



GETTING STARTED WITH COLLECTIVEACCESS ON OPENSIFT FLEX

DEPLOY, MONITOR AND AUTO-SCALE JAVA EE AND PHP APPLICATIONS
IN THE CLOUD

A RED HAT EVALUATION GUIDE

BRAM WIERCX | FARO. FLEMISH INTERFACE CENTRE FOR CULTURAL HERITAGE

DECEMBER 2011



WHAT IS OPENSIFT FLEX?

OpenShift Flex is a free, cloud-based application platform for Java EE and PHP applications delivered in a dedicated hosting model. OpenShift Flex is an intelligent auto-scaling solution for applications in the cloud that doesn't require templates to be tweaked and scripts to be written or maintained. With this guide you'll get familiar with the OpenShift Flex platform, how it works, how to deploy CollectiveAccess, plus how to take advantage of the management, monitoring, log analysis and auto-scaling capabilities of the platform.

Please note that although the OpenShift Flex platform is free to use, you will incur charges from your cloud provider for the underlying resources your application consumes.

ASSUMPTIONS

For the purposes of this guide, we assume that you have already created an OpenShift account and have been granted access to OpenShift Flex. It is also assumed that you have already set up your cloud account in OpenShift Flex as well as, created a server cluster on which to deploy CollectiveAccess. If you are brand new to OpenShift Flex, we suggest you get familiar with the platform by reviewing the OpenShift Flex **Evaluation Guide** at: <http://www.redhat.com/openshift/documents>.

STEP 1: CREATE A CLUSTER

- Click on the **Cluster** button in OpenShift Flex's navigation bar to bring up the **Cluster** page
- Click on **Add Cluster** button on the right-hand side



DEFINE A SERVER CLUSTER

CLUSTER NAME	<input type="text" value="faronet"/>	INITIAL SERVERS	<input type="text" value="1"/>	MIN	<input type="text" value="1"/>	MAX	<input type="text" value="2"/>		PROJECTED AVG. 1
CLOUD PROVIDER ACCOUNT	<input type="text" value="faro (bram@faronet.be)"/>	SERVER CLUSTER LOCATION	<input type="text" value="Europe West"/>						
PROVIDER INSTANCE	<input type="text" value="Small"/>	\$0.095 /SERVER HOUR	INSTANCE COST \$69 /MONTH						
32 BIT 1 CORE(S) 1.66 GB									
DISK SIZE	<input type="text" value="25"/> GB	<input type="checkbox"/> LOAD BALANCER	< \$0.01 /HOUR	USAGE COST \$2 /MONTH					
ADMIN PASSWORD	<input type="password" value="*****"/>							PROJECTED COST \$72 /MONTH (1 SERVER CLUSTER)	
<small>The admin password is used for SSH access and communication from this console to your cloud servers. Your cloud servers will be exposed to the public internet so please choose a secure password that you can remember.</small>									
CONFIRM PASSWORD	<input type="password" value="*****"/>							APPROX. COST IN US\$	
<input type="button" value="CANCEL"/>				<input type="button" value="SAVE"/>					

- Define a **Server Cluster**
- **Save**

STEP 2: SELECT COLLECTIVEACCESS OPENSIFT APPLICATION

- Click on the **Applications** button in OpenShift Flex's navigation bar to bring up the **Applications** page
- Click on **Add Application** button on the right-hand side

WELCOME workshopt@redhat.com [LOGOUT](#)

OPENSIFT™ PaaS by Red Hat Cloud
INTRO CLOUDS CLUSTERS SERVERS **APPLICATIONS** PERFORMANCE LOGS EVENTS

APPLICATION LIST

▼ FILTER BY CLUSTERS You don't have any applications

	NAME (click to configure)	CLOUD	CLUSTE	VERSION	DEPLOYMEN	STATUS	ACTIONS
qe's:jimmygtes							

- Select **Server Cluster**
- Select **Deploy a Sample Application**
- Select **Community**
- Select one of the two **CollectiveAccess** packages:



- 1) **CollectiveAccess Providence** core cataloging and data management application
- 2) **CollectiveAccess Providence added with Pawtucket** public web-access application.

NEW APPLICATION

SERVER CLUSTER
faro:faronet

Import Choose File i
 Create New
 Deploy a Sample Application


We've provided some sample applications for you to get started. Deploying an application can take several minutes so please be patient. Each sample application comes with a load generator that will also be deployed so you can check out performance monitoring and log files. Feel free to poke around under the hood or modify these as a starting point for your own applications.

	NAME	VERSION	DESCRIPTION
<input checked="" type="radio"/>	collectiveaccess	1.0	CollectiveAccess Providence core cataloging and data management application
<input type="radio"/>	drupal	7.8	Drupal LAMPHP sample applicaton. Login is admin/4dm1n
<input type="radio"/>	emerging	1.0	MRG-Messaging "shell" application (for direct access) with port 5672

- You should now see your application is downloading, importing and deploying, this can take some time (+/- 10 minutes)

FILTER BY CLUSTERS		Select an application.					
NAME (click to configure)	CLOUD	CLUSTER	VERSION	DEPLOYMENT	STATUS	ACTIONS	
collectiveaccess	fero	faronet	1.0	0	Deploying	Deploying	

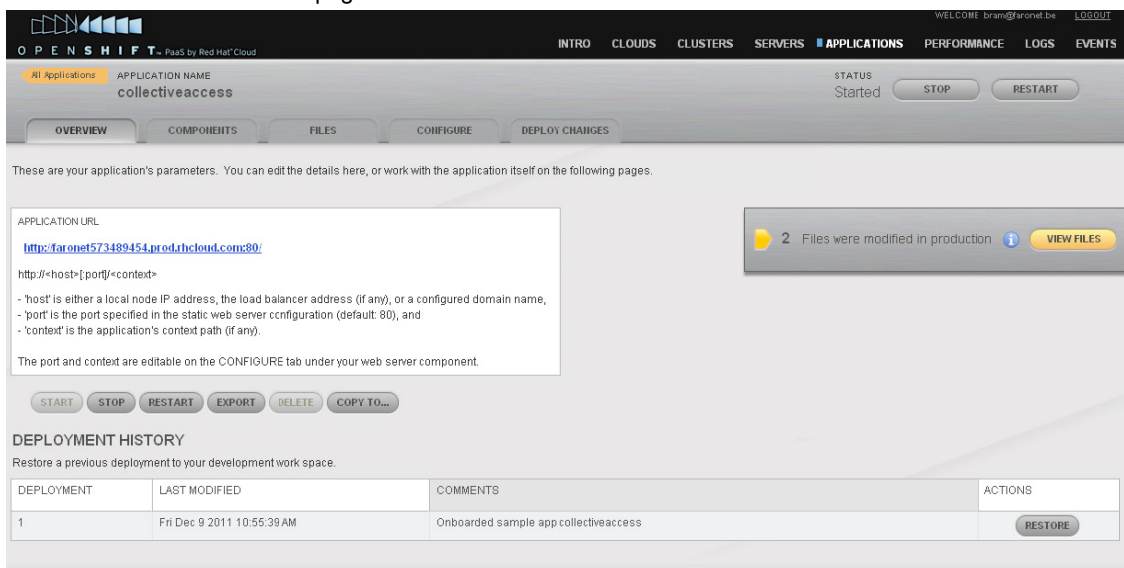
- You should now see your application (after importing and deploying) in the **Applications** page



STEP 3: START THE COLLECTIVEACCESS APPLICATION

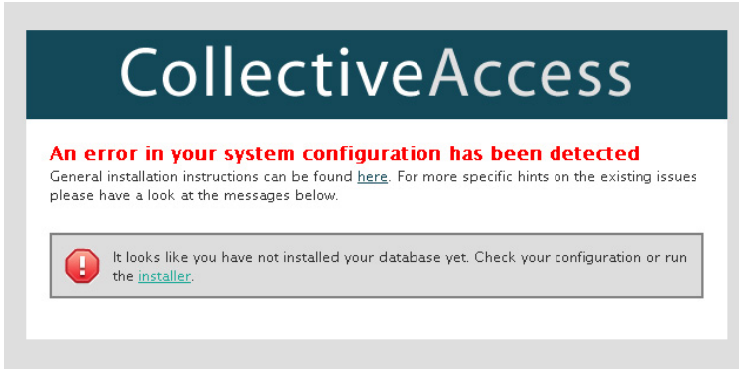
Now that CollectiveAccess is deployed, you can **Start**, **Stop**, **Restart**, **Export**, **Delete** or **Rollback** to a previous version from the **Applications** page. If you aren't already there, click on the **Overview** tab of the **Applications** page.

- Click the **application url** to finish the installation and configuration of CollectiveAccess by navigating to the CollectiveAccess home page





- Start now with the CollectiveAccess installation <http://faronet573489454.prod.rhcloud.com/install>



- For CollectiveAccess Pawtucket installations, Pawtucket is installed in the subdirectory /Pawtucket : <http://faronet573489454.prod.rhcloud.com/pawtucket>



- More information and support : <http://www.collectiveaccess.org> . Do not forget to **make backups**: http://collectiveaccess.org/docs/Backup_Guide_v1.pdf and snapshots of our data on amazon EBS instance: <http://aws.amazon.com/ebs/> (the data is not included in a openshift export). You can also make backups of our database in OpenShift.

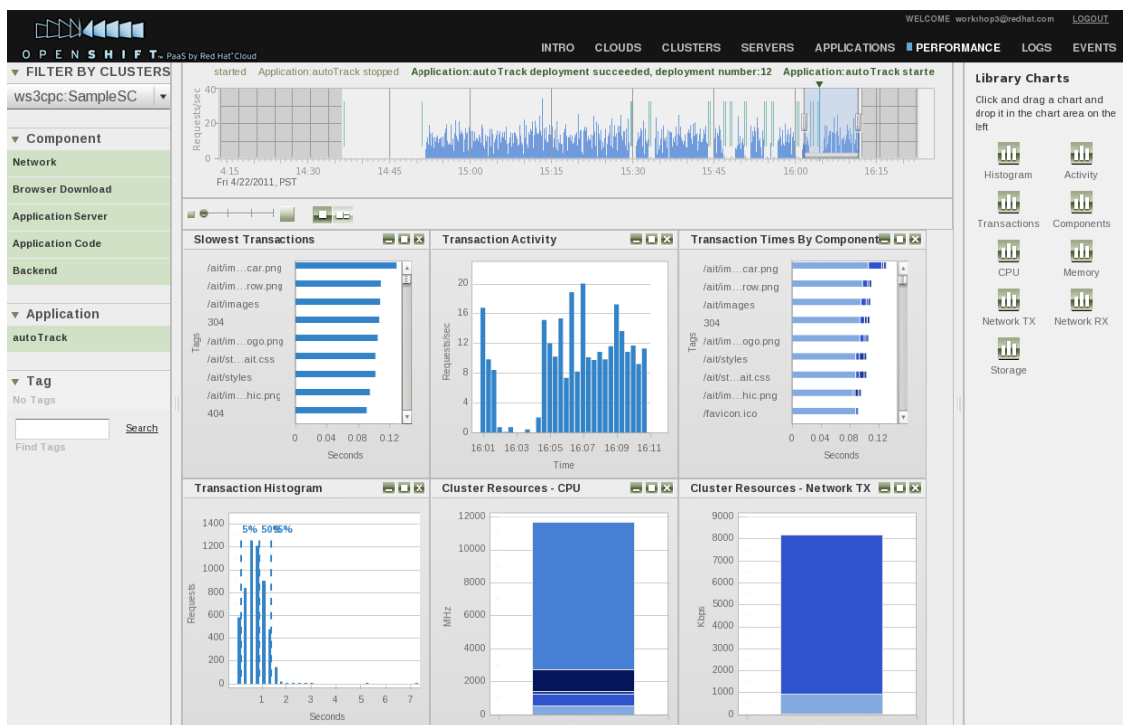


STEP 4: PERFORMANCE MONITORING

OpenShift Flex has performance monitoring configured straight out-of-the-box. To get started, click on the **Performance** tab to bring up the performance monitoring page. On this page you will find a variety of performance charts to view both historically and in real-time.

PERFORMANCE CHARTS

- You can add additional charts by dragging the icons found in the **Library Charts** section onto the main page.
- You can choose to include or exclude certain components or application from the performance views by selecting or deselecting them in the left hand-section of the page.




HISTORICAL PERFORMANCE

- To view performance data from a particular point in time, move the range selecting slider in the historical view.



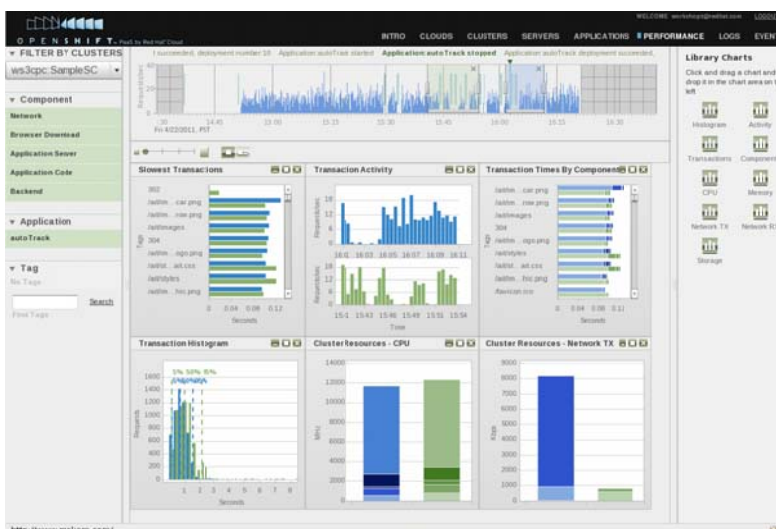
- The data in the performance charts will change to reflect the currently selected time slice.
- You will also notice that important system events like application changes and restarts are watermarked above the timeline.
- Daily performance data metrics are currently set to be archived for 2000 days, with 1 minute granularity being aggregated after 14 days, 15 minute granularity id aggregated after 100 days and 1 hour granularity is aggregated after 400 days



COMPARING HISTORICAL PERFORMANCE

You can also compare the performance data of two separate time slices.

- Click on the current selected time slice.
- This will produce another time slice selector labeled **2** on the left-hand side of the timeline.
- Click this icon to select the time slice you want to compare.
- You will now see the performance charts below indicating the respective metrics for each time slice being compared.

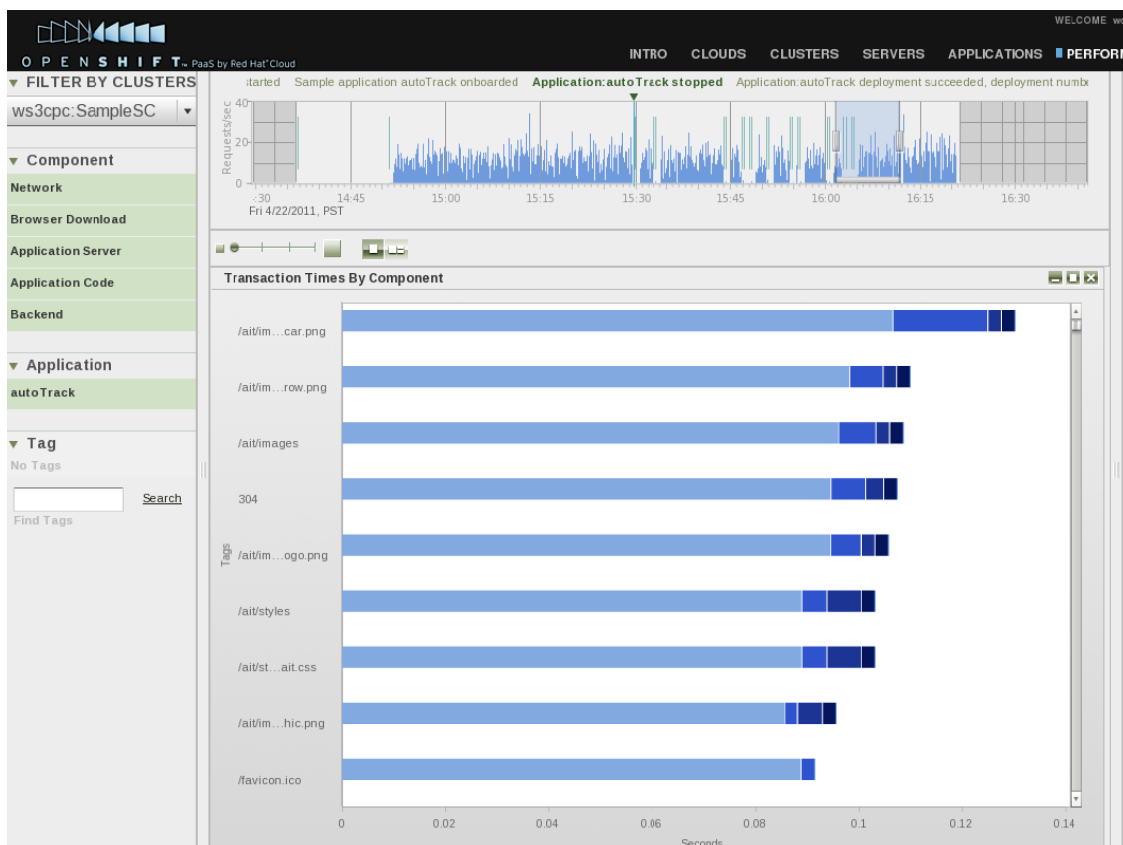




PERFORMANCE CHART DRILL DOWN

The performance charts also offer the ability to drill down to get more granular data about response times and transactions. For example: By drilling down into the Transactions Times By Component chart, you can identify which transactions are taking the longest to complete.

- Hovering over the various segments in the bar graphs allows you to identify how long the transaction spent in each software component.
- By clicking on the **Backend** portion of the graph, you can drill down even further to see exactly what SQL statements were being executed and the performance metrics associated with them.

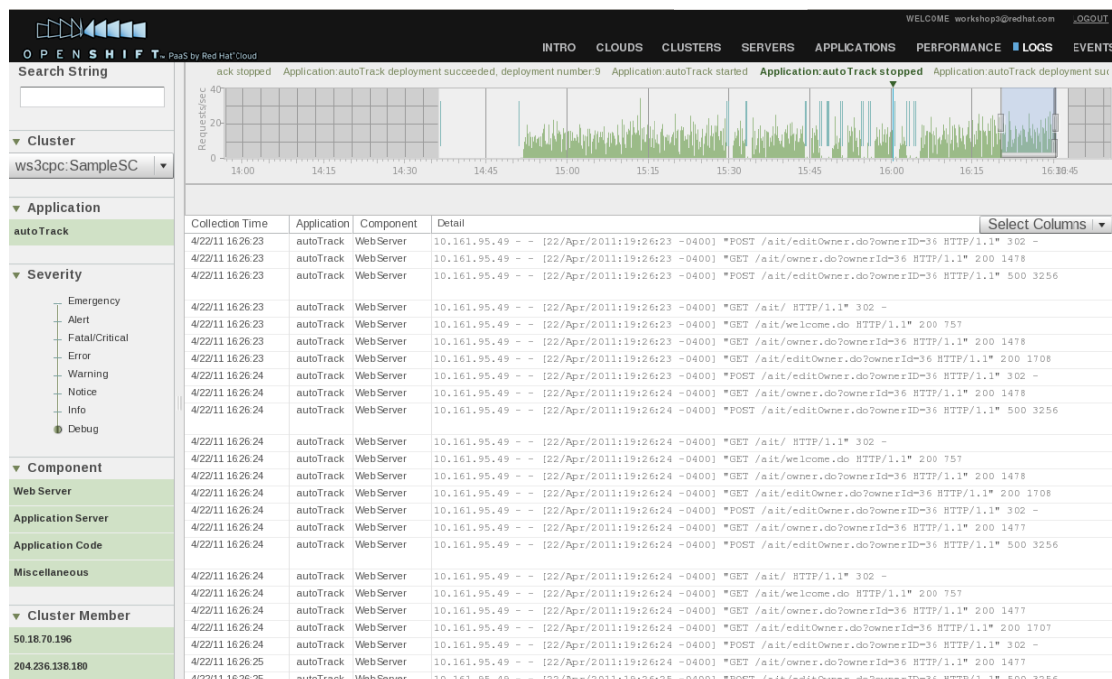




STEP 5: LOG ANALYSIS

OpenShift Flex automatically aggregates and persists relevant logs across software components for analysis.

- Click the **Logs** tab
- From this page you filter logs by a specific string, application, severity, component and cloud server.
- As you saw in the Performance page, you can use the historical view to select the range of historical data you are interested in analyzing.



Collection Time	Application	Component	Detail
4/22/11 16:26:23	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:23 -0400] "POST /sit/editOwner.do?ownerId=36 HTTP/1.1" 302 -
4/22/11 16:26:23	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:23 -0400] "GET /sit/owner.do?ownerId=36 HTTP/1.1" 200 1478
4/22/11 16:26:23	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:23 -0400] "POST /sit/editOwner.do?ownerId=36 HTTP/1.1" 500 3256
4/22/11 16:26:23	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:23 -0400] "GET /sit/ HTTP/1.1" 302 -
4/22/11 16:26:23	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:23 -0400] "GET /sit/welcome.do HTTP/1.1" 200 757
4/22/11 16:26:23	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:23 -0400] "GET /sit/owner.do?ownerId=36 HTTP/1.1" 200 1478
4/22/11 16:26:23	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:23 -0400] "GET /sit/editOwner.do?ownerId=36 HTTP/1.1" 200 1708
4/22/11 16:26:24	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:23 -0400] "POST /sit/editOwner.do?ownerId=36 HTTP/1.1" 302 -
4/22/11 16:26:24	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:24 -0400] "GET /sit/owner.do?ownerId=36 HTTP/1.1" 200 1478
4/22/11 16:26:24	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:24 -0400] "POST /sit/editOwner.do?ownerId=36 HTTP/1.1" 500 3256
4/22/11 16:26:24	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:24 -0400] "GET /sit/ HTTP/1.1" 302 -
4/22/11 16:26:24	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:24 -0400] "GET /sit/welcome.do HTTP/1.1" 200 757
4/22/11 16:26:24	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:24 -0400] "GET /sit/owner.do?ownerId=36 HTTP/1.1" 200 1478
4/22/11 16:26:24	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:24 -0400] "GET /sit/editOwner.do?ownerId=36 HTTP/1.1" 200 1708
4/22/11 16:26:24	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:24 -0400] "POST /sit/editOwner.do?ownerId=36 HTTP/1.1" 302 -
4/22/11 16:26:24	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:24 -0400] "GET /sit/owner.do?ownerId=36 HTTP/1.1" 200 1477
4/22/11 16:26:24	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:24 -0400] "POST /sit/editOwner.do?ownerId=36 HTTP/1.1" 500 3256
4/22/11 16:26:24	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:24 -0400] "GET /sit/ HTTP/1.1" 302 -
4/22/11 16:26:24	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:24 -0400] "GET /sit/welcome.do HTTP/1.1" 200 757
4/22/11 16:26:24	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:24 -0400] "GET /sit/owner.do?ownerId=36 HTTP/1.1" 200 1477
4/22/11 16:26:24	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:24 -0400] "GET /sit/editOwner.do?ownerId=36 HTTP/1.1" 200 1707
4/22/11 16:26:24	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:24 -0400] "POST /sit/editOwner.do?ownerId=36 HTTP/1.1" 302 -
4/22/11 16:26:25	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:24 -0400] "GET /sit/owner.do?ownerId=36 HTTP/1.1" 200 1477
4/22/11 16:26:25	autoTrack	Web Server	10.161.95.49 -- [22/Apr/2011:19:26:25 -0400] "POST /sit/editOwner.do?ownerId=36 HTTP/1.1" 500 3256

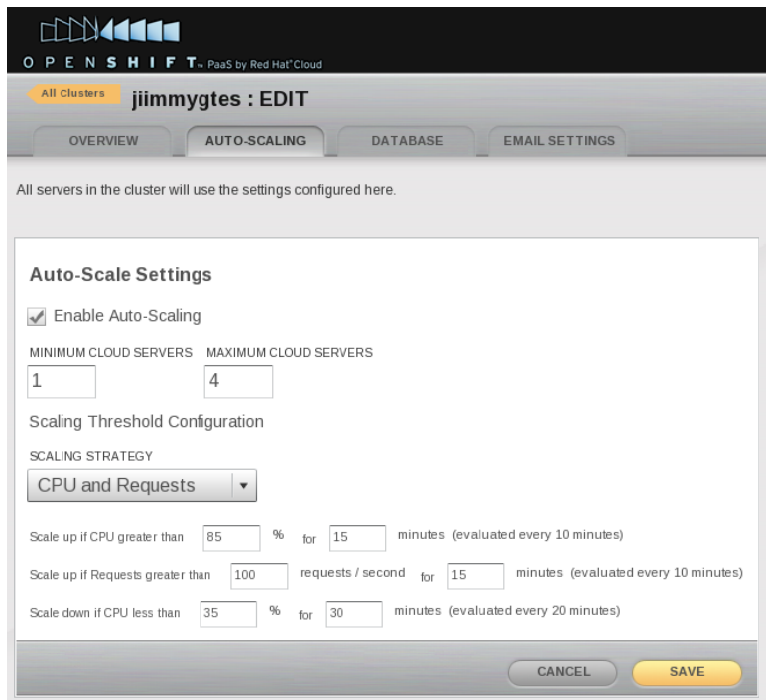
STEP 6: AUTO-SCALE YOUR APPLICATION

To configure your application for auto-scaling up and down please do the following:

- Click on the **Clusters** tab
- Click the **Name** of the **Server Cluster** you are interested in configuring for auto-scaling
- Click on the **Auto-Scaling** tab



- Check the **Enable Auto-Scaling** box
- Configure your desired auto-scaling rules by selecting the following parameters
 - Minimum cloud servers
 - Maximum cloud servers
 - Scaling strategy
 - Scale up if CPU greater than
 - Scale up if requests greater than
 - Scale down if CPU less than
- Click **Save** when you are finished



The screenshot shows the 'Auto-Scaling Settings' page in the OpenShift console. The page title is 'jjimmygtes : EDIT' and it has tabs for 'OVERVIEW', 'AUTO-SCALING', 'DATABASE', and 'EMAIL SETTINGS'. The 'AUTO-SCALING' tab is selected. Below the tabs, there is a note: 'All servers in the cluster will use the settings configured here.' The main content area is titled 'Auto-Scale Settings' and contains the following configuration options:

- Enable Auto-Scaling
- MINIMUM CLOUD SERVERS:
- MAXIMUM CLOUD SERVERS:
- Scaling Threshold Configuration
- SCALING STRATEGY:
- Scale up if CPU greater than % for minutes (evaluated every 10 minutes)
- Scale up if Requests greater than requests / second for minutes (evaluated every 10 minutes)
- Scale down if CPU less than % for minutes (evaluated every 20 minutes)

At the bottom of the form, there are two buttons: 'CANCEL' and 'SAVE'.

WHAT'S NEXT?

At this point you should have a basic understanding of how to deploy, manage, monitor and scale an application with OpenShift Flex. If you still have questions about OpenShift Flex, please visit our forum or ping us on IRC. Additional resources such as knowledge base articles, application deployment guides, and how-to videos can be found on the OpenShift Flex product page located at:

<http://openshift.redhat.com/app/flex>



TERMS AND CONDITIONS

Please be advised of the following terms and conditions related to OpenShift Flex.

- This offering does not come with support, upgrades, or warranties.
- Any unauthorized use of OpenShift Flex is strictly prohibited.
- In order to receive OpenShift Flex, you will be required to accept the terms and conditions.
- This offering currently is not suitable for production use, so it is advised to not store sensitive or confidential data.

To view the terms and conditions document, please visit:

http://openshift.redhat.com/app/legal/services_agreement

RESOURCES

Forum

<http://www.redhat.com/openshift/forums/flex/>

<http://www.collectiveaccess.org/forum/>

IRC Channel

irc.freenode.net

#openshift

How-To Videos

<http://www.redhat.com/openshift/videos>

Documentation

http://docs.redhat.com/docs/en-US/OpenShift/1.0/html/OpenShift_Flex_User_Guide/index.html

<http://wiki.collectiveaccess.org/>

<http://www.collectiveaccess.org/support/documentation>

Knowledge Base

<http://www.redhat.com/openshift/kb>

Evaluation Guides

<http://www.redhat.com/openshift/documents>



RED HAT SALES AND INQUIRIES

NORTH AMERICA
 1-888-REDHAT1
www.redhat.com
sales@redhat.com

EUROPE, MIDDLE EAST
 AND AFRICA
 00800 7334 2835
www.europe.redhat.com
europe@redhat.com

ASIA PACIFIC
 +65 6490 4200
www.apac.redhat.com
apac@redhat.com

LATIN AMERICA
 +54 11 4341 6200
www.latam.redhat.com
info-latam@redhat.com