



# Red Hat:

**The “Top 5 Things” Customers need to know in order to be successful with Red Hat Infrastructure and Technologies**

Chris Henderson  
Technical Account Manager  
15 May 2013

# Red Hat

## From a TAM's Perspective

### Top 5 Things:

- RHEL Lifecycle Understanding And Customer Lifecycle Alignment
- Centralization, Core Builds, And Componentization
- Infrastructure Testing
- Important Configuration For Troubleshooting Issues (RCA)
- Customer Portal Value

# RHEL Lifecycle Understanding And Customer Lifecycle Alignment

# RHEL Lifecycle Understanding And Customer Lifecycle Alignment

- **You need to understand RHEL release schedules**

- Major RHEL releases come out approximately every 2-3 years
- Minor RHEL releases come out approximately every 6 months
- Z-stream Releases come out approximately every 6 weeks
- Important Links

<https://access.redhat.com/support/policy/updates/errata/>

[https://access.redhat.com/support/policy/update\\_policies.html](https://access.redhat.com/support/policy/update_policies.html)

- **EUS (Extended Update Support)**

- Allows you to remain on a minor RHEL Release
- Reduces your testing requirements
- Aligns to ISVs supportability matrices
- Once the next minor RHEL release comes out the previous minor release would now become EUS
- Requires Add-on Entitlement

# RHEL Lifecycle Understanding And Customer Lifecycle Alignment

- **RFE (Request For Enhancement)**
  - This is important to you and us, it makes RHEL and other RH products better
  - If you see something that could be added or improved upon create a RFE bugzilla
- **Hot Fix Requests**
  - If you have found a bug and it's critical to your business, you can request an accelerated hotfix that will come out in a z-stream release
  - Faster than waiting on a minor release
- **ELS (Extended Lifecycle)**
  - Allows customer to run on a major RHEL release much longer
  - Keeps customers in check with security fixes
  - Add-on
- **Red Hat vulnerabilities**
  - CVE Database  
<https://access.redhat.com/security/cve/>

# RHEL Lifecycle Understanding And Customer Lifecycle Alignment

- **Customer Lifecycle Hardware Perspective**
  - Typically driven by hardware refresh 3-5 years
  - Times are moving faster these days cpu lifecycle has sped up
    - About every 24 months
  - Hardware ISVs fall in line with cpu lifecycle
    - At some point they will stop supporting older Operating Systems  
e.g. HP G8 and RHEL4
- **Customer Lifecycle Application Perspective**
  - Holiday Season rules the Retail Customer's lifecycles
  - Business/Application deployments weigh in heavily

# RHEL Lifecycle Understanding And Customer Lifecycle Alignment

- **Take All Of These Lifecycles Into Account To Align Appropriately**
  - Specific to RHEL it allows you to position errata/patching goals before application testing starts
  - Specific to RHEL it allows you to know when potential errata will fall if application deployment/testing is delayed, might allow for another round of patching
- **Why is alignment so critical with RHEL releases/lifecycles**
  - As mentioned above hardware cycles are much faster these days which means there will inevitably be OS bugs
  - Keeping a close eye where RHEL is with patches in regards to your own testing and release cycles can help reduce downtime when production changes are finally rolled out

# Centralization, Core Builds, And Componentization



# Centralization, Core Builds, And Componentization

- **Use A Centralized Build and Management System**
  - Satellite Server (it's really not that expensive)
    - Cobbler is integrated
  - New RH Systems Management Coming Soon
  - Large Companies Work In Silos
    - Repetitive work being done
      - e.g. Project X admins write a script to install websphere app server
      - Project Y admins write a script to install websphere app server
    - Using extra storage space for the same stuff
    - Failure to Collaborate
      - e.g. Oh, cool you are using aide to monitor for intrusion detection on linux
      - and I get to see how you configured it too, awesome

# Centralization, Core Builds, And Componentization

- **Use A Centralized Build and Management System**
  - Reduces admin footprint, if you do things the “Right Way”
    - Cost Reduction
    - Can also be seen as giving admins time to do more proactive things that are often overlooked
      - Troubleshooting configuration
      - Monitoring configuration
  - Eases the ability to make server builds homogeneous
    - Is my hardware firmware the same across all my XYZ servers?
    - Are all my rpm versions the same?
    - Are all my configuration files the same?
    - Do I have the same set of services running on all my servers?
    - Do I have the same number of rpms installed?
    - Have I applied all available patches?
    - How am I validating all of this?

# Centralization, Core Builds, And Componentization

- **Core Builds**
  - It all starts with a kickstart
    - Ksmeta variables give me flexibility
      - Dynamic filesystem partitioning schemes
      - Tells me what additional components to install
    - Templates are great in the virtualization world
      - But build them using kickstart
      - Makes it easy to create a new template
        - Maybe you need a new partitioning scheme
      - Don't bake in additional components
        - Like our websphere example, installing a new version is easy using a base template and a new websphere install script

# Centralization, Core Builds, And Componentization

- **Core Builds**
  - It all starts with a kickstart
    - Limit your rpm install base to as few packages needed as necessary
      - Reduces unneeded errata
      - Reduces services that you don't need running
      - Reduces your security attack surface
      - Normally this rpm install base can and should support 99% of your applications, make it generic

# Centralization, Core Builds, And Componentization

- **Componentization**
  - Kickstart %posts
    - Use separate scripts for each component installed and/or configured
      - java configuration
      - ntp configuration
      - sysctl configuration
      - syslog configuration
      - serial console configuration
      - kdump configuration
      - limits.conf configuration
      - sysstat configuration
      - postfix/sendmail configuration
      - websphere/jboss/weblogic
      - snmp
      - etc.

# Centralization, Core Builds, And Componentization

- **Componentization**

- Let's look at kdump configuration
  - For a vm I may not want to do local dumps and instead configure kdump to dump via net, so if I separate this out, it makes it easy for me to choose just the right one for the job
- I want to move forward with testing a new Java version on my QA builds but maintain my existing production builds all using the same kickstart
  - Again if I have this script separated out, I can pick and choose the right one based on ksmeta
- Basically componentization allows me to easily replace single parts of my build keeping the rest in tact

# Infrastructure Testing

# Infrastructure Testing

- **100% Dependence On ISVs Is Bad**
  - What!! I can only get 30MB/s read performance with 6Gb/s sas backplane
  - What!! There was a bnx2 regression that causes panics with high load
- **100% Dependence On Application Testing Is Bad**
- **We rush to get the OS configured to get the application loaded**
  - Now it's too late, we've made the hand off
  - Application tests don't always have good simulators to real life production work loads, so we hit problems when we go to production
- **Stress Test Your Infrastructure**
  - What's your read/write speeds
    - Server internal array
    - Shared Storage (NAS/SAN)
  - What's your max throughput
  - What's your limiting factor (good for capacity and scalability planning)



# Infrastructure Testing

- **Stress Test Your Infrastructure**
  - What's your IP network throughput/latency numbers?
  - Know your production workload and align your stress test accordingly
    - Is it random or sequential IO?
    - Is it dominate reads or writes?
    - What is the typical read write size?
  - Did you play around with different raid types in your testing?
  - Did you play around with block sizes on the array size?
  - Are you partitioned aligned to the array block size?
  - Will your filesystem have many small files or fewer large ones
  - Are you tcp tuned for the amount of traffic you'll need to support?
  - Can you break it, panic a box, reach a limit or hang a box?

# Infrastructure Testing

- **Stress Testing Tools**

- iotzone (filesystem)
- stress (cpu, memory, i/o, disk)  
<http://weather.ou.edu/~apw/projects/stress/>
- netperf (network throughput)
- iperf (network throughput)
- EMC analyzer for monitoring
- sysstat/sar for monitoring

# Infrastructure Testing

- **Stress Testing iotest example (see man page for options)**

```
#!/bin/bash
```

```
mkdir -p /netbackup/vital_db_dssu/tmp
```

```
testnum=$1
```

```
right_now=`date +"%Y%m%d%H%M%S"``
```

```
iotest -Rb iotest-$right_now-$(hostname)-vxfstest$testnum.xls -l -s2G -r256 -t 32 -i0 -F  
/netbackup/vital_db_dssu/tmp/iotest00 /netbackup/vital_db_dssu/tmp/iotest01  
/netbackup/vital_db_dssu/tmp/iotest02 /netbackup/vital_db_dssu/tmp/iotest03  
/netbackup/vital_db_dssu/tmp/iotest04 /netbackup/vital_db_dssu/tmp/iotest05  
/netbackup/vital_db_dssu/tmp/iotest06 /netbackup/vital_db_dssu/tmp/iotest07  
/netbackup/vital_db_dssu/tmp/iotest08 /netbackup/vital_db_dssu/tmp/iotest09  
/netbackup/vital_db_dssu/tmp/iotest10 /netbackup/vital_db_dssu/tmp/iotest11  
/netbackup/vital_db_dssu/tmp/iotest12 /netbackup/vital_db_dssu/tmp/iotest13  
/netbackup/vital_db_dssu/tmp/iotest14 /netbackup/vital_db_dssu/tmp/iotest15  
/netbackup/vital_db_dssu/tmp/iotest16 /netbackup/vital_db_dssu/tmp/iotest17  
/netbackup/vital_db_dssu/tmp/iotest18 /netbackup/vital_db_dssu/tmp/iotest19  
/netbackup/vital_db_dssu/tmp/iotest20 /netbackup/vital_db_dssu/tmp/iotest21  
/netbackup/vital_db_dssu/tmp/iotest22 /netbackup/vital_db_dssu/tmp/iotest23  
/netbackup/vital_db_dssu/tmp/iotest24 /netbackup/vital_db_dssu/tmp/iotest25  
/netbackup/vital_db_dssu/tmp/iotest26 /netbackup/vital_db_dssu/tmp/iotest27  
/netbackup/vital_db_dssu/tmp/iotest28 /netbackup/vital_db_dssu/tmp/iotest29  
/netbackup/vital_db_dssu/tmp/iotest30 /netbackup/vital_db_dssu/tmp/iotest31
```

# Infrastructure Testing

- **Stress Testing iperf and stress combo**

Trying to crash SystemA when netconsole is in use

SystemA (be2net card 10Gb)

SystemB (10Gb card)

1) start iperf server on SystemB

2) start iperf client on System A

3) start iperf server different port on System A

4) start iperf client on System B matching port from #3

5) System A is configured with netconsole to send to System B

-while 1 loop doing echo 't' > /proc/sysrq-trigger to generate netconsole traffic

6) start nfs on SystemA

7) mount nfs from SystemB

8) cd to nfs mount dir on SystemB and run stress

stress --hdd 15 --hdd-bytes 2G

# Infrastructure Testing

- **Stress Testing iperf and stress combo details**

```
iperf -p 8001 -s
```

```
iperf -p 8001 -B 10.16.47.214 -c 10.16.47.68 -t 100000 -i 2
```

```
[root@dell-t5500-01]# cat netconsole_repeat.txt
```

```
#!/bin/bash
```

```
while [ 1 ]
```

```
do
```

```
echo 'm' > /proc/sysrq-trigger
```

```
sleep 2
```

```
done
```

# Infrastructure Testing

- **Build Your Own Infrastructure Regression Test Repository**
  - If you hit a BZ and in the BZ a reproducer is supplied keep this handy
    - Retest it yourself after the next OS upgrade
  - I would build a hammer test for each piece
    - Hammer those nic cards, firmware and drivers
    - Hammer that internal storage array hardware, firmware and drivers
    - Hammer that shared storage backend (NAS/SAN0, hba firmware and driver
  - The point that I'm trying to make here is that Red Hat will try to do the best job possible but it's very difficult to test out every hardware, firmware, and os release combination, so it's important that customers do their own testing.

# Important Configuration For Troubleshooting Issues (RCA)

# Important Configuration For Troubleshooting Issues (RCA)

- **Make sure all your Linux systems (including production) are properly configured and tested to create vmcore's**
  - <https://access.redhat.com/knowledge/solutions/6038>
  - <https://access.redhat.com/knowledge/solutions/5943>
  - <https://access.redhat.com/knowledge/solutions/39408>
- **Make sure all your RHEL systems (including production) have core dump collection enabled for userspace programs including suid daemons**
  - <https://access.redhat.com/knowledge/solutions/56021>
  - <https://access.redhat.com/knowledge/solutions/4896>
  - <https://access.redhat.com/knowledge/node/5352>
  - <https://access.redhat.com/knowledge/solutions/21787>



# Important Configuration For Troubleshooting Issues (RCA)

- **HA Clusters make it hard to capture vmcores, what else can you do**

- <https://access.redhat.com/knowledge/solutions/17382>
- <https://access.redhat.com/knowledge/solutions/4258>
- How about setting up serial console with vsp (HP, but should be a comparable for any hardware vendor) and then using screen command to ssh to serial console. Make sure to enable logging in screen ctrl-a then :logfile screen\_<hostname>log then ctrl-a then :log on I just tested this with a virsh console <somemachine> but this should work for a physical system as well.

<https://access.redhat.com/knowledge/solutions/45964>

This should give a full stack trace which may help in mapping to a Red Hat bugzilla

# Important Configuration For Troubleshooting Issues (RCA)

- **Accessing a hung box to force a vmcore**
  - <https://access.redhat.com/knowledge/node/2023>
  - <https://access.redhat.com/knowledge/solutions/45964>
  - Netconsole can help too
  - If you have a repeated problem with a server get a DRAC/RSA/ILO hooked up to it to potentially capture a stacktrace message
  - Make use of conserver to record serial console messages for all systems
    - <http://www.conserver.com/> (we're trying to get this into RHEL7)

# Red Hat Customer Portal

# Red Hat Customer Portal

[access.redhat.com](https://access.redhat.com)

The screenshot shows the Red Hat Customer Portal homepage. At the top, there is a dark navigation bar with the Red Hat logo, the text "CUSTOMER PORTAL", a search bar, and links for "Help" and "Register or Log in". Below this is a secondary navigation bar with tabs for "Knowledge", "Groups", "Support", "Downloads", "Security", and "Subscriptions".

The main content area features a central banner with the text "Access the industry's best resources for your FULL TECHNOLOGY LIFE CYCLE" and a flow diagram with three steps: "PLAN" (document icon), "DEPLOY" (box icon), and "CONNECT" (plug icon). Below the banner are four large buttons with icons: "GET STARTED WITH YOUR SUBSCRIPTION" (power icon), "ACTIVATE A NEW SUBSCRIPTION" (lightning bolt icon), "DOWNLOAD PRODUCTS AND UPDATES" (download icon), and "OPEN A NEW SUPPORT CASE" (folder icon).

A "FEATURED CONTENT" section highlights "Migrating Red Hat Enterprise Virtualization Manager from Version 2.2 to 3.0". Below this are three content blocks: "ABOUT SUBSCRIPTION MANAGEMENT" with a video thumbnail, "RED HAT STORAGE" with a "LIBERATE YOUR INFORMATION" graphic, and "TRENDING GROUP DISCUSSIONS" listing topics like "RHEVM : Update from 3.0 to 3.1beta failing", "RHEL vs Oracle Linux/Oracle Unbreakable Enterprise Kernel on large x86\_64 servers", "RHN Satellite advice sought: Integrating EPEL and other repos", and "tftp server".

At the bottom, there are four columns of "FREQUENT ACTIONS": "SUPPORT" (Login problems, Phone numbers & hours, Contact us), "SITE INFO" (Browser support policy, Site map, Awards & recognition), and "LEGAL" (Terms of use, Privacy policy, About Red Hat). The Red Hat logo is in the bottom right corner.

Log in with your Red Hat username and password.

# Using your subscription benefits

Red Hat enterprise product information from a source you know and trust - planning through production.



The Red Hat Customer Portal delivers the depth of knowledge and resources on Red Hat enterprise products that only Red Hat can provide.


- Expert videos
- Tech briefs
- Reference architectures
- Knowledgebase articles
- Downloads and demos
- Groups
- Activations
- Security updates
- Support case management
- Source browser
- Product lifecycle policies
- Subscription management
- Product documentation
- More...

# Groups

Collaborate with peers and Red Hat experts to share ideas on Red Hat enterprise products.

[access.redhat.com/groups](https://access.redhat.com/groups)

## High availability configuration

 **Daryl Rose**  
Community Member - 65 Points - 29 Jun 2011

Hello,

I am in the process of replacing two aging web servers. The platform of choice is RHEL. These servers must be in a high availability configuration. I'm am looking at suggestions on how to do this. Is clustering the right answer? A college of mine several months ago started looking into ricci and luci as his

This environment will consist of  
updated as well and then the pri

Anyone have any experiece w

Thank you.

Daryl

★★★★★  
Your rating: None Average: 4.5 (2 votes)

### a few questions

0 votes

What kind of content will be s  
you have backend applicatio  
Luci and ricci are front-end n  
themselves.  
Red Hat has an excellent clu  
that will fit the primary/sec  
single server.  
Some other tools that we use  
If you need to scale "horizon  
balancing solutions as well.

[reply](#) [Flag as offensive](#)

### Thank you, that helps

0 votes

Thank you, that helps out tremendously.

I believe that most of the content will be static, but I'm not really sure. Weblogic has been mentioned as running on the servers as well.

I'll review the link that you posed and read up on the documentation. Thank you for setting me straight on the ricci and luci.


[reply](#) [Flag as offensive](#)

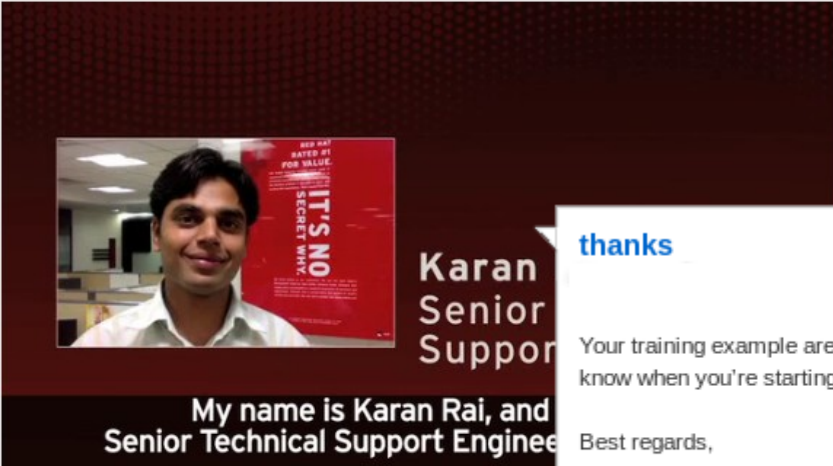
# Videos

[access.redhat.com/knowledge](https://access.redhat.com/knowledge)

Get vital solutions, useful tips, product demonstrations, and inside information straight from the source: Red Hat engineers and industry experts.

## Introduction to LDAP - Part 1

 **Karan Rai**  
Associate - Community Member - 75 Points - 21 Sep 2011



**Karan Senior Support**

**My name is Karan Rai, and Senior Technical Support Engineer**

Senior Technical Support Engineer Karan Rai gives an overview of LDAP - L how LDAP stores data, examples of LDAP databases, breakdown of common Interchange Format) works.

★★★★★  
Your rating: None Average: 4.4 (8 votes)

**thanks**

1  
vote

+ -

Your training example are very simple and very effective. I like your coaching method. Please let us know when you're starting other topic.

Best regards,

[reply](#) [Flag as offensive](#)

# Tech briefs

Tech briefs provide practical advice to help you solve real world problems. Each tech brief provides a detailed use case covering recommended practices, how-to instructions, or detailed discussions on Red Hat technology. Tech briefs are reviewed and tested by Red Hat engineers.

[access.redhat.com/knowledge](https://access.redhat.com/knowledge)

## How to Build an RPM



**Chris Negus**

Associate - Pro - 495 Points - 16 Sep 2011

 [rpm\\_building\\_howto.pdf](#)

If you create software to run in Red Hat Enterprise Linux, packaging it into an RPM makes it easier for you to manage, distribute, verify and update that software. This Tech Brief provides a basic HOWTO for creating an RPM package, signing it, distributing it (via a yum repository), and verifying its contents. It does this by illustrating the process using an existing source code package.

If you are taking the RHCE exam, this Tech Brief aims at covering the information needed for the RPM building requirement.

**Average:**



Your rating: None Average: 4.3 (4 votes)

### Great article

Wow, this is definitely the most recent and comprehensive article on building RPMS. Thanks Chris. I will study it carefully over the coming days. You're a brilliant techie.

0  
votes



[reply](#) [Flag as offensive](#)



# Reference architectures

[access.redhat.com/knowledge](https://access.redhat.com/knowledge)

## 2012 - RHEV 3.0 Disaster Recovery: Site-to-Site Fail-over



John Herr

Associate - Community Member - 50 Points - 18 Jan 2012

[2012\\_rhev\\_3\\_dr.pdf](#)

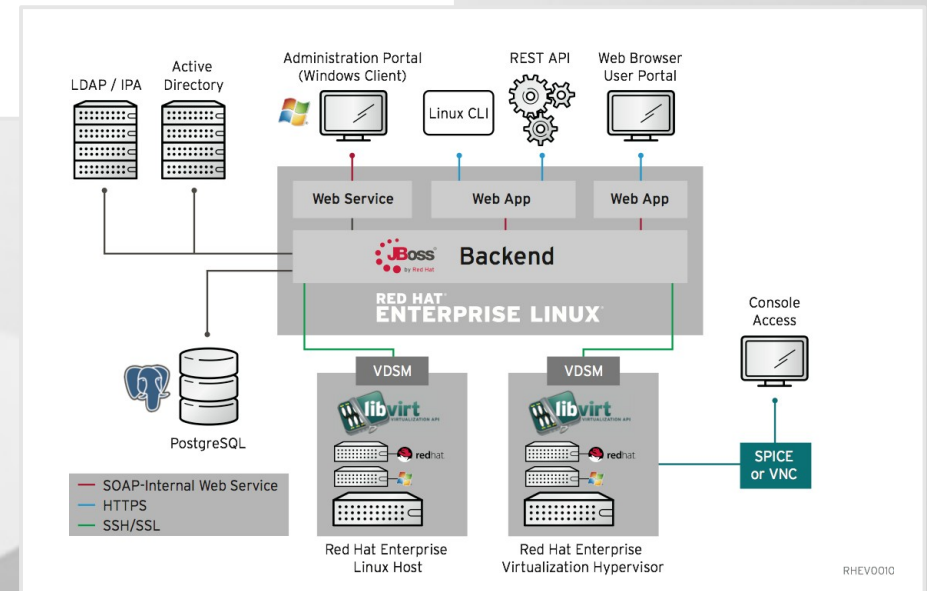
This paper demonstrates the process of failing a RHEV environment to a backup location during a site level disaster. Simulating a site to site LUN mirroring solution, three separate environments with running virtual machines are failed over to a backup location and brought up with minimal downtime.

Average:



Your rating: None Average: 5 (4 votes)

Detailed technical case studies of solutions built, tested, and benchmarked in Red Hat labs by Red Hat engineers.



# Knowledgebase

Find answers to support questions before you open a ticket using the same Knowledge-centered Support system that our engineers use to find and document resolved issues.

[access.redhat.com/knowledge](https://access.redhat.com/knowledge)

## How can core dumps be enabled for applications in Red Hat Enterprise Linux?

last modified by Ray Dassen on 03/30/12 - 06:56

### Issue

How can core dumps be enabled for applications in Red Hat Enterprise Linux?

### Environment

- Red Hat Enterprise Linux 3
- Red Hat Enterprise Linux 4
- Red Hat Enterprise Linux 5

### Resolution

A core dump is a on disk representation of a state of a user space application at the time of its abnormal termination. Core dumps can be used to diagnose the cause of a crash of an application.

Core dumps are not useful to end users, as they require specialised skills to analyze and determine the cause of the crash. In most cases, enabling a core dump is not necessary until there is a specific need.

In Red Hat Enterprise Linux, core file creation is disabled for interactive shells. This is configured in the `/etc/profile` by the `ulimit` command.

The `/etc/profile` contains system wide environment variables and commands that will be run for any interactive shell. By default, it contains this line:

#### Verified

This solution has been verified to work by Red Hat Customers and Support Engineers for the specified product version(s).

This solution is part of Red Hat's fast-track publication program, providing a huge library of solutions that Red Hat engineers have created while supporting our customers. To give you the knowledge you need the instant it becomes available, these articles may be presented in a raw and unedited form.

#### Rate

○ ★ ★ ★ ★

No votes yet

#### Email Notifications

 [Enable email notifications for this document](#)

# Adding users

Unlimited access to the Red Hat Customer Portal for every Red Hat subscription makes it easy to share knowledge throughout your organization with **user management** tools.

User List

Displaying users 1-15 of 146 total      View: all      Search:

User name	First name	Last name	Position	Dept	Created	Updated	Deactivated date	Deactivated by
jsmith	John	Smith	Automation Expert		1/13/10 1:58 PM EST	2/19/12 2:03 AM EST	2/19/12 2:03 AM EST	Peter Sanchez
jdoe	Jane	Doe	Senior Business Analyst		9/1/10 8:43 AM EST	10/9/11 12:10 AM EST		
psanchez	Peter	Sanchez			12/8/11 2:04 PM EST	12/8/11 2:06 PM EST		
jquirk	Joe	Quirk			12/11/08 10:26 AM EST	10/9/11 12:11 AM EST	1/12/10 2:57 PM EST	Peter Sanchez
cbrooks	Charlie	Brooks			12/11/08 10:53 AM EST	10/9/11 12:09 AM EST		
kbrown	Kathryn	Brown	Product Manager		10/24/07 8:18 AM EST	10/9/11 12:08 AM EST		
nojukwu	Naomi	Ojukwu			11/6/10 10:56 PM EST	10/15/11 10:30 PM EST		
ehidenberg	Emily	Hidenberg			11/14/07 3:35 AM EST	10/9/11 12:08 AM EST		



# Questions

Chris Henderson  
[chenders@redhat.com](mailto:chenders@redhat.com)