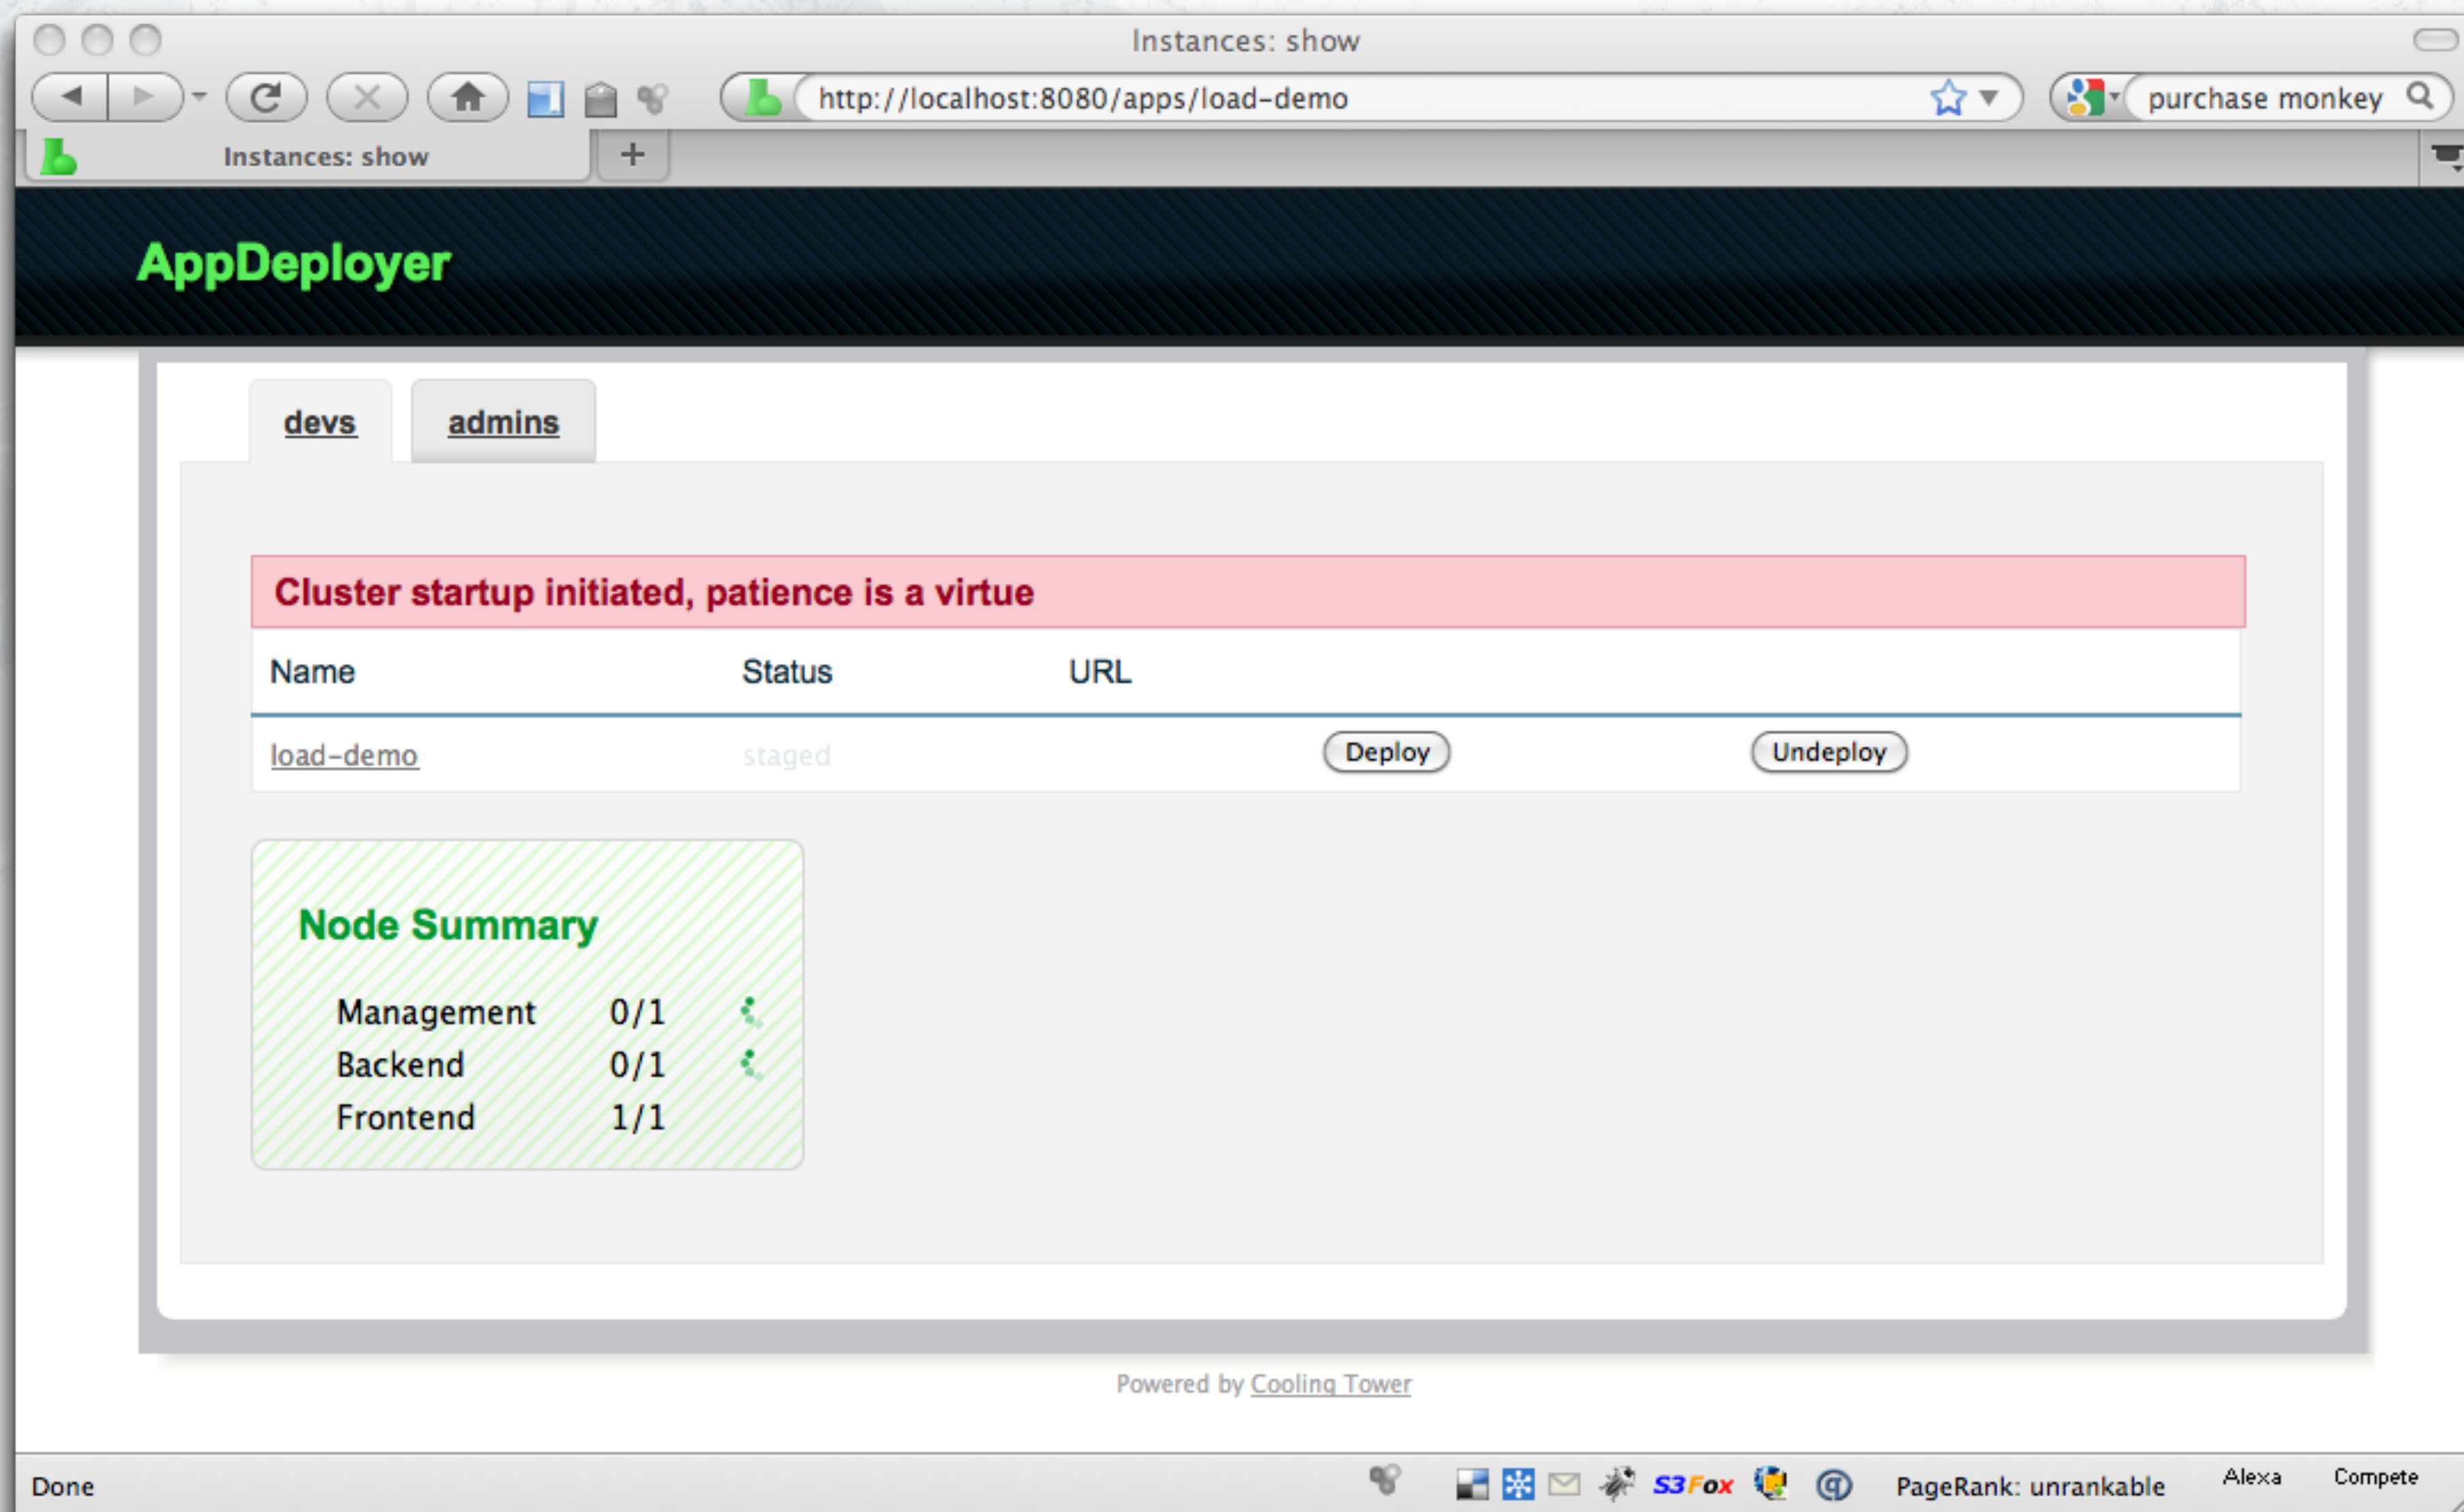


Developer Portal



Remember!

Besides **CoolingTower** and the **Deltacloud** adapter, there is **no AS** or anything running... **yet**.



Instances: show

http://localhost:8080/apps/load-demo

purchase monkey

AppDeployer

devs admins

Cluster startup initiated, patience is a virtue

Name	Status	URL
load-demo	staged	Deploy Undeploy

Node Summary

Management	0/1	⋮
Backend	0/1	⋮
Frontend	1/1	⋮

← Cold start!

Powered by [Cooling Tower](#)

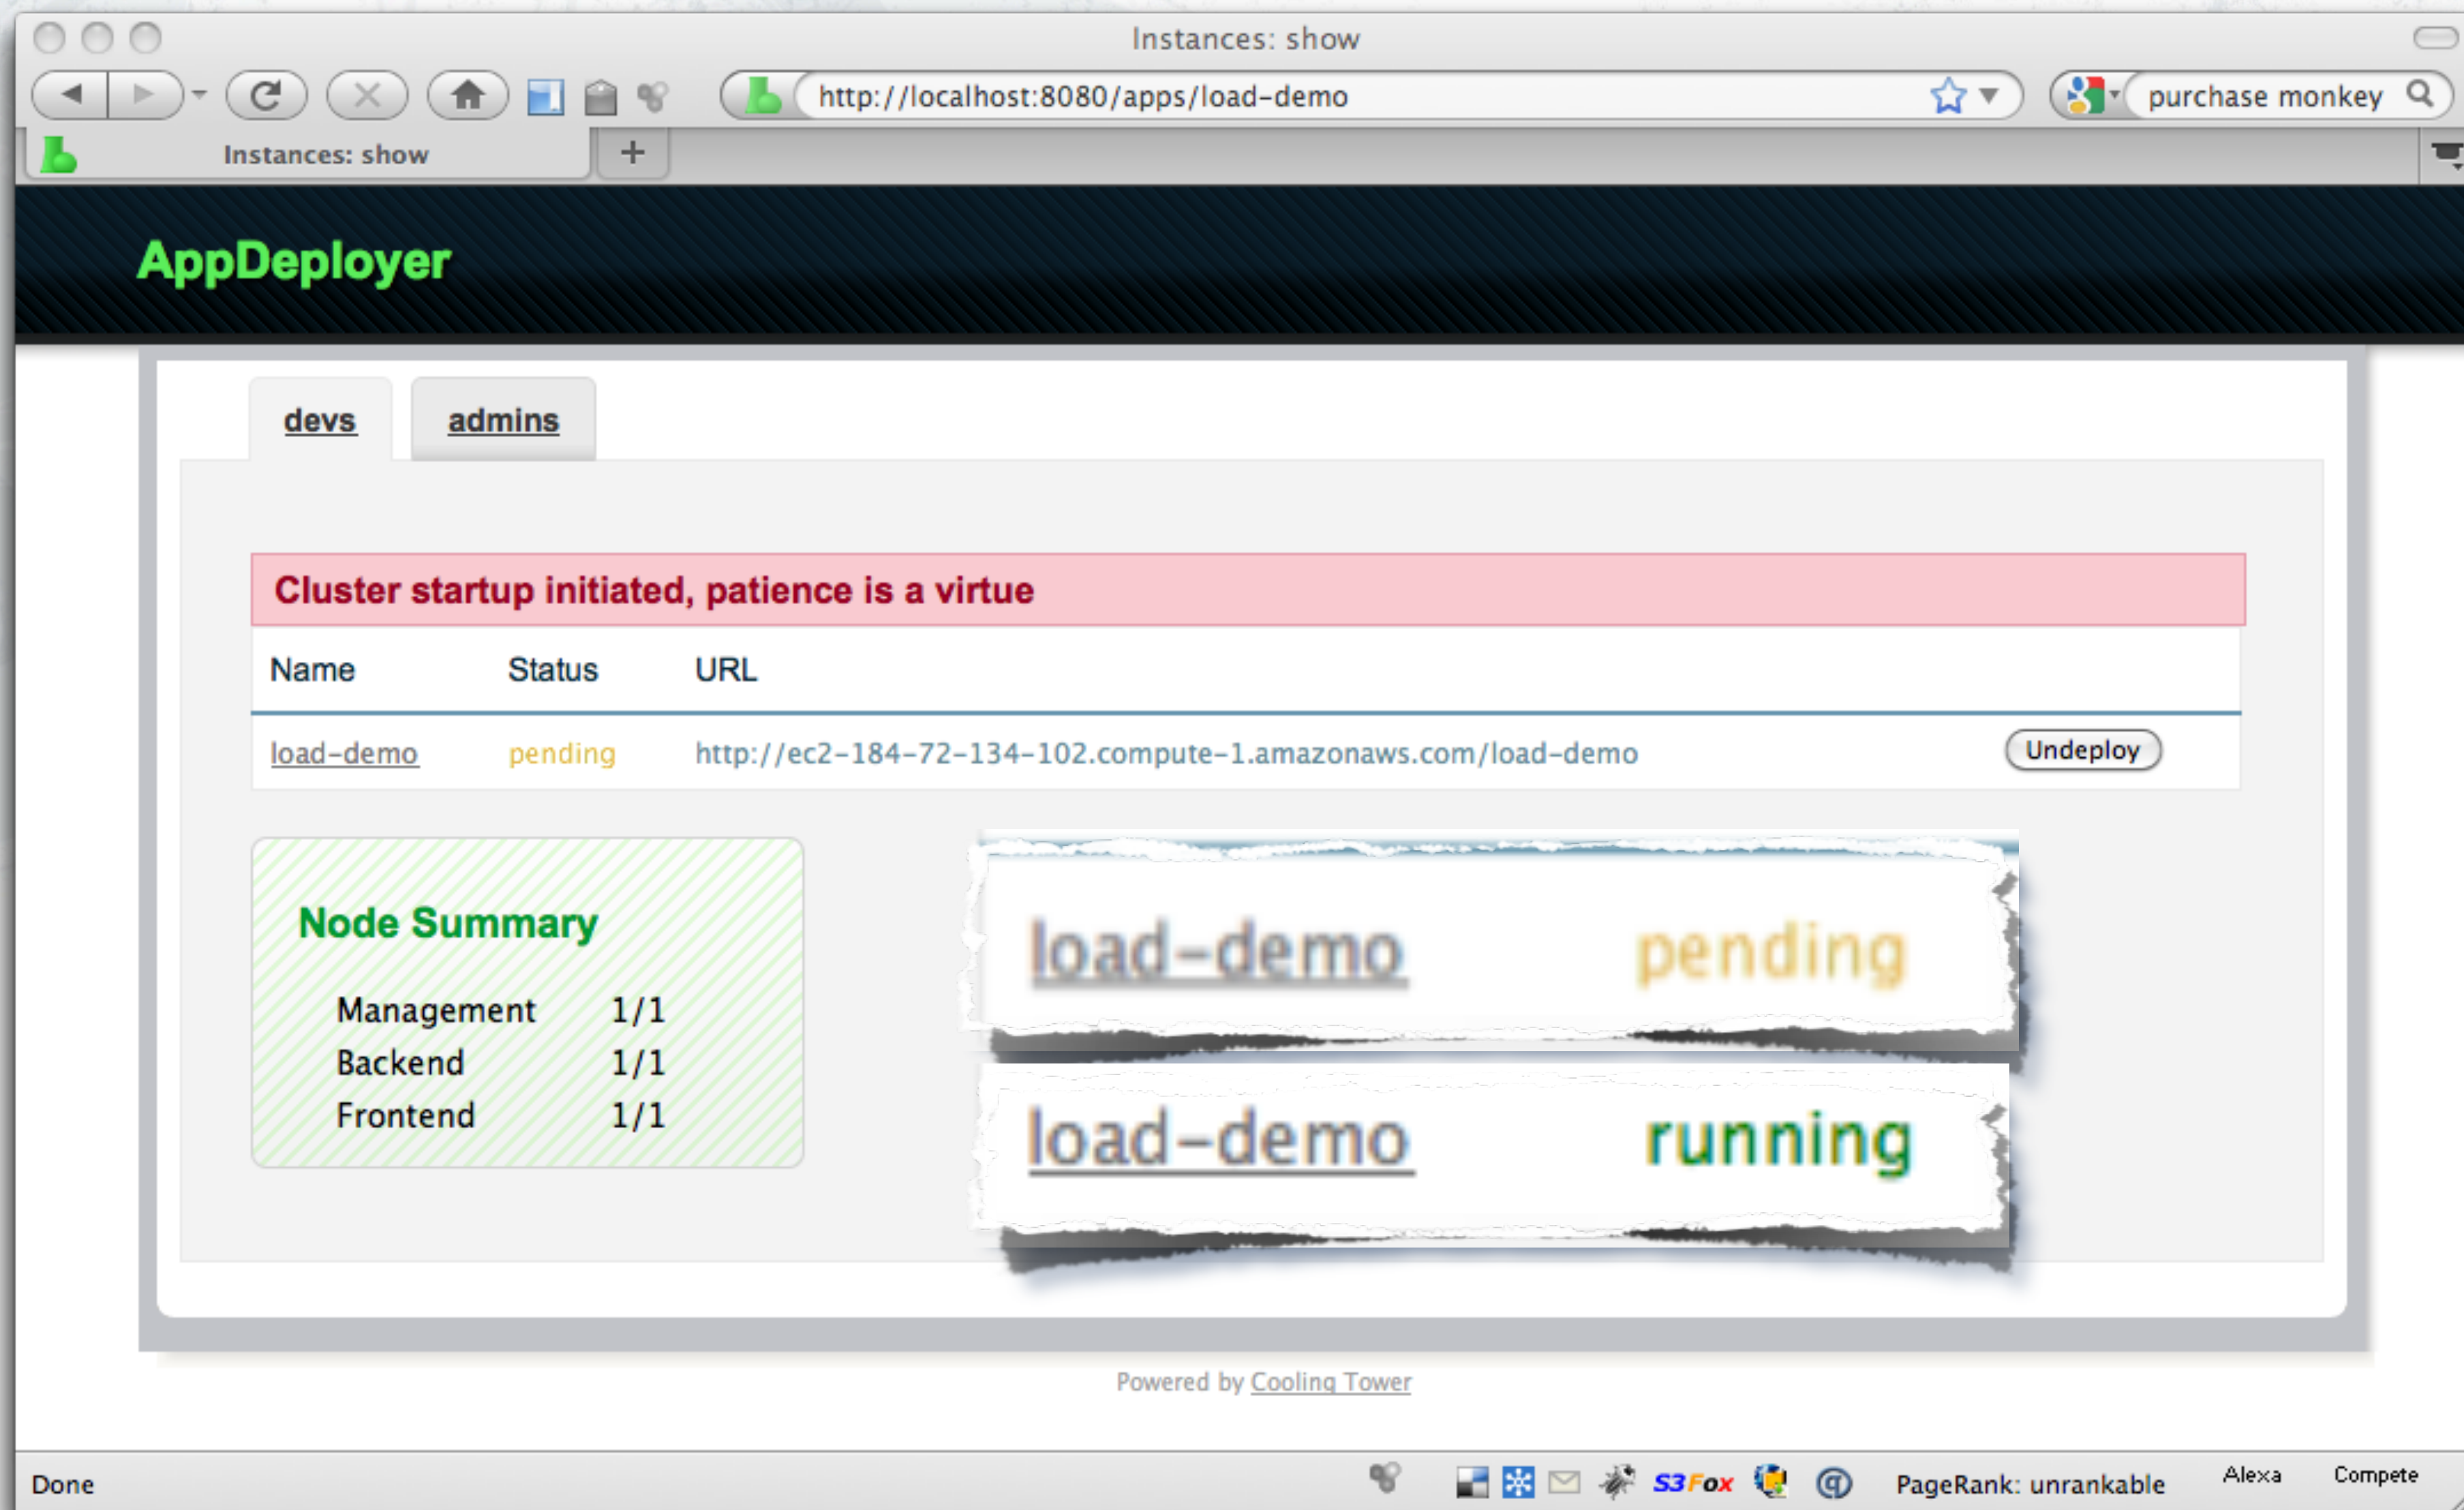
Done

S3Fox

PageRank: unrankable

Alexa

Compete





Your app is **now deployed.**



What just happened?

CirrAS

CoolingTower just used **Deltacloud** to launch a **JBoss AS** cluster, plus a front-end tier of **mod_cluster**.

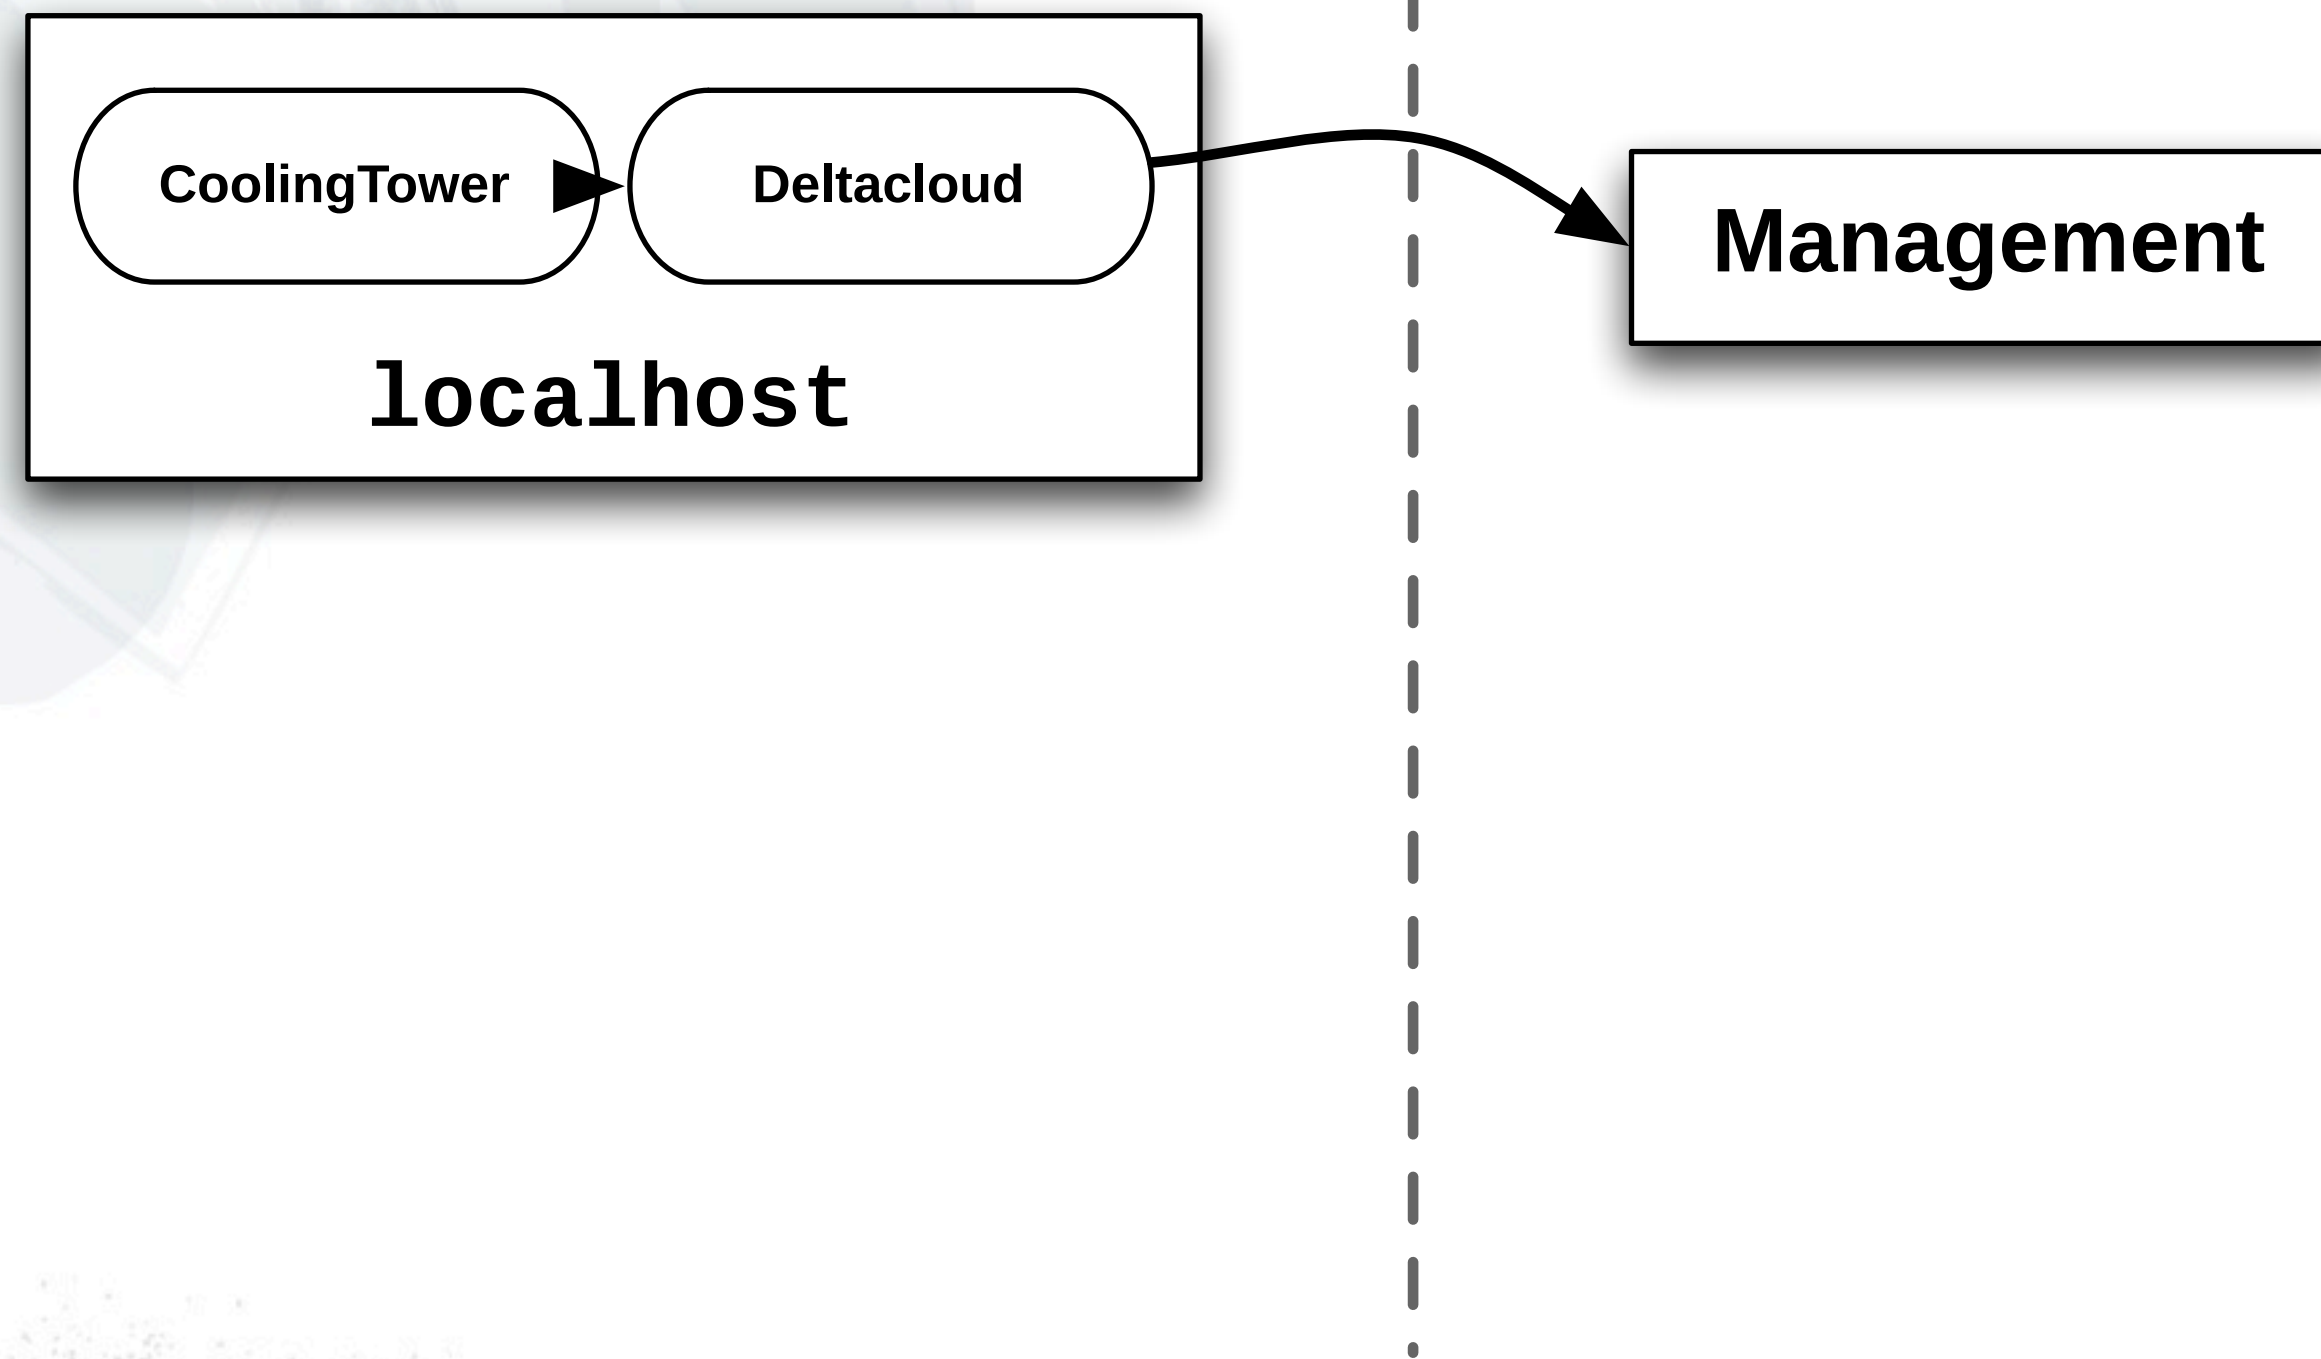
These are **CirrAS** appliances.



CIRRAS

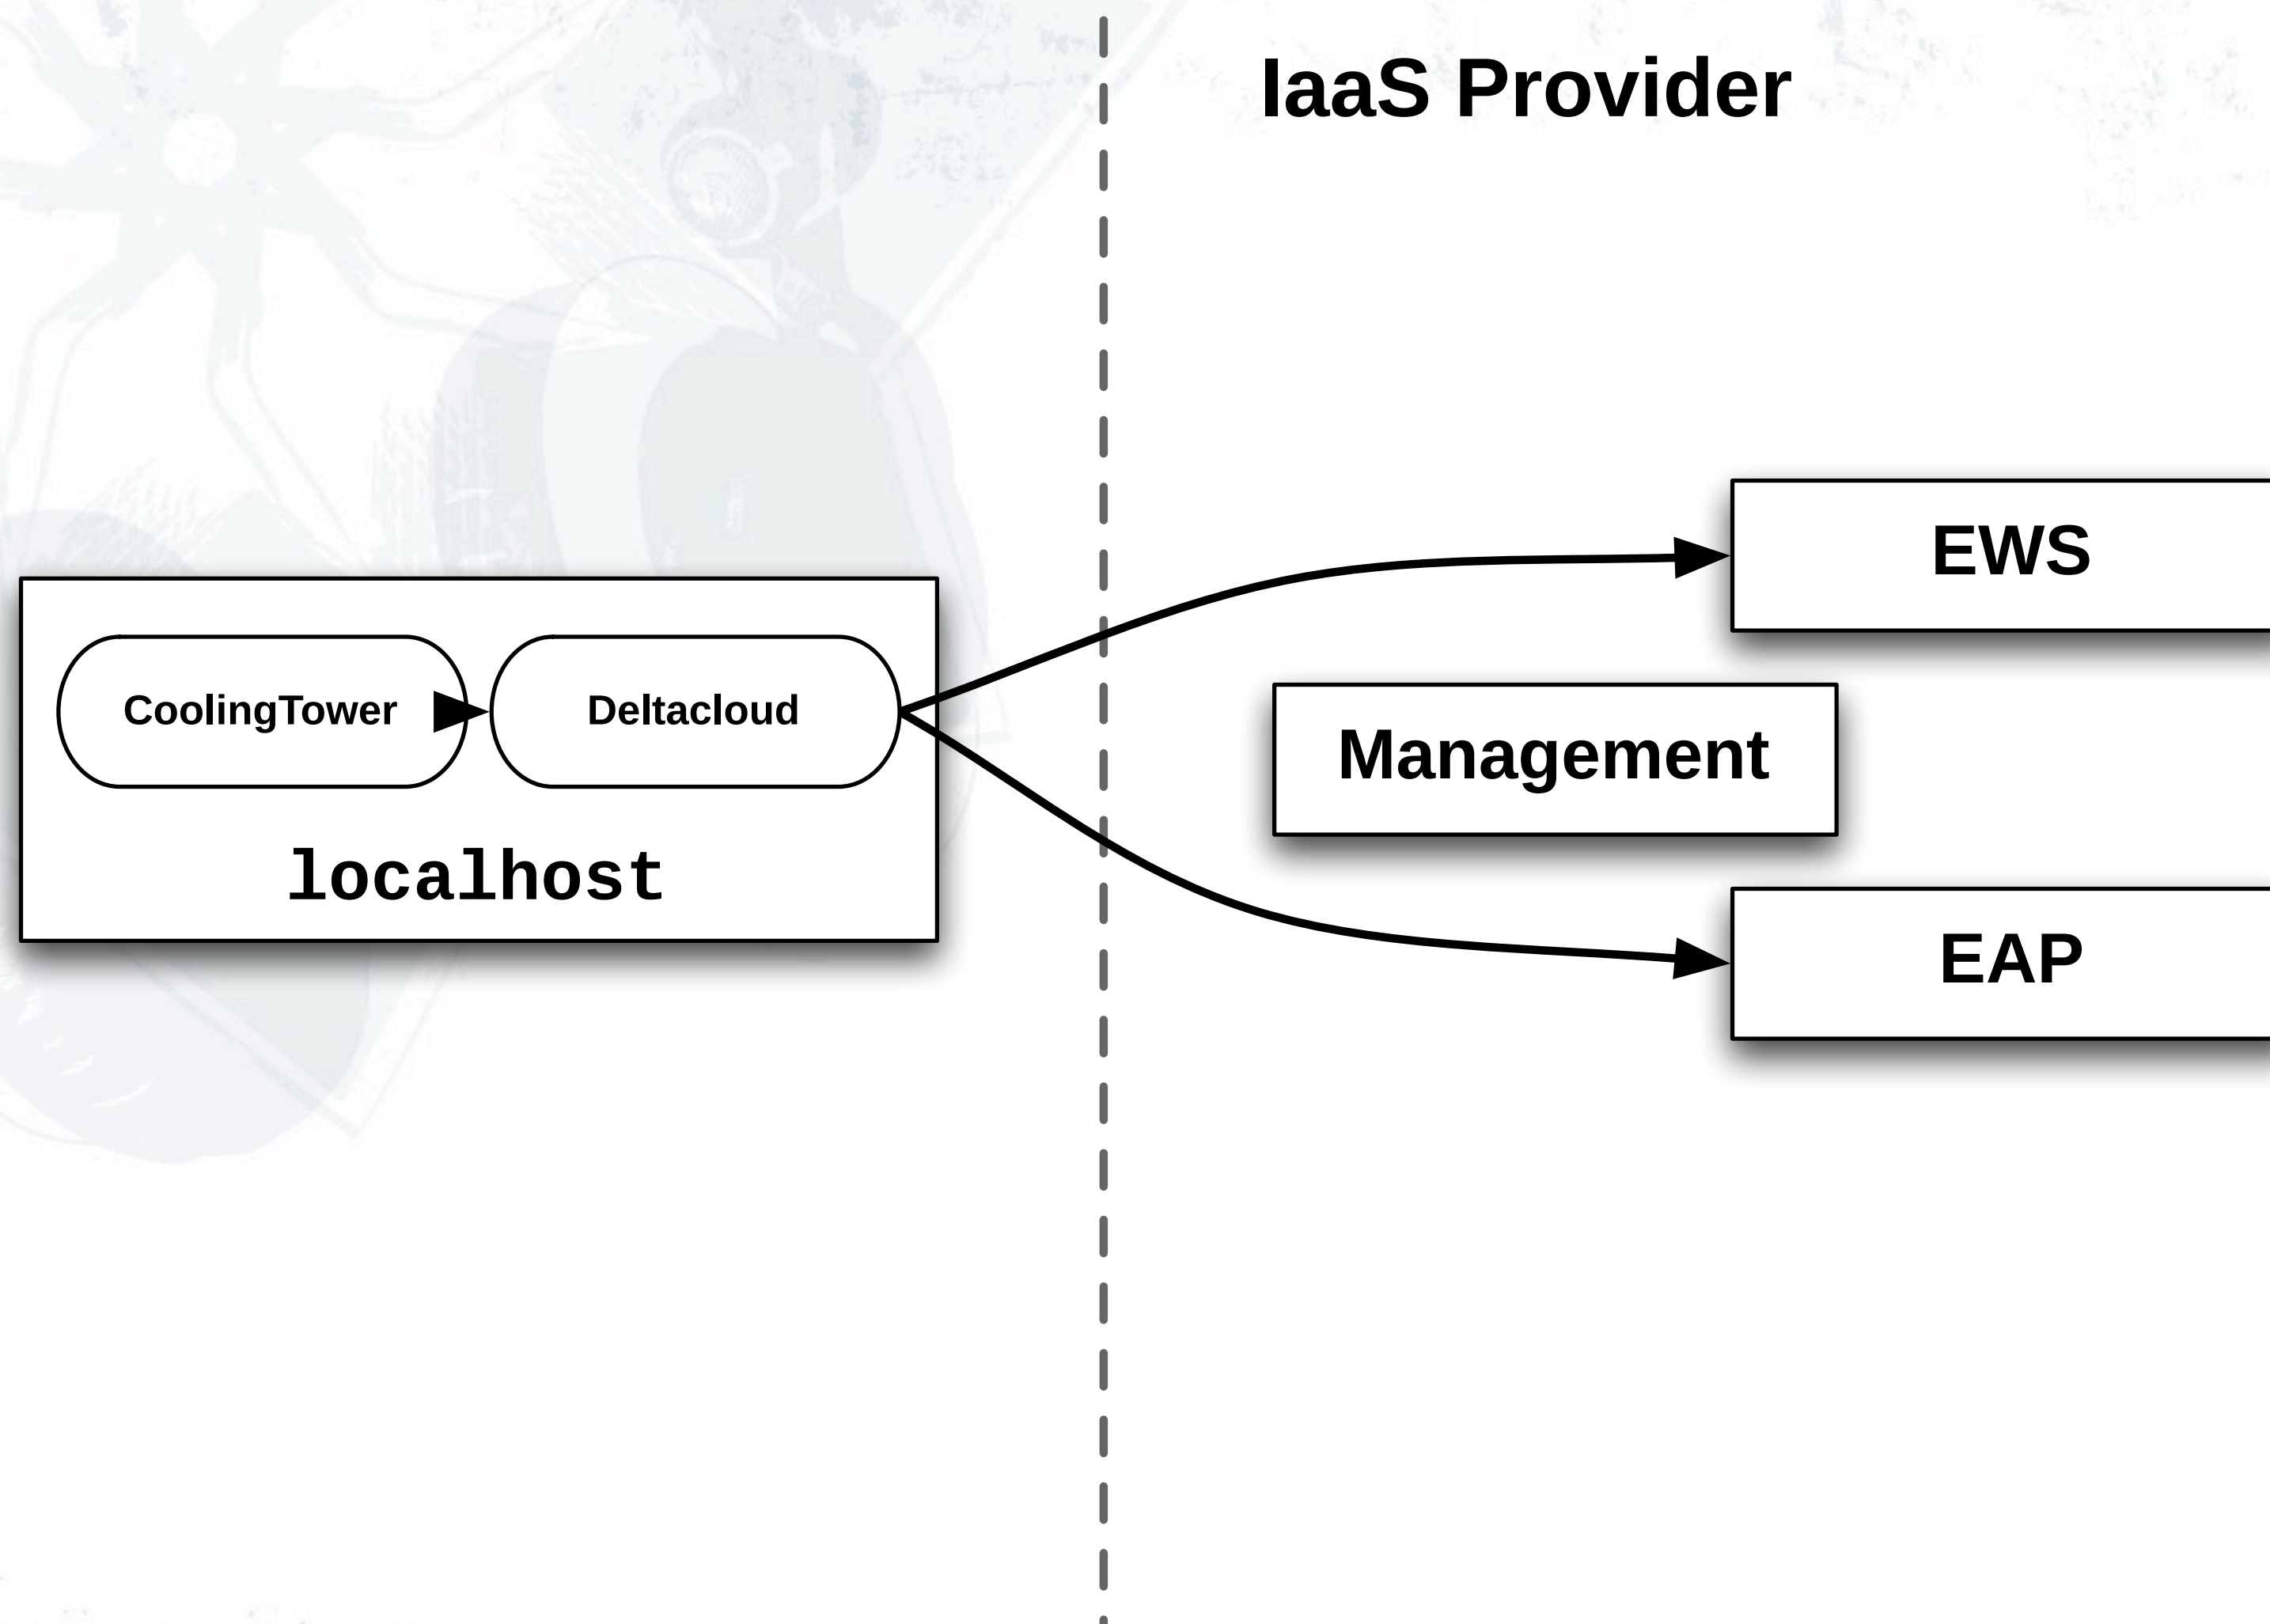
First, send out a scout...

- Launch a **management node**, the spine upon which the PaaS rests.
- Cluster-specific configuration passed to management node.
 - AWS credentials
 - S3 bucket for clustering
- Coordinates the weaving of all other nodes.
- Includes **RHQ/JON/JOPR**.



Then launch the tiers...

- Launch 1 node of **EWS**
 - **mod_cluster**
 - Apache **httpd**
- Launch 1 node of **EAP**
 - **JBossAS**




Management Node, **activate!**


Management node **watches the
IaaS**, notices new nodes appearing.

Cluster Configuration

- Tells each **EAP** node about each **EWS** node
- Tells **EAP** nodes about the S3 bucket for clustering.
- Restarts services on nodes as appropriate.




IaaS Provider




Cluster Formation

- **EAP** nodes rendezvous on **S3** for cluster discovery.
- **EAP** node informs **EWS** nodes of cluster view.



IaaS Provider





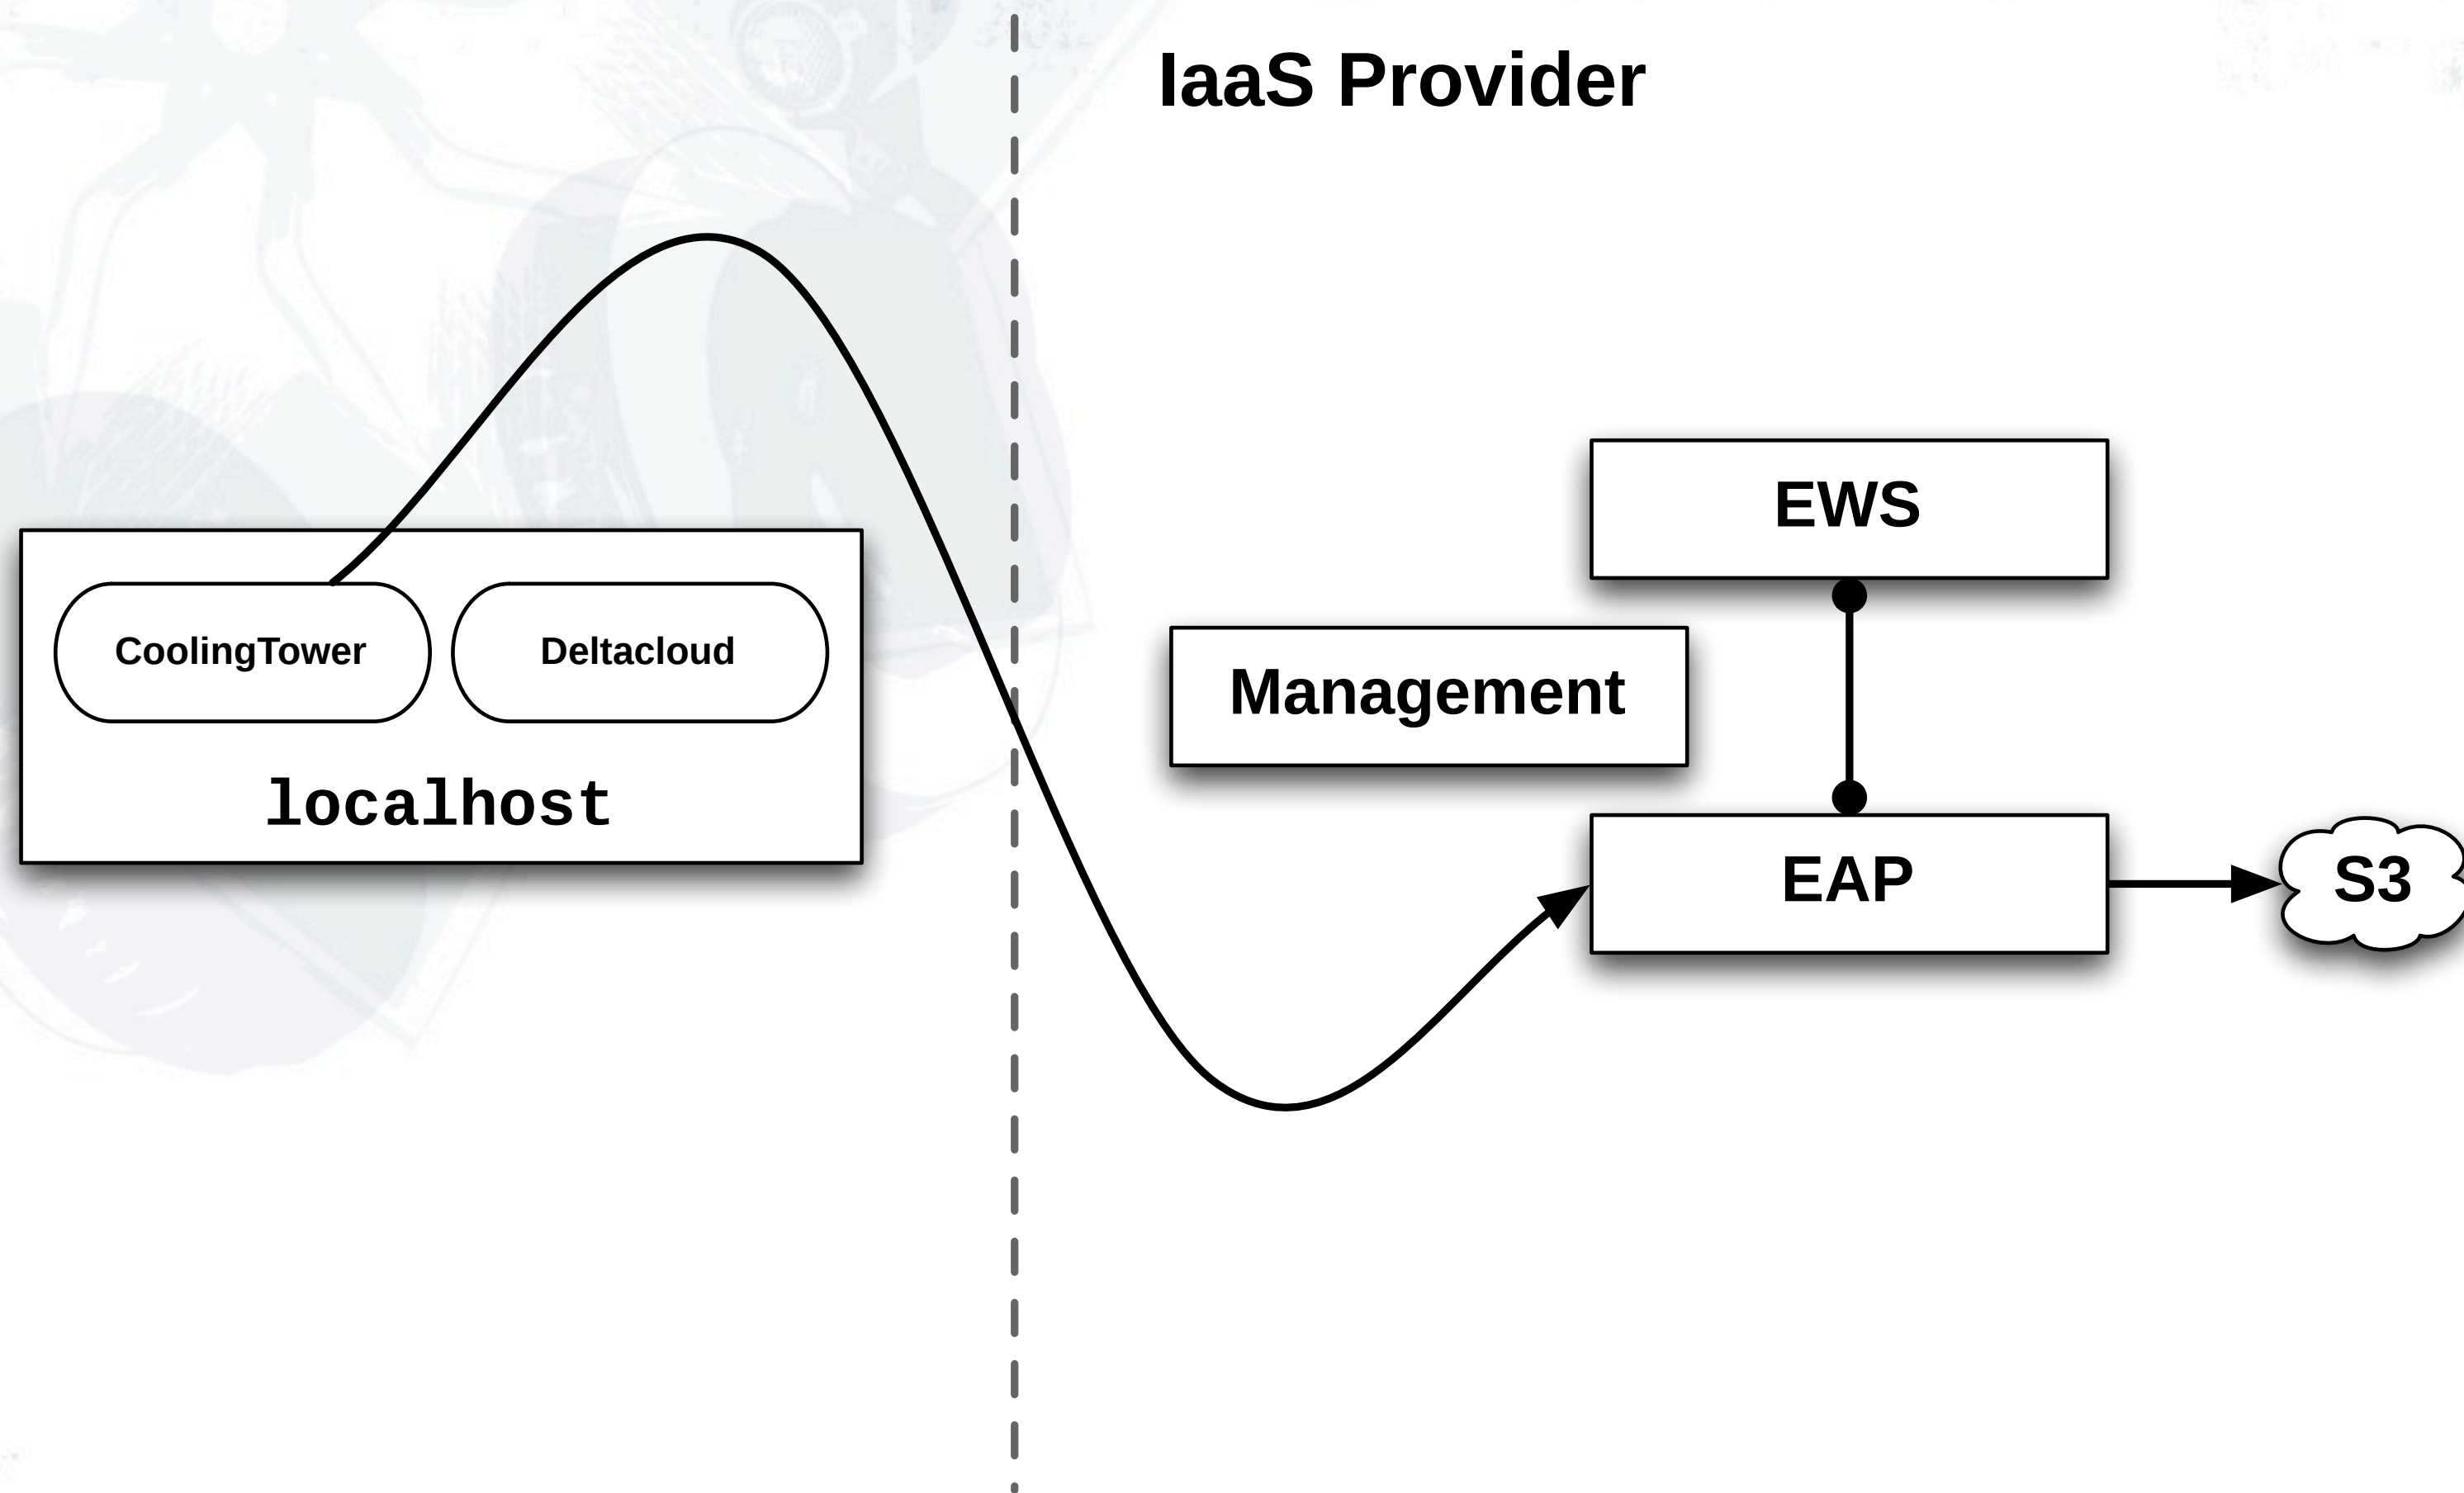
There ya go...

**From nothing to minimal 2-tier
cluster *in 1 click.***

but that's not all!

Your app gets deployed

CoolingTower wakes up, notices an **EAP** node is ready, and deploys your app.



Gives you a URL

The screenshot shows a web browser window with the address bar at `http://localhost:8080/apps/load-demo`. The page title is "Instances: show". The main content area has a dark blue header with the text "AppDeployer" in green. Below the header, there are two tabs: "devs" and "admins". The "devs" tab is selected. Under the "devs" tab, there is a table with the following data:

Name	Status	URL	
load-demo	running	http://ec2-184-72-134-102.compute-1.amazonaws.com/load-demo	Undeploy

Below the table, there is a green and white striped bar. At the bottom of the page, there is a footer that says "Powered by [Cooling Tower](#)". The browser's status bar at the bottom shows "Done" and various icons.

[load-demo](#)

running

<http://ec2-184-72-134-102.compute-1.amazonaws.com/load-demo>

mod_cluster status

Mod_cluster Status

[Auto Refresh](#) [show DUMP output](#) [show INFO output](#)

Node node-10.196.143.159 (ajp://10.196.143.159:8009):

[Enable Contexts](#) [Disable Contexts](#)

Balancer: mycluster,Domain: ,Flushpackets: Off,Flushwait: 10000,Ping: 10000000,Smax: 1,Ttl: 60000000,Elected: 0,Read: 0,Transferred: 0,Connected: 0,Load: 92,Num sessions: 0

Virtual Host 1:

Contexts:

/load-demo, Status: ENABLED

Aliases:

localhost

SessionIDs:

Contexts:

/load-demo, Status: ENABLED [Disable](#)

Done

PageRank: unranked Alexa Compete



Role #2

“I gotta keep all of
this stuff running”

PaaS Administrator

Ultimately a PaaS is a **shared resource**. **One administrator** for **N** applications.

Instances: index

http://localhost:8080/admin

fried watermelon

Instances: index

AppDeployer

devs **admins**

Cluster Management

Shutdown All

Frontend

ID	Hostname	Status
i-75d3361f	ec2-184-72-134-102.compute-1.amazonaws.com [ssh]	running
To create a new instance, click		Down
		Up

Backend

ID	Hostname	Status
i-77d3361d	ec2-184-73-100-57.compute-1.amazonaws.com [ssh]	running
To create a new instance, click		Down
		Up

Management

ID	Hostname	Status
i-49d33623	ec2-184-73-111-126.compute-1.amazonaws.com [ssh]	running
To create a new instance, click		Down
		Up

Done

S3Fox

PageRank: unrankable

Alexa

Compete

Instances: index

http://localhost:8080/admin

fried watermelon

Instances: index

AppDeployer

devs admins

Cluster Management

Shutdown All

Frontend

ID	Hostname	Status
i-75d3361f	ec2-184-72-134-102.compute-1.amazonaws.com [ssh]	running Down

To create a new instance, click [Up](#)

i-49d33623 ec2-184-73-111-126.compute-1.amazonaws.com [ssh] running [Down](#)

To create a new instance, click [Up](#)

Done

S3Fox PageRank: unrankable Alexa Compete

Instances: index

http://localhost:8080/admin

fried watermelon

AppDeployer

devs admins

Cluster Management

Shutdown All

Frontend

ID	Hostname	Status
i-75d3361f	ec2-184-72-134-102.compute-1.amazonaws.com [ssh]	running Down

Backend

ID	Hostname	Status
i-77d3361d	ec2-184-73-100-57.compute-1.amazonaws.com [ssh]	running Down
		To create a new instance, click Up

i-49d33623	ec2-184-73-111-126.compute-1.amazonaws.com [ssh]	running Down
		To create a new instance, click Up

Done

S3Fox

PageRank: unrankable

Alexa

Compete

Instances: index

http://localhost:8080/admin

fried watermelon

AppDeployer

devs admins

Cluster Management

Shutdown All

Frontend

ID	Hostname	Status
i-75d3361f	ec2-184-72-134-102.compute-1.amazonaws.com [ssh]	running Down
		To create a new instance, click Up

Management

ID	Hostname	Status
i-49d33623	ec2-184-73-111-126.compute-1.amazonaws.com [ssh]	running Down
		To create a new instance, click Up

Done

S3Fox

PageRank: unrankable

Alexa

Compete

Instances: index

http://localhost:8080/admin

fried watermelon

AppDeployer

devs admins

Cluster Management

Shutdown All

Frontend

ID	Hostname	Status
i-75d3361f	ec2-184-72-134-102.compute-1.amazonaws.com [ssh]	running Down
To create a new instance, click		Up

App Management

Name	Status	URL
load-demo	running	http://ec2-184-72-134-102.compute-1.amazonaws.com/load-demo Undeploy

i-49d33623	ec2-184-73-111-126.compute-1.amazonaws.com [ssh]	running Down
To create a new instance, click		Up

Done

S3Fox PageRank: unrankable Alexa Compete

Scale

To scale a tier, administrator **simply launches more nodes.**

CirrAS appliances have been intentionally designed to be **automatically horizontally scalable.**

Instances: index

http://localhost:8080/admin

fried watermelon

AppDeployer

devs admins

Cluster Management

Shutdown All

Frontend

ID	Hostname	Status
i-75d3361f	ec2-184-72-134-102.compute-1.amazonaws.com [ssh]	running Down

Backend

ID	Hostname	Status
i-77d3361d	ec2-184-73-100-57.compute-1.amazonaws.com [ssh]	running Down
		To create a new instance, click Up

i-49d33623	ec2-184-73-111-126.compute-1.amazonaws.com [ssh]	running Down
		To create a new instance, click Up

Done

S3Fox PageRank: unrankable Alexa



Instances: index

http://localhost:8080/admin

fried watermelon

AppDeployer

devs admins

Cluster Management

Shutdown All

Frontend

ID	Hostname	Status
i-75d3361f	ec2-184-72-134-102.compute-1.amazonaws.com [ssh]	running Down

To create a new instance, click [Up](#)

Backend

ID	Hostname	Status
i-77d3361d	ec2-184-73-100-57.compute-1.amazonaws.com [ssh]	running Down
i-075ebb6d		pending Down

To create a new instance, click [Up](#)

Done

PageRank: unrankable Alexa Compete



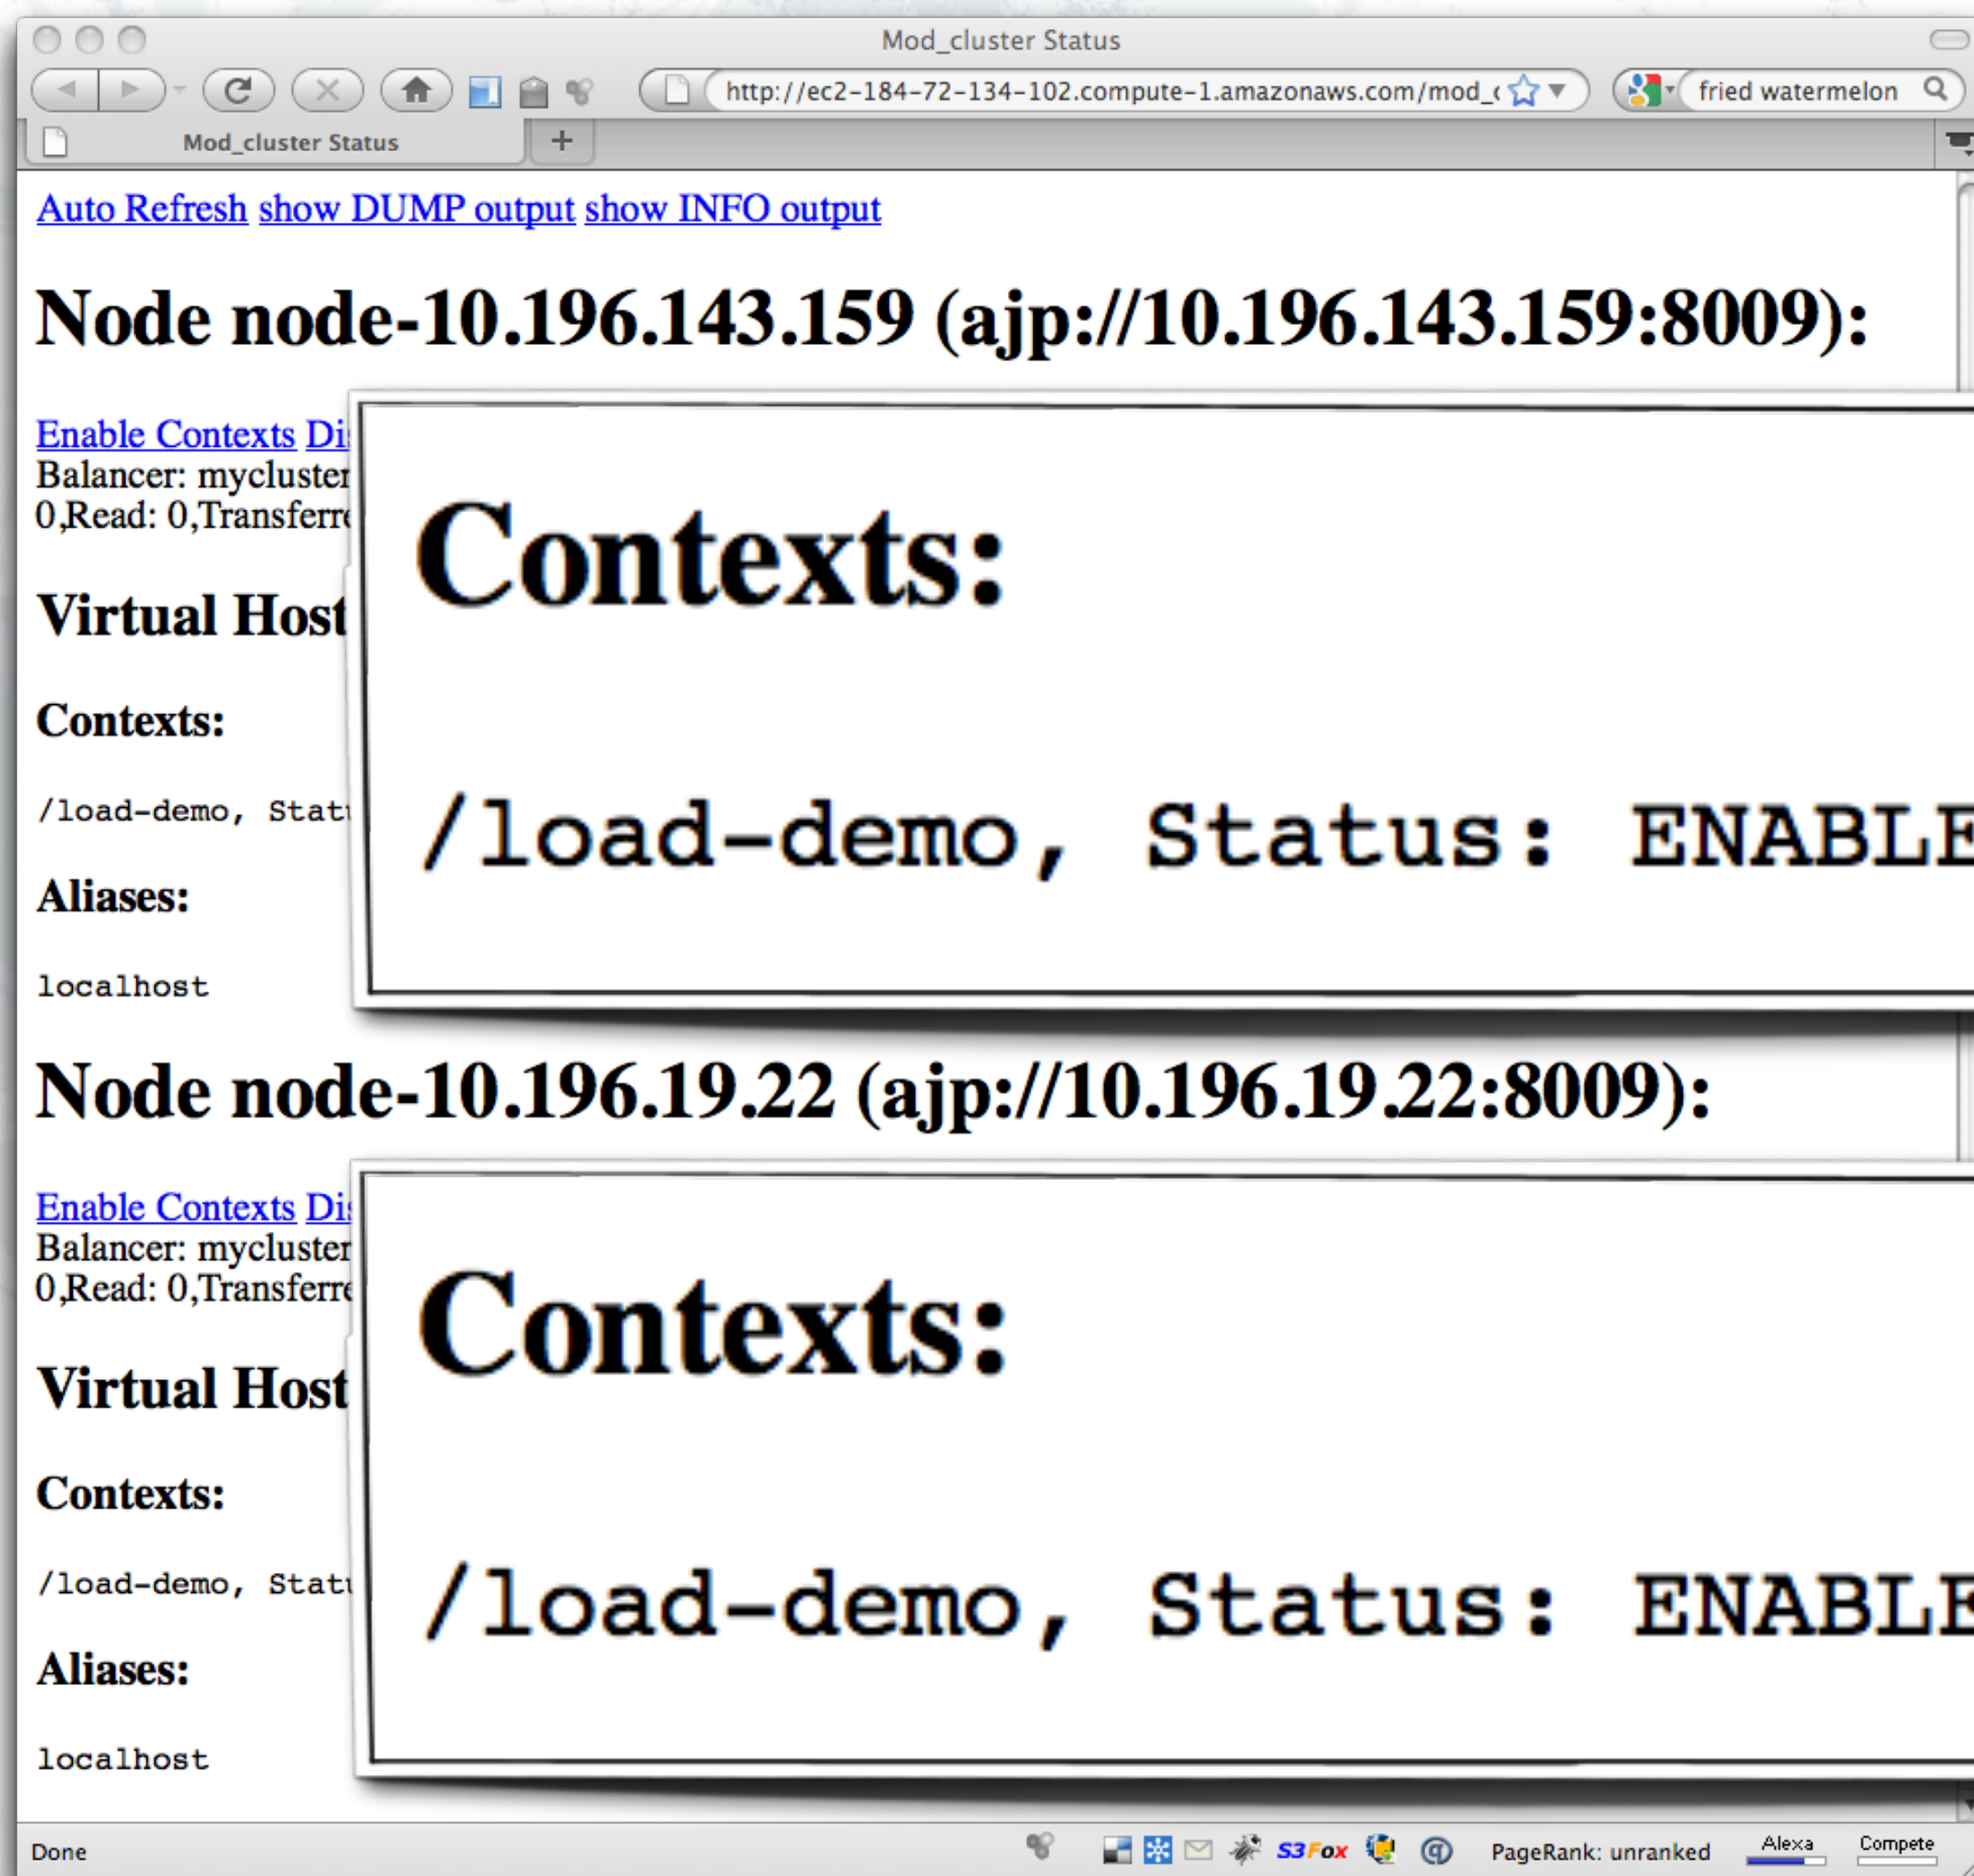
I love coffee.

Now would be a good time
to enjoy some.

What's going on?

While we enjoy our beverage...

- Another CirrAS node starts up.
- Management node notices.
- Shoots configuration (S3, EWS) to it.
- It joins the cluster.
- Neighbor hands the app to it.
- It becomes an available worker.



Oh yeah, JON

The screenshot shows the JON Dashboard web application running in a browser. The browser's address bar displays the URL `http://ec2-184-73-111-126.compute-1.amazonaws.com:7080/Das`. The application's navigation bar includes links for Overview, Resources, Groups, Administration, and Help, along with a 'Logged in as rhqadmin' status. The main content area is divided into several sections:









































- Dashboard**: Contains a 'Search Resources' widget with input fields for 'Resource Name' and 'Platforms', and a 'Saved Charts' widget showing 'No charts to display'.
- Summary Counts**: A table providing an overview of system metrics.

Metric	Count
Platform Total	4
Server Total	7
Service Total	304
Compatible Group Total	0
Mixed Group Total	0
Group Definition Total	0
Average Metrics per Minute	173
- Auto-Discovery**: A list of discovered resources, all with a status of 'COMMITTED'.

Resource Name	Status
ip-10-196-142-230.ec2.internal - Linux Operating System	COMMITTED
back-end-ip-10-196-143-159 - Linux Operating System	COMMITTED
back-end-ip-10-196-19-22 - Linux Operating System	COMMITTED
front-end-ip-10-196-142-180 - Linux Operating System	COMMITTED
- Recently Added Resources**: A table showing the most recently added resources.

Resource Name	Date / Time
back-end-ip-10-196-19-22	06/19/2010 10:32:05 PM
ip-10-196-142-230.ec2.internal	06/19/2010 03:12:08 PM
back-end-ip-10-196-143-159	06/19/2010 03:12:08 PM
front-end-ip-10-196-142-180	06/19/2010 03:12:07 PM
- Favorite Resources**: A section with an 'XML' link and the message 'No resources to display'.
- Recent Alerts**: A section with an 'XML' link and the message 'No recent alerts to display'.
- Operations**: Contains two sub-sections, 'Recent Operations' and 'Scheduled Operations', both showing 'No operations to display'.
- Has Alerts or Currently Unavailable**: A section with an 'XML' link and the message 'No resources to display'.

The browser's status bar at the bottom shows 'Done' and various utility icons, including PageRank (unranked) and Alexa (Complete).

	Server ▲ 1	Parent	Server Type	Plugin	Description	Availability
     	ip-10-196-142-180 Apache 2.2.14 (/etc/httpd/)	front-end-ip-10-196-142-180	Apache HTTP Server	Apache	Apache Web Server	✓
   	ip-10-196-142-230 Embedded JBossWeb Server 2.0.1.GA (0.0.0.0)	ip-10-196-142-230 RHQ Server, JBoss AS 4.2.3.GA default (0.0.0.0:2099)	Embedded Tomcat Server	JBossAS	JBossAS- Embedded JBossWeb Web Server (jboss- web.deployer/)	✓
      	ip-10-196-142-230 RHQ Server, JBoss AS 4.2.3.GA default (0.0.0.0:2099)	ip-10-196-142-230.ec2.internal	JBossAS Server	JBossAS	JBoss Application Server hosting the RHQ Server	✓
      	ip-10-196-143-159 JBoss EAP 5.0.1 cluster-ec2 (10.196.143.159:1099)	back-end-ip-10-196-143-159	JBossAS Server	JBossAS5	JBoss Enterprise Application Platform	✓
      	ip-10-196-19-22 JBoss EAP 5.0.1 cluster-ec2 (10.196.19.22:1099)	back-end-ip-10-196-19-22	JBossAS Server	JBossAS5	JBoss Enterprise Application Platform	✓
   	JBoss AS JVM	ip-10-196-142-230 RHQ Server, JBoss AS 4.2.3.GA default (0.0.0.0:2099)	JBoss AS JVM	JBossAS	JVM of the JBossAS	✓
   	JBoss Cache subsystem	ip-10-196-142-230 RHQ Server, JBoss AS 4.2.3.GA default (0.0.0.0:2099)	JBossCacheSubsystem	JBossCache	JBoss Cache subsystem	✓

Embedded JOPR

The screenshot displays the JBoss EAP Admin Console interface. The browser address bar shows the URL `http://ec2-184-73-100-57.compute-1.amazonaws.com:8080/admin-`. The console title is "JBoss EAP Admin Console". The left sidebar shows a tree view of the system structure, including "JBossAS Servers", "JBoss EAP 5 (cluster-ec2)", "Service Binding Manager", "JBoss Caches", "Applications", "Resources", and "JBoss Web". The main content area shows the configuration for `ip-10-196-143-159.ec2.internal`. The status is "Available". The "General Properties" section lists the Name, Version, and Description. The "Traits" section lists the Hostname, OS Name, OS Version, Architecture, Distribution Name, and Distribution Version.

JBoss EAP Admin Console

Welcome admin [Logout]

ip-10-196-143-159.ec2.internal

Summary Configuration Metrics Control Content Status: ✔ Available

General Properties

Name: ip-10-196-143-159.ec2.internal
Version: Linux 2.6.21.7-2.fc8xen
Description: Linux Operating System

Traits

Hostname: ip-10-196-143-159.ec2.internal
OS Name: Linux
OS Version: 2.6.21.7-2.fc8xen
Architecture: i386
Distribution Name: Red Hat Enterprise Linux Server
Distribution Version: release 5.5 (Tikanga)

JBoss EAP Admin Console 1.3.2.GA (r784) - Powered by [Embedded Jopr](#)
© 2002-2009 Red Hat Middleware, LLC. All rights reserved. JBoss is a registered trademark of Red Hat, Inc.

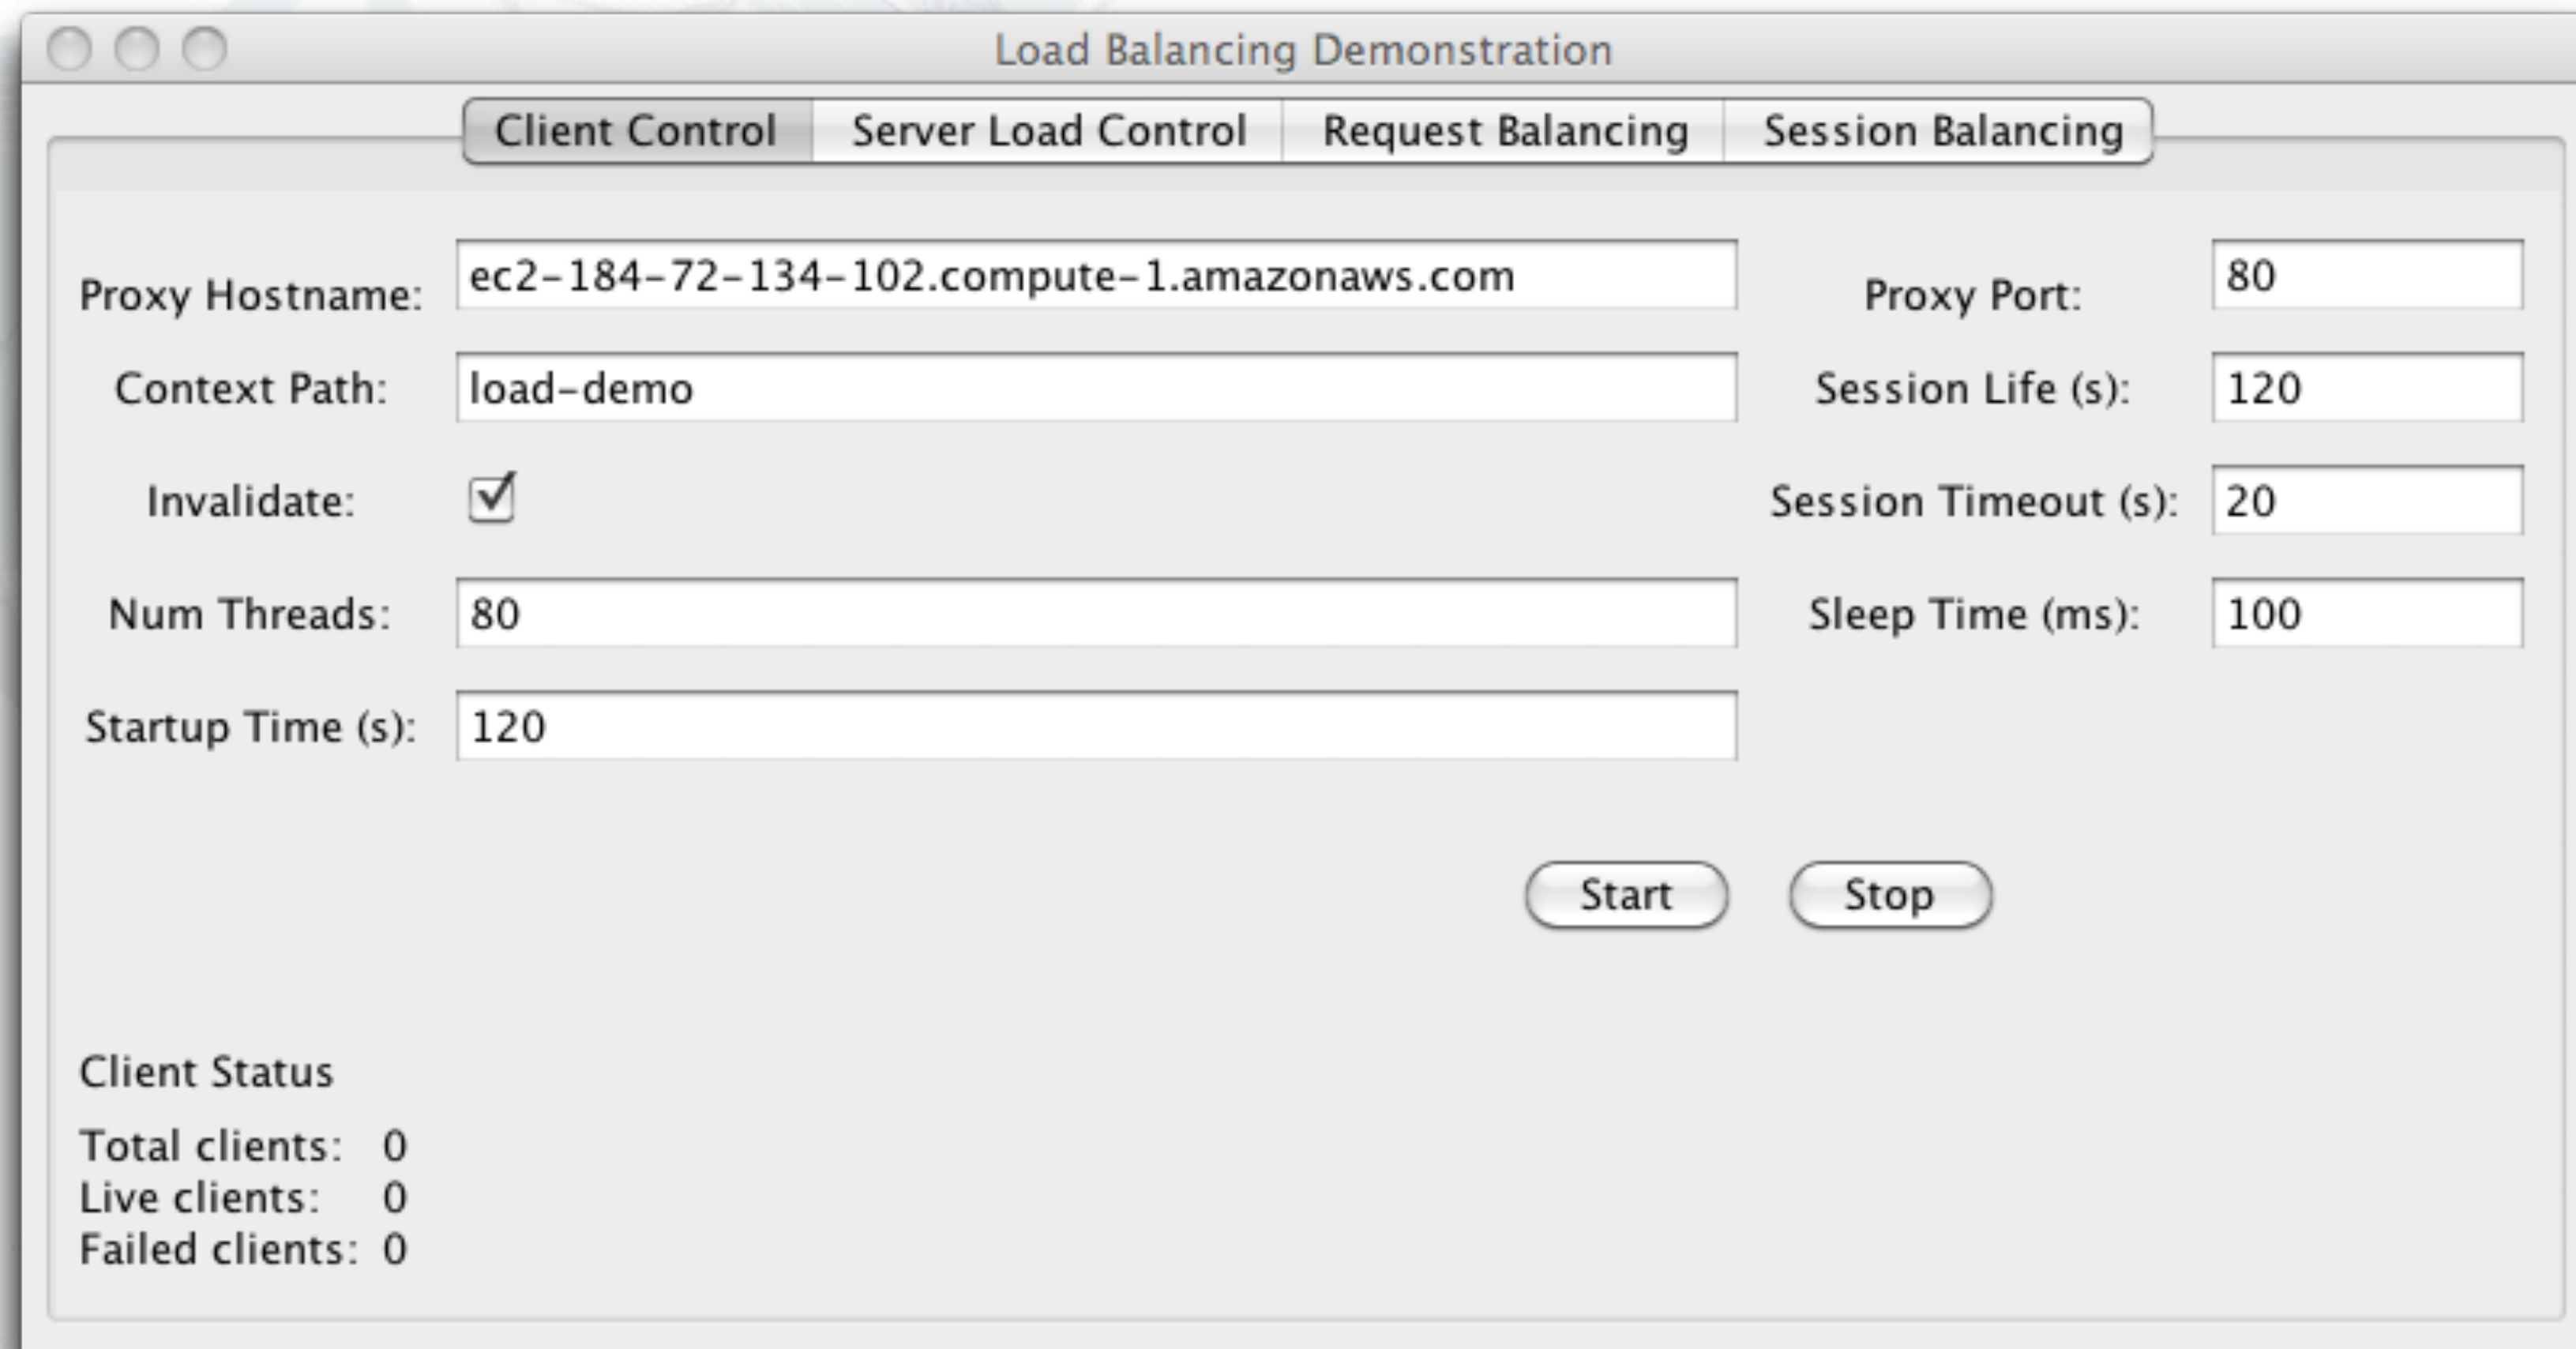
Done



Role #3

**“I’m a user, and I want to
use your app a whole lot.”**

Our user...



The screenshot shows a Java Swing window titled "Load Balancing Demonstration". It has four tabs: "Client Control" (selected), "Server Load Control", "Request Balancing", and "Session Balancing". The "Client Control" tab contains several input fields and checkboxes. The "Proxy Hostname" field is set to "ec2-184-72-134-102.compute-1.amazonaws.com". The "Proxy Port" field is set to "80". The "Context Path" field is set to "load-demo". The "Invalidate" checkbox is checked. The "Num Threads" field is set to "80". The "Startup Time (s)" field is set to "120". The "Session Life (s)" field is set to "120". The "Session Timeout (s)" field is set to "20". The "Sleep Time (ms)" field is set to "100". At the bottom right of the tab are "Start" and "Stop" buttons. At the bottom left, under the heading "Client Status", are three lines of text: "Total clients: 0", "Live clients: 0", and "Failed clients: 0".

Load Balancing Demonstration

Client Control Server Load Control Request Balancing Session Balancing

Proxy Hostname: ec2-184-72-134-102.compute-1.amazonaws.com Proxy Port: 80

Context Path: load-demo Session Life (s): 120

Invalidate: ☒ Session Timeout (s): 20

Num Threads: 80 Sleep Time (ms): 100

Startup Time (s): 120


Start Stop

Client Status

Total clients: 0

Live clients: 0

Failed clients: 0



Load Balancing Demonstration

Client Control Server Load Control Request Balancing Session Balancing


Target Hostname: Target Port:

Load Creation Action:

Duration:


Ratio:

Create Load






**Take a node
out of service.**





Scale **more!**



For less than the cost of a cup of coffee...

1	EWS	\$0.08
5	EAP	\$0.40
1	Management	\$0.08
<hr/>		\$0.56

Roadmap

- Application versions
- Multiple clusters/environments
- Per-application configuration
- ELB/HAProxy support



Thanks!



Q&A

Resources

<http://jboss.org/stormgrind/>
<http://cloudpress.org/>

Projects

Blog

#stormgrind

IRC

@stormgrind

@marekgoldmann

@jcrossley3

@bobmcwhirter

Twitter