Avaya Aura® Collaboration Environment FAQ and Troubleshooting for Service Developers

Release 3.0
August 2014

© 2014 Avaya Inc.

All Rights Reserved.

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

“Documentation” means information published by Avaya in varying mediums which may include product information, operating instructions and performance specifications that Avaya may generally make available to users of its products and Hosted Services. Documentation does not include marketing materials. Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of documentation unless such modifications, additions, or deletions were performed by Avaya. End User agrees to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Avaya is not responsible for the contents or reliability of any linked websites referenced within this site or documentation provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Avaya provides a limited warranty on Avaya hardware and software. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya’s standard warranty language, as well as information regarding support for this product while under warranty is available to Avaya customers and other parties through the Avaya Support website: http://support.avaya.com or such successor site as designated by Avaya. Please note that if you acquired the product(s) from an authorized Avaya Channel Partner outside of the United States and Canada, the warranty is provided to you by said Avaya Channel Partner and not by Avaya.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, HTTP://SUPPORT.AVAYA.COM/LICENSEINFO OR SUCH SUCCESSOR SITE AS DESIGNATED BY AVAYA, ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AVAYA CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AVAYA CHANNEL PARTNER. UNLESS OTHERWISE AGREED TO BY AVAYA IN WRITING, AVAYA DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN AVAYA, AN AVAYA AFFILIATE OR AN AVAYA CHANNEL PARTNER; AVAYA RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS “YOU” AND “END USER”), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND AVAYA INC. OR THE APPLICABLE AVAYA AFFILIATE (“AVAYA”).

Avaya grants you a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to you. “Software” means Avaya’s computer programs in object code, provided by Avaya or an Avaya Channel Partner, whether as stand-alone products, pre-installed , or remotely accessed on hardware products, and any upgrades, updates, bug fixes, or modified versions thereto. “Designated Processor” means a single stand-alone computing device. “Server” means a Designated Processor that hosts a software application to be accessed by multiple users. “Instance” means a single copy of the Software executing at a particular time: (i) on one physical machine; or (ii) on one deployed software virtual machine (“VM”) or similar deployment.

License types

CPU License (CP). End User may install and use each copy or Instance of the Software on a number of Servers up to the number indicated in the order provided that the performance capacity of the Server(s) does not exceed the performance capacity specified for the Software. End User may not re-install or operate the Software on Server(s) with a larger performance capacity without Avaya’s prior consent and payment of an upgrade fee.

Named User License (NU). You may: (i) install and use the Software on a single Designated Processor or Server per authorized Named User (defined below); or (ii) install and use the Software on a Server so long as only authorized Named Users access and use the Software. “Named User”, means a user or device that has been expressly authorized by Avaya to access and use the Software. At Avaya’s sole discretion, a “Named User” may be, without limitation, designated by name, corporate function (e.g., webmaster or helpdesk), an e-mail or voice mail account in the name of a person or corporate function, or a directory entry in the administrative database utilized by the Software that permits one user to interface with the Software.

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, Hosted Service, or hardware provided by Avaya. All content on this site, the documentation, Hosted Service, and the Product provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software unless expressly authorized by Avaya. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

Virtualization

Each product has its own ordering code and license types. Note that each Instance of a product must be separately licensed and ordered. For example, if the end user customer or Avaya Channel Partner would like to install two Instances of the same type of products, then two products of that type must be ordered.

Third Party Components

“Third Party Components” mean certain software programs or portions thereof included in the Software or Hosted Service may contain software (including open source software) distributed under third party agreements (“Third Party Components”), which contain terms regarding the rights to use certain portions of the Software (“Third Party Terms”). As required, information regarding distributed Linux OS source code (for those Products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party

Components and the Third Party Terms that apply is available in the Documentation or on Avaya's website at: http://support.avaya.com/Copyright or such successor site as designated by Avaya. You agree to the Third Party Terms for any such Third Party Components.

Note to Service Provider

The Product or Hosted Service may use Third Party Components subject to Third Party Terms that do not allow hosting and require a Service Provider to be independently licensed for such purpose. It is your responsibility to obtain such licensing.

Preventing Toll Fraud

"Toll Fraud" is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud intervention

If you suspect that you are being victimized by Toll Fraud and you need technical assistance or support, call Technical Service Center Toll Fraud Intervention Hotline at +1-800-643-2353 for the United States and Canada. For additional support telephone numbers, see the Avaya Support website: http://support.avaya.com or such successor site as designated by Avaya. Suspected security vulnerabilities with Avaya products should be reported to Avaya by sending mail to: securityalerts@avaya.com.

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the Documentation, Hosted Service(s), and Product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the Documentation, Hosted Service(s) and Product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party.

Avaya is a registered trademark of Avaya Inc.

All non-Avaya trademarks are the property of their respective owners. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support website: http://support.avaya.com, or such successor site as designated by Avaya.

Contact Avaya Support

See the Avaya Support website: http://support.avaya.com for Product or Hosted Service notices and articles, or to report a problem with your Avaya Product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: http://support.avaya.com (or such successor site as designated by Avaya), scroll to the bottom of the page, and select Contact Avaya Support.

Contents

Chapter 1: SDK 7
- Reference Documents for Service Developers 7
- Why did my SDK installation fail? 7
- How do I know if the SDK install script succeeded? 7
- Where do I find Collaboration Environment Javadoc? 8
- Why does Eclipse show errors after I import the helloservice? 12
- Why does my project build fail with “Unable to locate the Javac Compiler in …”? 13
- Why does eclipse maven plugin install fail? 15
- How do I find the helloservice-svar target directory after building the service? 16
- Why don’t I see service-archetype in the list when I try to start my project from within Eclipse? 16
- Why do I have different project directories when I create a project from the archetype? 17

Chapter 2: Eclipse plugin 19
- Why can’t I see the new functionality availed from the newer Eclipse plugin version? 19
- My cluster list is not getting refreshed in the System Manager Browser – in plugin UI 19
- Cluster status is Green but there is a Red icon next to some of my Collaboration Environment servers. Why is that? 20
- I was building the service and tailing logs, but now I can’t see any logs in the console for my other service. 20
- I have multiple Collaboration Environment servers selected in a cluster and now I enabled CE debug mode. Can I debug my application on multiple Collaboration Environment servers? 21

Chapter 3: Services 23
- Can I create a service which has Call related features and can be invoked through HTTP/REST interface? How can I create this? 23
- How can I replace my service with newer iterations where I did not change the version number? 23
- How do I replace my service with a newer version where I changed the version number? 24
- How do I remove/delete my service? 26
- Why don’t I see my (re)deployed service? 28
- Why do attributes defined by my service not appear under Service Globals? 28
- How can I make HTTPS requests using the WebSphere certificate from my service? 29
- How can I define alarms in my service? 29
- How do I define notification OID? 30
- How do I define enterprise product id and product identifier? 30
- How do I define event codes? 30
- How do I enable clearing of a defined alarm? 31
- How do I raise defined alarms? 31
- How do I convert my Collaboration Environment 2.0 service to a 3.0 service? 31

Chapter 4: Service debugging 35
- How can I tell if my service got installed properly? 35
- Why if my service is not installed even after I created a cluster and added my service in the cluster? 37
- Why do I have to manually select my service for installation while creating a cluster especially since I can see some services which are already selected for installation? 37
- How can I find why my service installation failed? 37
- How can I tell if my service is running or gets invoked? 38
- Why is my service not getting invoked? 39
- Why are the changes from System Manager to service attributes of my service not reflected during execution or runtime? 52
- Why is my service not accessible over http but accessible over https? 54

How to use postman for https invocation to my service? 55
How do I import collections bundled with Sample Services to Postman client in Chrome? 55
How do I invoke Multi-Channel Broadcast service and Dynamic Team Formation Service for my testing? 62
How do I invoke ClickToCall Sample service for my testing? 64
What are the pre-requisites for the Callingpolicies Sample service? 64
How do I test the Callingpolicies Sample service? 65
My service is installed and I see logs from my service, but things are not behaving like I expected. What should I check? 66
Where do I look for the debug logs I added to my service? 72
Why does my outgoing HTTP request from my service not use the security module IP? 73
Why does HTTPS access to my services give error like “Cannot connect to real <IP address>”? 73
How do I debug media-related issues? 74
I've set up an implicit sequence pattern so that SM will route calls to the Collaboration Environment, but calls are not being sent to the Collaboration Environment. What might cause this? 75
I've set up an implicit sequence pattern so that SM will route calls to the Collaboration Environment, but some of the calls that should be matching my pattern are not getting sent to the Collaboration Environment. What might cause this? 76
How can I stop Collaboration Environment from logging me out? 77
Why don't I see my service alarm on System Manager or Data Center? 78
Why is access to my service so slow or sometimes not accessible? 78
How can I have my service not show “Not Applicable” for my service? 78
Why does my service show a warning icon with a date? 79
Why do I have call failures when using APIs which use Avaya Media Server? 79
My service is invoked and it runs media operations, or indicates that it might run media operations. Why does the call seem to stall? 80
Chapter 5: Connector debugging 83
Why was email not sent from my service? 83
Why was SMS not sent from my service? 86
Why was Scopia conference not scheduled from my service? 89
Why do I see the following services – EventingConnector and CallEventControl installed? 91
Chapter 6: HTTP Load Balancing and High Availability 93
How to enable/disable HTTP load balancer in a cluster? 93
How to locate the HTTP load balancer nodes in a cluster? 93
How is the cluster IP address used on the load balancer? 94
Why hasn’t the cluster IP address been assigned to any of the servers? 94
What if planned maintenance needs to be done on active load balancer node? 95
What if an unplanned failover happened? 95
Why do I see both load balancers are in active state? 95
How can I tell when the HTTP Load Balancer failed over? 96
How does the load balancer distribute requests in a cluster? 96
How can I tell whether the load balancer distributes the traffic evenly in cluster? 97
How does Session Affinity work on load balancer? 97
Chapter 7: Eventing Framework 99
Why do I stop receiving events after 8-9 hours? 99
How can I tell which subscriptions were matched when an event is sent? 99
Why am I getting multiple event notifications for a single event? 99
Why am I not getting my expected events? 100
Why are there log statements saying the Eventing System is still initializing 100
Why does my subscription not seem to activate? 101

Chapter 1: SDK

Reference Documents for Service Developers

1. *Collaboration Environment Service Development Guide*
2. *Multi-Channel Broadcast Sample Service*
3. *Dynamic Team Formation Sample Service*
4. *Calling Policies Sample Service*

Why did my SDK installation fail?

The SDK install script runs a few Maven commands. Verify that you have Maven correctly installed by running the command “mvn –version”. The output should show version information similar to this:

```
Apache Maven 3.2.1 (ea8b2b07643dbb1b84b6d16e1f08391b666bc1e9;
2014-02-14T10:37:52-07:00)
Maven home: /opt/apache-maven-3.2.1
Java version: 1.7.0_55, vendor: Oracle Corporation
Java home: /usr/java/jdk1.7.0_55/jre
Default locale: en_US, platform encoding: UTF-8
OS name: "linux", version: "2.6.32-358.el6.x86_64", arch: "amd64", family: "unix
```

Of course, your Maven home or Java home (as well as Java version and OS name) may vary from this example. If you do not see version information (or an error regarding the “mvn” command), review the Maven installation instructions to assure you have set up Maven correctly.

How do I know if the SDK install script succeeded?

The SDK install script runs a few separate commands. If any of these commands fail, the SDK install script will exit and indicate a failure. If you scan the last few lines of output and see “success”, then the install script completed successfully.

You can also check for certain SDK artifacts in your local Maven repository. To find your local repository, run this command: `mvn help:effective-settings` Near the end of the output, you will see some XML. Go to the top of the XML and look for the `localRepository`element:

```
<settings xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/
2001/XMLSchema-instance"xsi:schemaLocation="http://maven.apache.org/SETTINGS/1.1.0
http://maven.apache.org/xsd/settings-1.1.0.xsd">
<localRepository xmlns="http://maven.apache.org/SETTINGS/1.1.0 /maven/
repositorylocalRepository xmlns="http://maven.apache.org/SETTINGS/1.1.0 /maven/
repository>
<proxies xmlns="http://maven.apache.org/SETTINGS/1.1.0">
    <proxy>">/usr/local
        <port>8000</port>
```

In this example, the local repository starts at the directory /usr/local/maven/repository. Go to "/usr/local/maven/" directory; you should be able to see "archetype-catalog.xml" file and "repository"

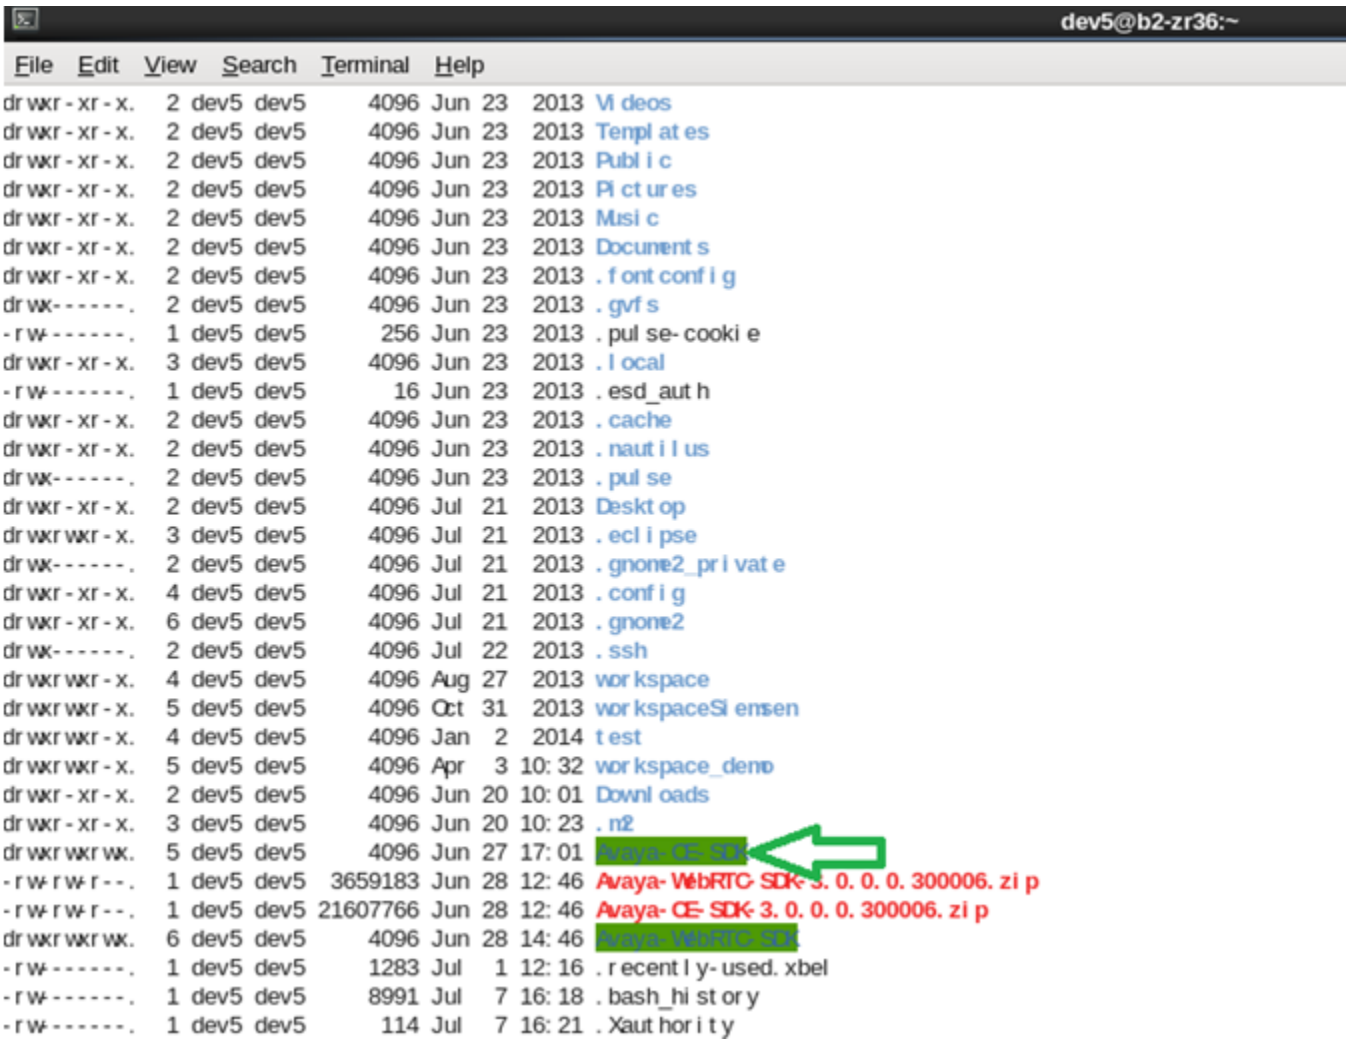
Navigate to com/avaya/. You should see the following directory structure:

```
[dev5@b2-zr36 avaya]$ ls -lart
total 32
drwxrwxr-x. 3 dev5 dev5 4096 Aug 25 21:30 ..
drwxrwxr-x. 4 dev5 dev5 4096 Aug 25 21:30 collaboration
drwxrwxr-x. 3 dev5 dev5 4096 Aug 25 21:30 common
drwxrwxr-x. 3 dev5 dev5 4096 Aug 25 21:30 csp
drwxrwxr-x. 6 dev5 dev5 4096 Aug 25 21:30 commons
drwxrwxr-x. 3 dev5 dev5 4096 Aug 25 21:31 ce
drwxrwxr-x. 8 dev5 dev5 4096 Aug 25 21:31 .
drwxrwxr-x. 9 dev5 dev5 4096 Aug 25 21:31 zephyr
```

You can navigate further into these directories to view the various artifacts required for collaboration environment service development.

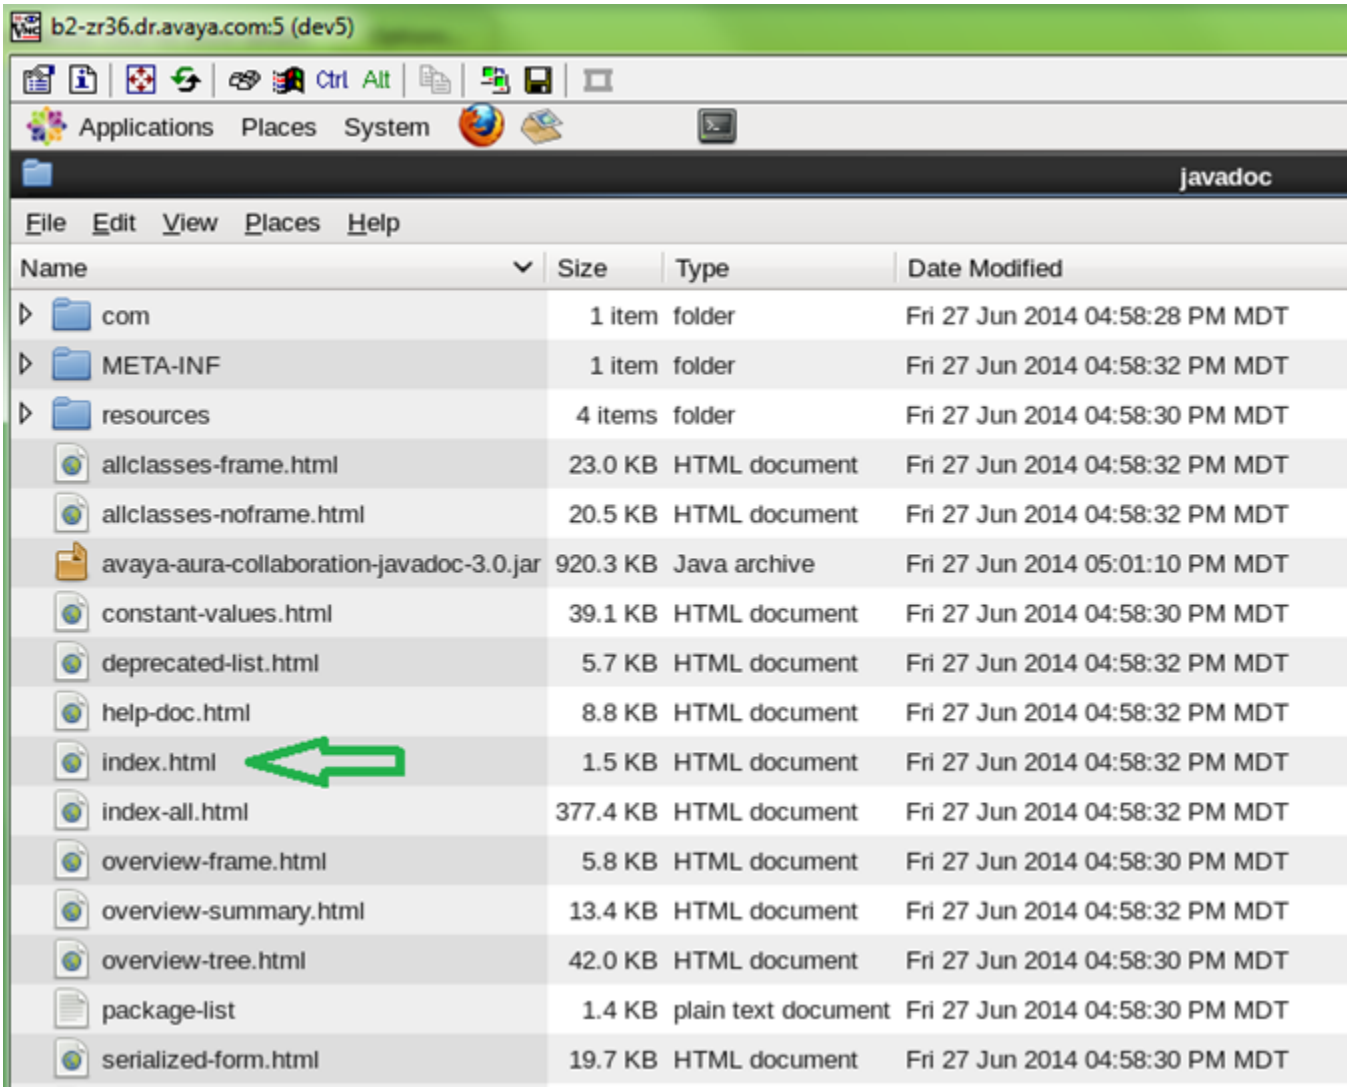
Where do I find Collaboration Environment Javadoc?

Login to your system where you have downloaded and copied the SDK. Assuming it is a linux system –

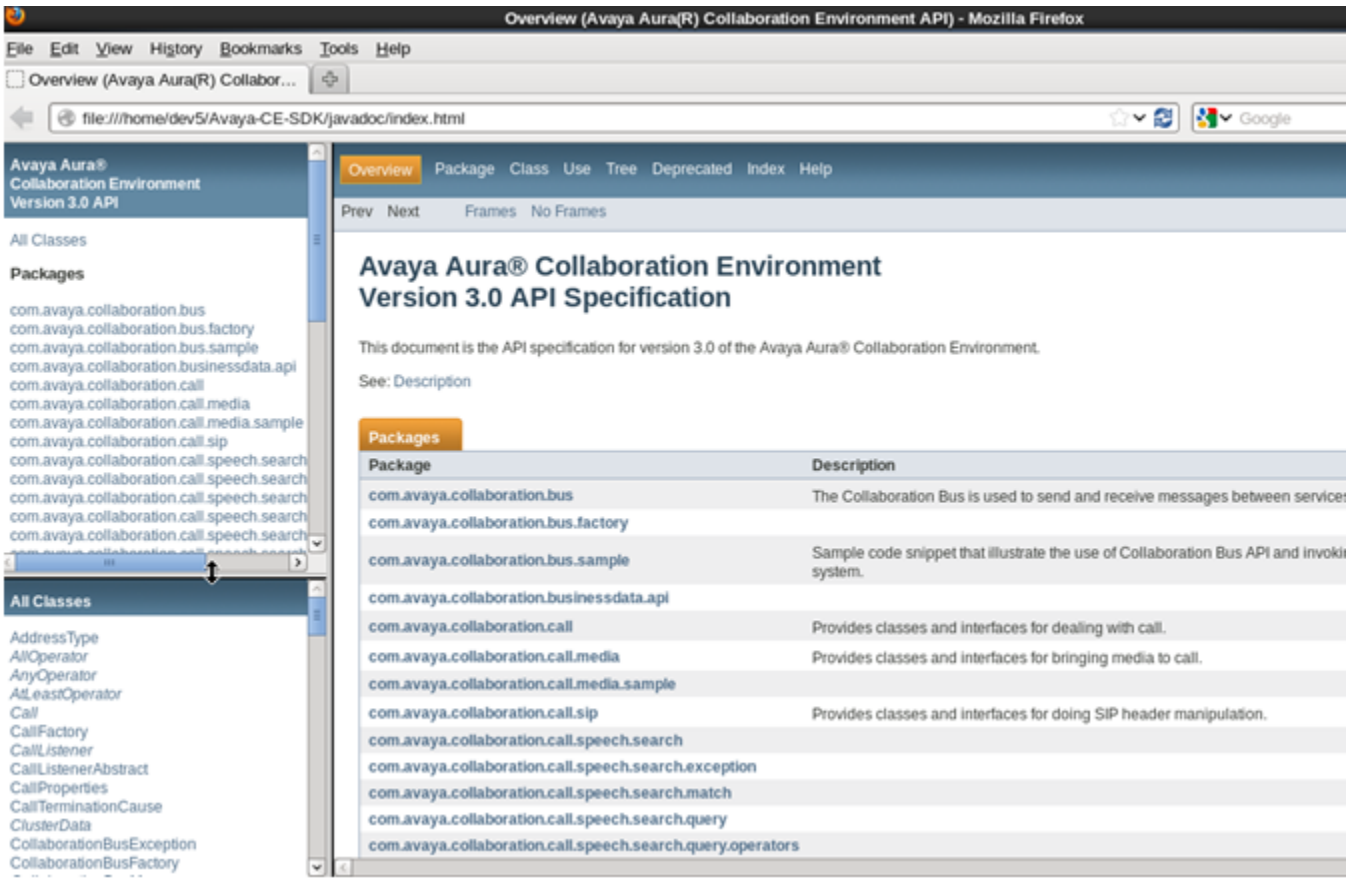


Execute **`cd <SDK-3.0 folder>/javadoc". avaya-aura-collaboration-javadoc-3.0.jar is the Javadoc jar`** is the Javadoc jar.

Execute **`jar -xvf avaya-aura-collaboration-javadoc-3.0.jar`**. Now open the index.html file.

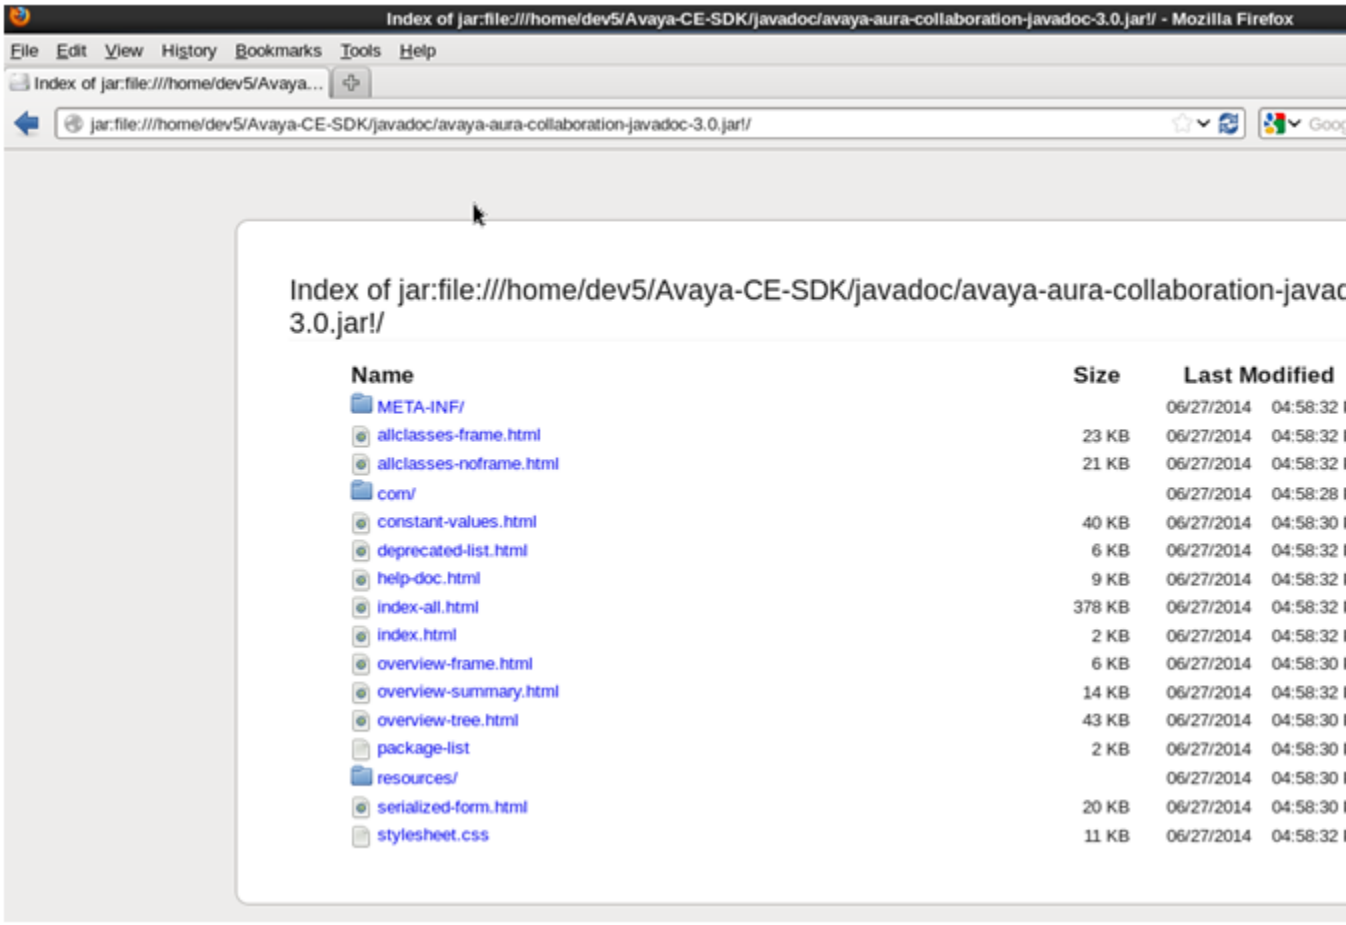


And you can see the Javadoc.



In case of a windows system, unzip the SDK. Unjar the avaya-aura-collaboration-javadoc-3.0.jar, and then open the index.html file.

Another neat trick is to use Firefox to open the Javadoc jar instead of exploding the jar yourself. You can use the following command to do it: **`jar:file:///home/dev5/<SDK-3.0 folder>/javadoc/avaya-aura-collaboration-javadoc-3.0.jar!/`**.



And then launch index.html by clicking the link.

Similar things can be done on windows systems too. The browser URL would look like –

jar:file:///C:/Users/amishra/Downloads/<SDK-3.0 folder>/javadoc/avaya-aura-collaboration-javadoc-3.0.jar!/

Click on index.html to look at the Javadoc

The Collaboration API javadocs can also be accessed by clicking on a class name in eclipse, pressing F2 and clicking Open Attached Javadoc in Browser.

Why does Eclipse show errors after I import the helloservice?

1. From the Eclipse menu, select Project – Clean, then select either “Clean all projects” or “Clean projects selected below” and select all the helloservice projects.
2. If that did not clear the errors, right-click on the top-level project (that would be helloservice without a suffix), then select Run As – Maven build… In the “Goals”

field of the Edit Configuration dialog, enter “clean package”, and then click Run. If errors still exist after the build, right-click on the top-level project and select Refresh.

3. Verify that you have configured Eclipse to use JDK 1.7 or later. From the top menu, select Project – Properties – Java Compiler, and check the Compiler compliance level.
4. In Eclipse, open the Markers window (from the Eclipse menu, select Window – Show view – Markers). The error text should give some clue as to the problem.
5. If you are still unable to resolve the errors within Eclipse, try building from the command line instead of in Eclipse (using the command “mvn clean package”). If this shows errors, and you believe your environment (Java, Maven, etc.) is set up correctly

   Try deleting the maven repository. To find your local repository, run this command: **`mvn help:effective-settings`**Near the end of the output, you will see some XML. Go to the top of the XML and look for the **localRepository** element:

   ```
   <settings xmlns=http://maven.apache.org/POM/4.0.0 xmlns:xsi=http://
   www.w3.org/2001/XMLSchema-instance xsi:schemaLocation="http://
   maven.apache.org/SETTINGS/1.1.0 http://maven.apache.org/xsd/
   settings-1.1.0.xsd">
   <localRepository xmlns="http://maven.apache.org/SETTINGS/1.1.0 /maven/
   repository</localRepository>
   <proxies xmlns="http://maven.apache.org/SETTINGS/1.1.0">
   <proxy>">/usr/local
   <port>8000</port>
   ```

   In this example, the local repository starts at the directory /usr/local/maven/repository. Go to “/usr/local/maven/” directory and delete everything under this directory.
6. Delete the SDK and then re-install the SDK according to the SDK installation instructions.

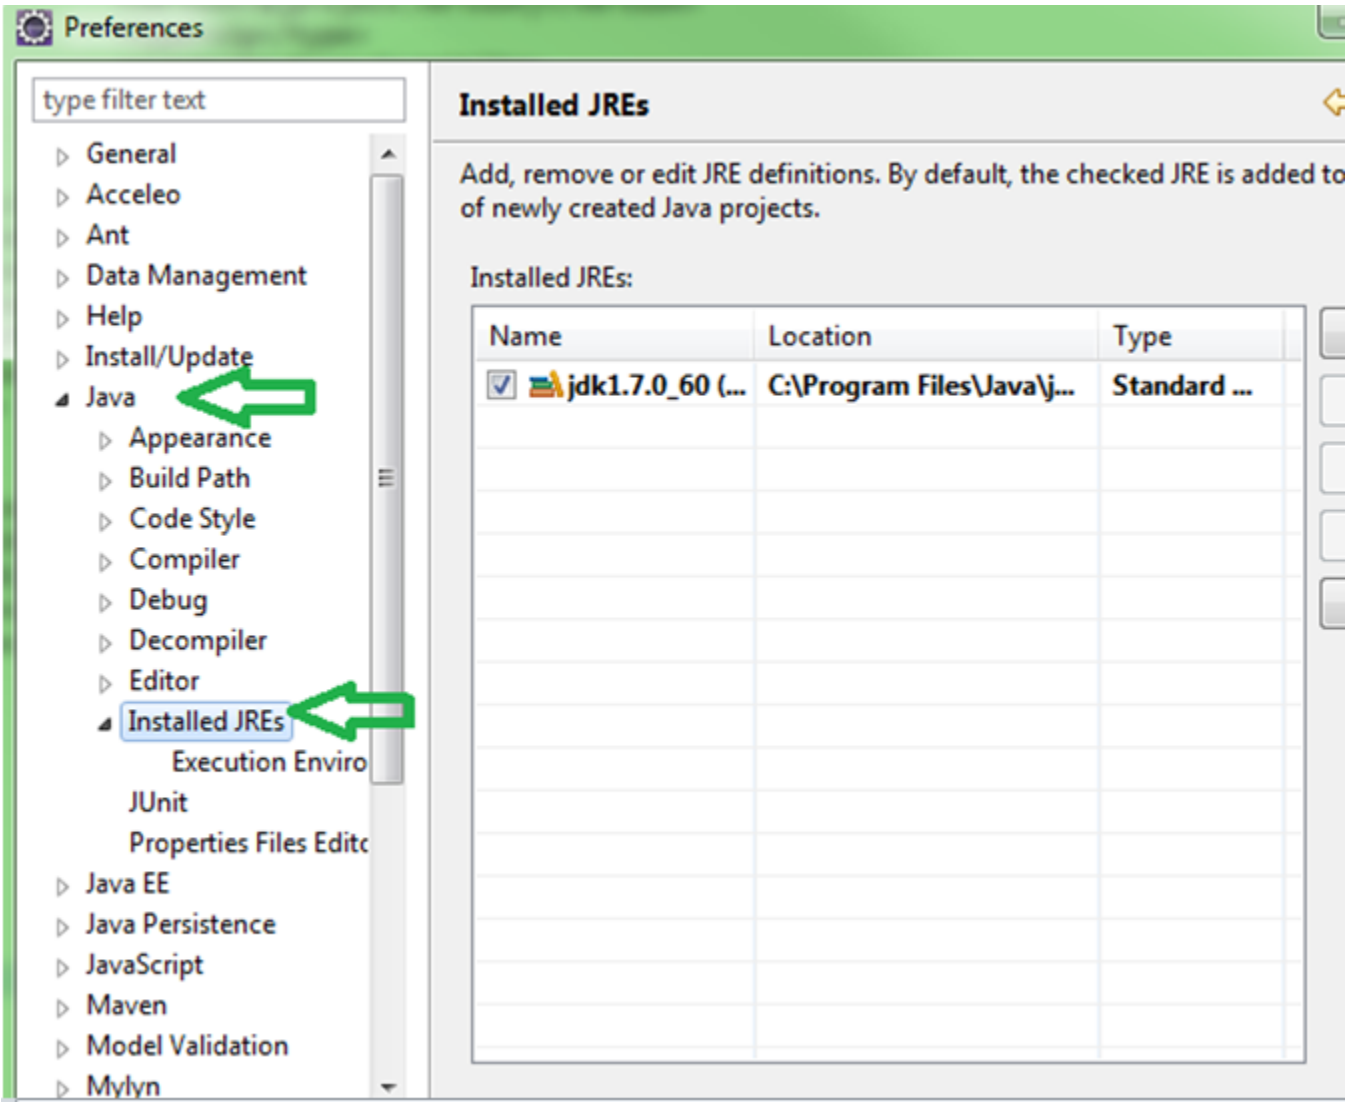
Why does my project build fail with “Unable to locate the Javac Compiler in …”?

If you see the following issue while building the project:

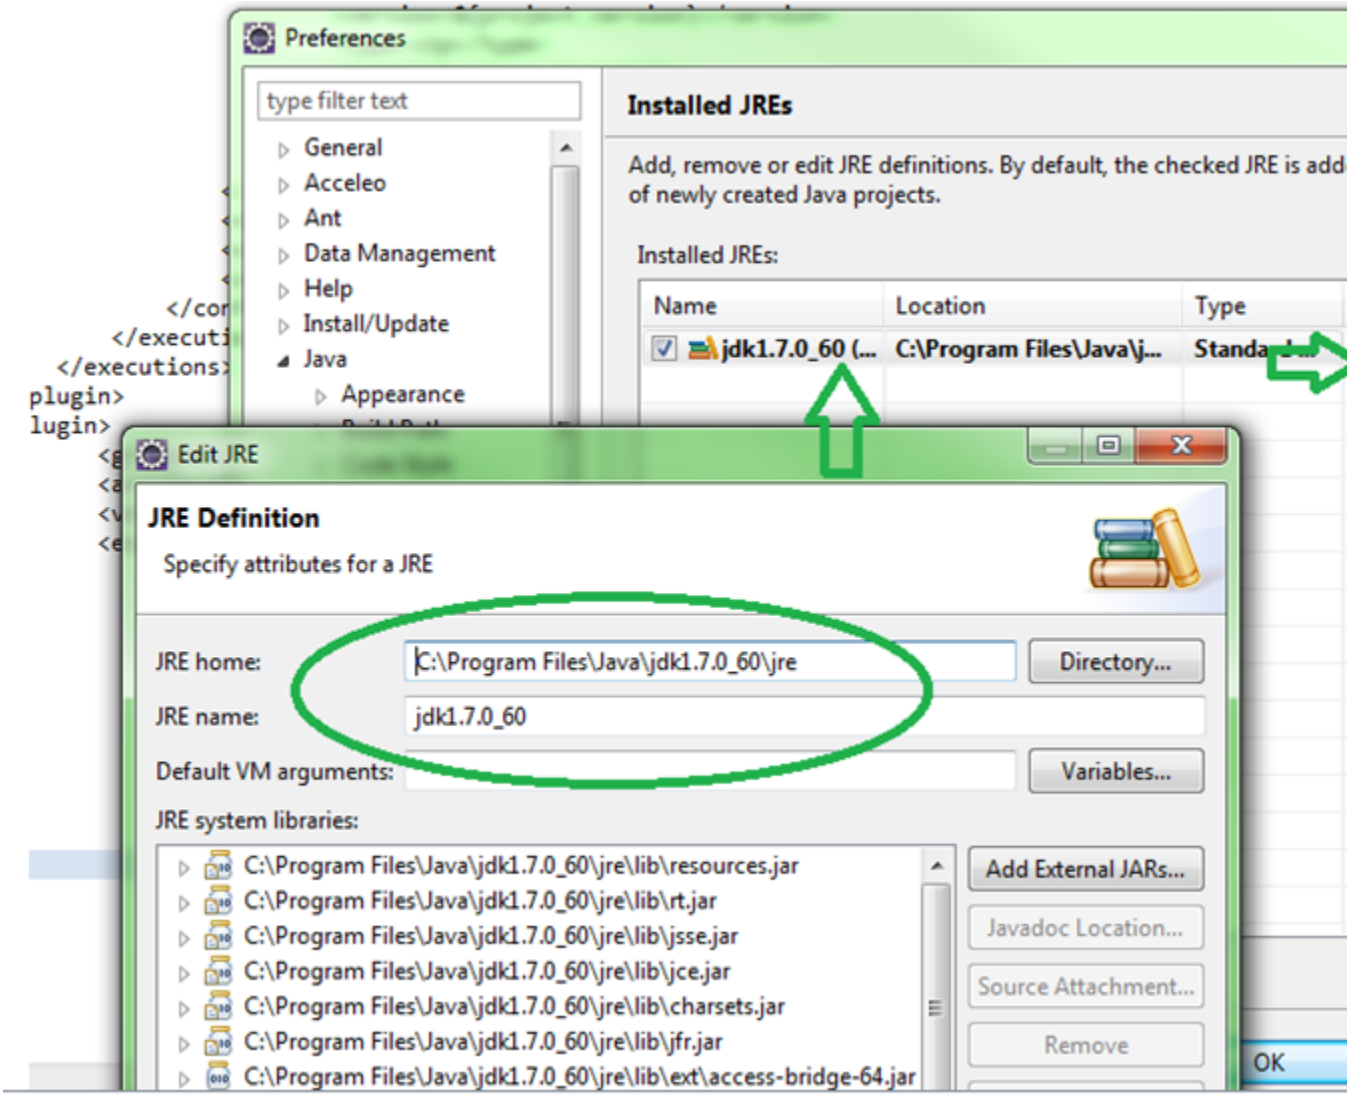
```
[INFO] BUILD FAILURE
[INFO] ------------------------------------------------------------------------
[INFO] Total time: 1.344s
[INFO] Finished at: Wed Sep 11 09:31:08 MDT 2013
[INFO] Final Memory: 5M/15M
[INFO] ------------------------------------------------------------------------
[ERROR] Failed to execute goal org.apache.maven.plugins:maven-compiler-plugin:
2.3.2:compile (default-compile) on project testtracy-war: Compilation failure
[ERROR] Unable to locate the Javac Compiler in:
[ERROR] C:\Program Files\Java\jre7\..\lib\tools.jar
[ERROR] Please ensure you are using JDK 1.4 or above and
[ERROR] not a JRE (the com.sun.tools.javac.Main class is required).
[ERROR] In most cases you can change the location of your Java
```

```
[ERROR] installation by setting the JAVA_HOME environment variable.
[ERROR] -> [Help 1]
```

Go to Eclipse -> Window -> Preferences -> Java -> Installed JREs :



Make sure that the JDK path is mentioned in the list and selected. If not select the entry and click on Edit:



Finish and save all the screens and build the project again.

Why does eclipse maven plugin install fail?

Check if you have Eclipse Kepler installed. If yes, Eclipse Kepler already includes a maven plugin and you do not need to install the maven plugin again.

Collaboration Environment recommends using Eclipse Luna.

How do I find the helloservice-svar target directory after building the service?

After the build completes, in Eclipse right-click on the top-level project (that would be helloservice without a suffix) and select Refresh.

After the build completes, the directory structure will look like this (some directories and files omitted to keep this smaller):

```
|-helloservice
 |---helloservice-ear
 |-----src
 |-----target
 |---helloservice-svar
 |-----src
 |-------main
 |---------assembly
 |---------resources
 |-----target
 |------- HelloWorld-0.0.0.1.5.svar  <- this is the svar to install
 |-------archive-tmp
 |-------dependency
 |-------dependency-maven-plugin-markers
 |---helloservice-war
 |-----src
 |-------main
 |---------java
 |-------test
 |---------java
 |-----target
```

Why don’t I see service-archetype in the list when I try to start my project from within Eclipse?

- When starting a new Maven project, verify that the checkbox for “Create a simple project (skip archetype selection)” is unchecked.
- In the Eclipse Maven archetype selection page, check that the Catalog field is set to All Catalogs.
- Verify that the SDK installer completed successfully. The SDK install script runs a few separate commands. If any of these commands fail, the SDK install script will exit and indicate a failure. If you scan the last few lines of output and see “success”, then the install script complete successfully.
- Go to the maven repository (see the question “How do I know if the SDK install script succeeded” to find your maven repository); you should be able to see “archetype-catalog.xml” file. If the file is not present, try deleting everything from maven repository and then re-installing the SDK according to the SDK installation instructions.

Why do I have different project directories when I create a project from the archetype?

Different project directories are different maven artifacts which together create a service archive file (.svar) for deployment on a CE. Taking an example of the helloservice, eclipse will show 4 project directories:

```
|-helloservice
|---helloservice-ear
|-----src
|-----target
|---helloservice-svar
|-----src
|-------main
|---------assembly
|---------resources
|-----target
|---helloservice-war
|-----src
|-------main
|---------java
|-------test
|---------java
|-----target
```

`|-helloservice –` This is the base maven project which contains “helloservice-ear”, “helloservice-svar” and “helloservice-war”.

`|-helloservice-ear –`This is the ear application which gets bundled into the svar for deploying. This contains the application.xml which defines the enterprise application archive.

`|-helloservice-svar –` This maven project is used to bundle the ear into a svar. It contains the manifest.xml, properties.xml which defines a service archive.

`|-helloservice-war –` This maven project can contain all the required servlets, REST services and classes along with the core logic of a Collaboration Environment service.

It is also possible to have an ejb jar in your service in case you would like to execute some EJB related operations (session beans, JPA etc). For this, you will need to have an appropriate code structure and modify the pom.xml for the ear to bundle the ejb jar into your project. E.g the email connector has a dependency similar to the following in the pom.xml of the ear project:

```
<dependencies>
  <dependency>
     <groupId>com.xyz.emailconnector</groupId>
     <artifactId>emailconnector-ejb</artifactId>
     <version>XYZ</version>
     <type>ejb</type>
  </dependency>
</dependencies>
```

pom.xml of the ejb project should have lines similar to:

```
<parent>
  <artifactId>emailconnector</artifactId>
```

```
  <groupId> com.xyz </groupId>
</parent>
<groupId>com.xyz.emailconnector</groupId>
<artifactId>emailconnector-ejb</artifactId>
<packaging>ejb</packaging>

<name>emailconnector-ejb</name>
```

Build section of pom.xml of the EJB project should have following lines:

```
<build>
   <plugins>
    ……
    ……
      <plugin>
        <groupId>org.apache.maven.plugins</groupId>
        <artifactId>maven-ejb-plugin</artifactId>
        <version>2.3</version>
        <configuration>
          <ejbVersion>3.1</ejbVersion>
        </configuration>
      </plugin>
      …...
      …...
   </plugins>
</build>
```

pom.xml of the project root should have lines similar to:

```
<modules>
   <module>emailconnector-ejb</module>
   ….
   ….
</modules>
```

Chapter 2: Eclipse plugin

Why can’t I see the new functionality availed from the newer Eclipse plugin version?

About this task

The older version eclipse plugin jar was maintained in a cache in Eclipse environment. Follow the steps below to handle this problem:

Procedure

1. Close Eclipse.
2. Delete the jar com.avaya.zephyr.eclipse.plugin from <Eclipse install path>/plugins/ directory.
3. Restart the Eclipse and verify that no Eclipse plugin UI and button is visible on toolbar.
4. Copy this same jar from <unzipped SDK location>/lib/ com.avaya.zephyr.eclipse.plugin.jar to the plugins directory above.
5. Start Eclipse and verify that the plugin icon is displayed in toolbar.

My cluster list is not getting refreshed in the System Manager Browser – in plugin UI.

In case user edits a cluster or adds / deletes clusters in System Manager, those changes are not refreshed in the plugin UI.

Procedure

1. You need to select the System Manager added, right click the entry and select Delete SMGR.

2. Re-add the System Manager entry using menu options to view the correct , updated cluster information.

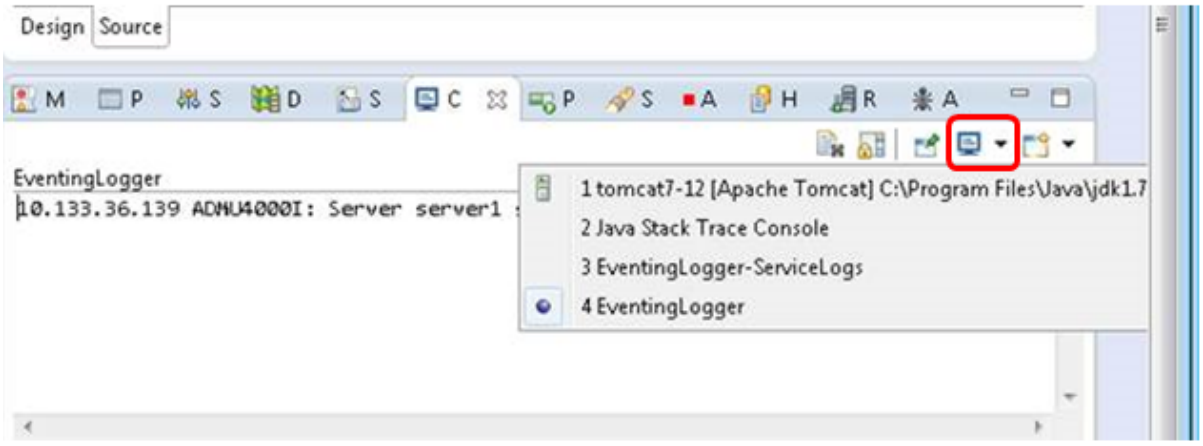
Cluster status is Green but there is a Red icon next to some of my Collaboration Environment servers. Why is that?

If a certain icon is Red, it means that the particular Collaboration Environment server in currently unreachable.

The icons displayed adjacent to the Collaboration Environment servers in a cluster display the availability of the Collaboration Environment server.

I was building the service and tailing logs, but now I can’t see any logs in the console for my other service.

The Eclipse plugin uses the Standard Eclipse console to display logs, build results and status of any operations. However at times due to Eclipse plugin development limitations, the right instance of the console cannot be auto focused. You need to click on the icon highlighted below to display all the console windows and select the relevant one.



I have multiple Collaboration Environment servers selected in a cluster and now I enabled CE debug mode. Can I debug my application on multiple Collaboration Environment servers?

No. The eclipse plugin by default selects the first selected Collaboration Environment in the cluster and enables debug mode.

Chapter 3: Services

Can I create a service which has Call related features and can be invoked through HTTP/REST interface? How can I create this?

Yes, this is possible. When you use service-archetype to create a new service, the archetype creates skeleton code which you can then modify (including renaming):

- MyCallListener.java – This provides a base file to make call related requests using the Call apis.
- MyApplication.java – This is needed if you want your service to be invoked by a REST web service. This class is needed to provide the base application path for your REST webservice.
- MyResource.java – This class contains the GET/POST/PUT/DELETE methods for your REST web service.
- MyServlet.java – This class is needed if you want to invoke a service using a servlet.

More details on how to create a new service can be found in – *Collaboration Environment Service Development Guide*. Look at Trimming the Project section as well.

How can I replace my service with newer iterations where I did not change the version number?

During development, you will probably go through cycles of coding and testing. For multiple iterations of your service, you do NOT need to change the service version number. You provided the version number in the Maven service-archetype when first starting development of the service. This version number is kept as the **serviceVersion** property in the project POM. For the “helloservice” sample service provided in the SDK, for example, the directory structure is (trimmed down for context):

```
|-helloservice
|---pom.xml      <- project pom
|---helloservice-ear
|-----pom.xml
|---helloservice-svar
|-----pom.xml
|-----src
|-------main
```

```
|---------assembly
|---------resources
|------------properties.xml
|---helloservice-war
|-----pom.xml
```

This project pom contains this definition:

```
[properties]
        [serviceName]HelloWorld[/serviceName]
        [serviceVersion]0.0.0.1.5[/serviceVersion]
    [/properties]
```

You do NOT need to change the **serviceVersion** just to progressively add to and test your service. When developing in this mode, you may find it more convenient to replace your service on the Collaboration Environment server, rather than through the System Manager interface.

These instructions assume you are merely iterating your service rather than changing the version. Note that you cannot use this approach if you have changed service attributes (in the properties.xml file) that you want to test on System Manager. You will instead need to re-install your service through System Manager. To do this, copy the SVAR from your build to your Collaboration Environment server.

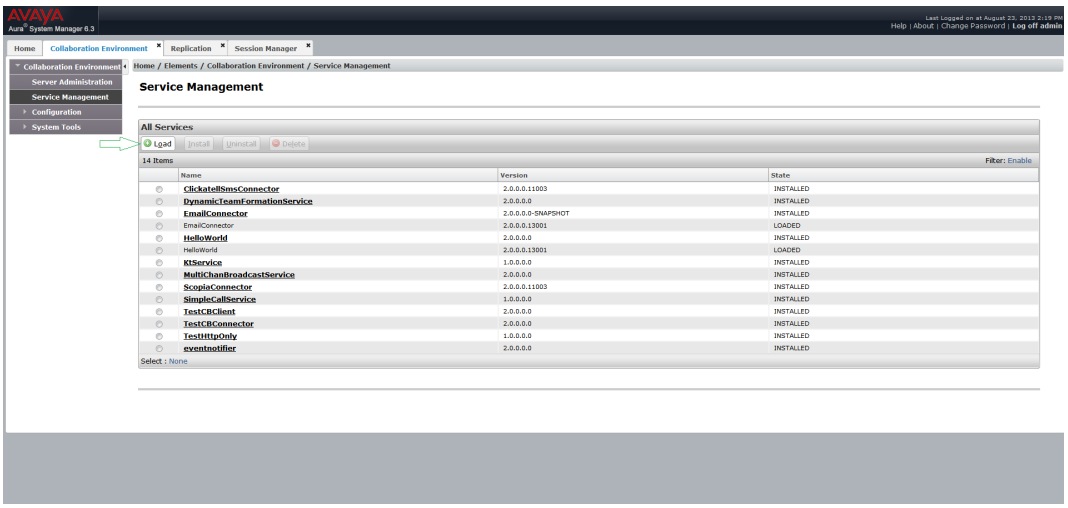
On the Collaboration Environment server, run this command: **`deploy_service <path-to-your-svar>`**. For example, if you copied HelloWorld-0.0.0.1.5.svar to the /tmp directory on your Collaboration Environment server, the command would be: **`deploy_service /tmp/HelloWorld-0.0.0.1.5.svar`**The output of this command should indicate whether the service deployed successfully. You can always check using the **`deploy_service -lv`** command. The command **`deploy_service -h`** describes the command in general, including how to interpret the output of the **`deploy_service -lv`** command.

This approach would still require you to install your service one time through System Manager. All subsequent replacements, as long as the service name and version remain the same, could be done with this command.

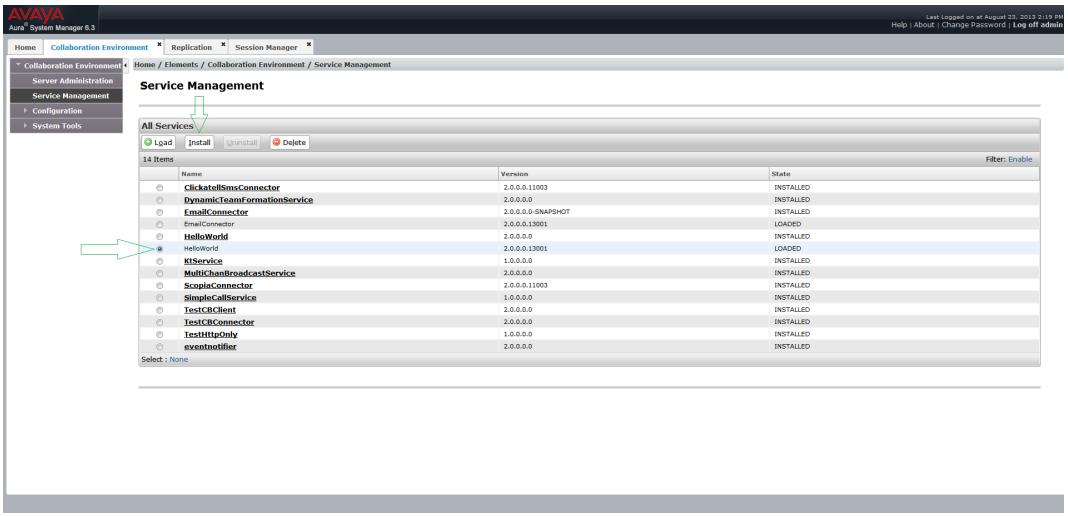
How do I replace my service with a newer version where I changed the version number?

If you changed the version number of your service (that is, you changed the **serviceVersion** property in your project pom), you will need to install the new version on System Manager.

1. Log in to System Manager. On the System Manager Home page, select Collaboration Environment from the Elements panel.
2. In the Collaboration Environment page, select Service Management from the left menu.
3. Load the newer version of your service. Press the Load button and follow the prompts.



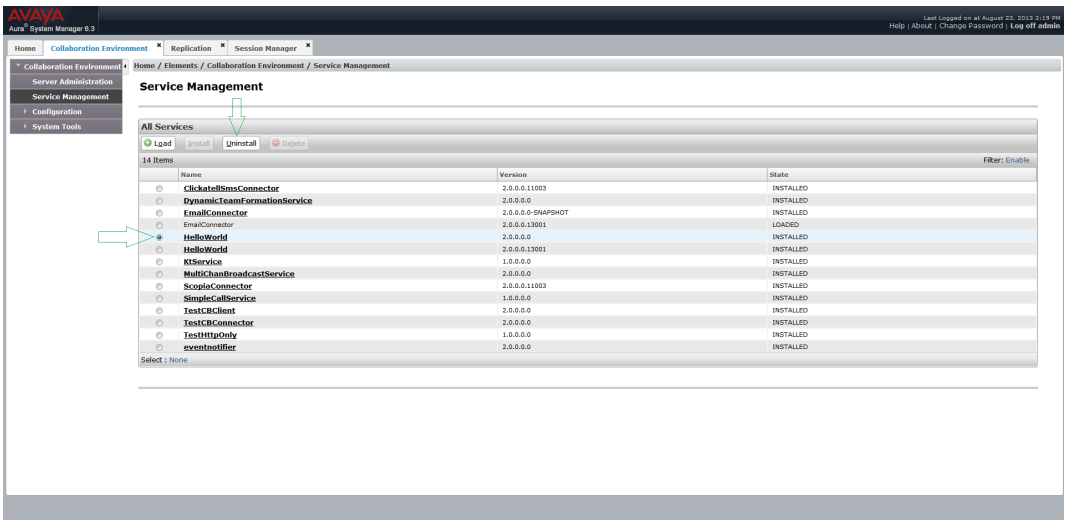
4. After the load completes, check the radio button by the newer version of your service and press the Install button. On click of install, you would be asked to select the cluster/clusters of Collaboration Environment servers where the service should be installed. Choose appropriate clusters. Note: You can also use cluster Administration screens to install services.



5. You can run multiple versions of your service concurrently. You might find this useful, for example, to run the newer version with a small subset of users until you feel the service is ready to roll out to a larger community. You can use the Set Preferred Version to achieve this. However, if you want to remove an older version of your service, find the older version of your service in the list on Service Management page, check the radio button next to your service, and press the Uninstall button. Select appropriate cluster/clusters from where you would like to uninstall the service.

Note:

You can also use cluster administration screens to uninstall services from specific clusters



6. Once uninstall completes, select your service and press the Delete button.

How do I remove/delete my service?

There are 2 distinct operations for removing a service:

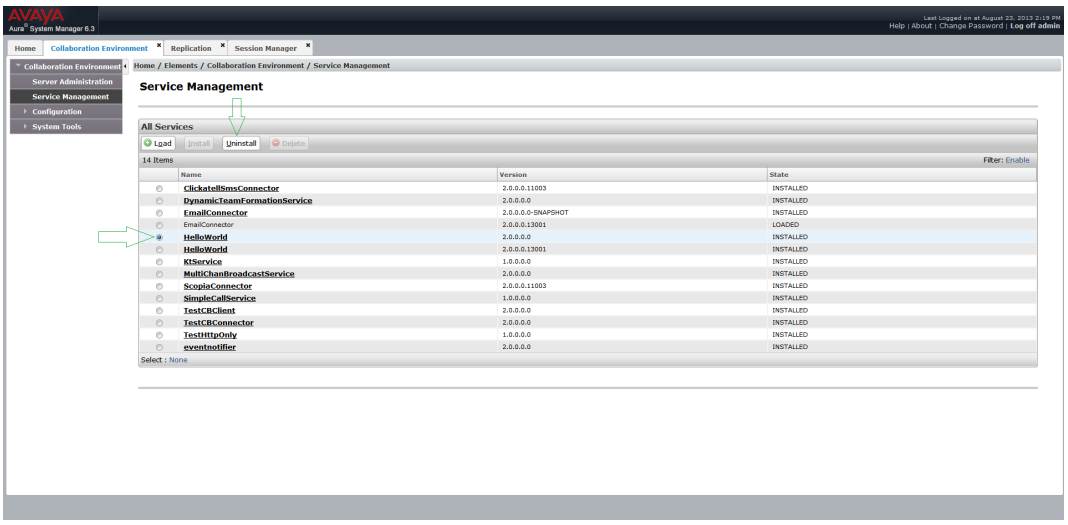
1. Uninstall – This step does an un-deploy of the service on all configured Collaboration Environment servers, but leaves entries in the database.
2. Delete – This step removes entries from the database. This means that the service will be deleted from any service profiles that included the service; you should configure users with other service profiles, or add this service back to the service profiles if you decide to re-install this service.

Warning:

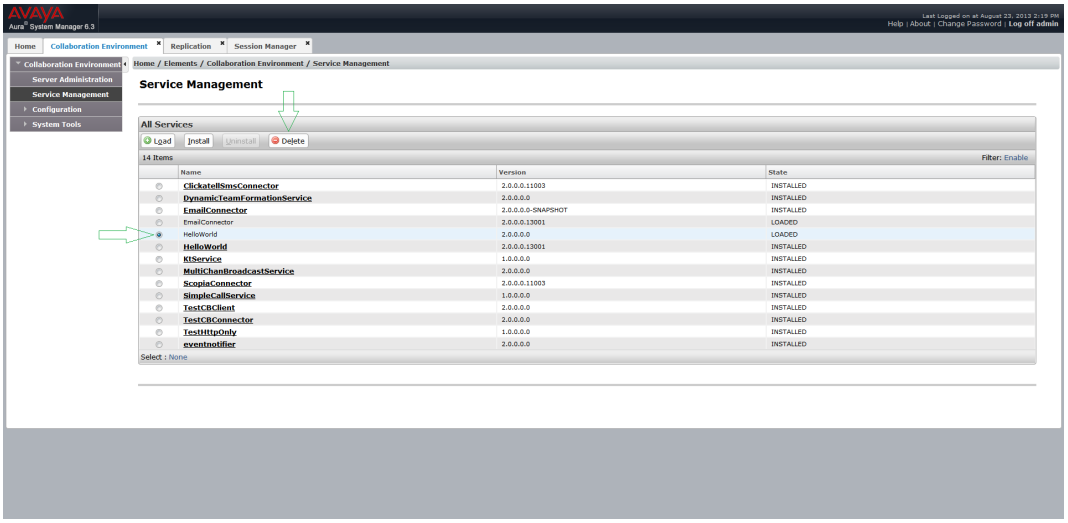
On deleting the last remaining version of a service, all the attributes are wiped out. If you want to upgrade your service rather than a true deletion and want keep old attribute data intact, you must consider adding the new version of the service BEFORE removing the old one.

Procedure

1. Log in to System Manager. On the System Manager Home page, select Collaboration Environment from the Elements panel.
2. In the Collaboration Environment page, select Service Management from the left menu.
3. Find your service, check the radio button next to your service, and press the Uninstall button.



4. When uninstall completes, select your service and press the Delete button.



5. To verify that your service was removed, log in to the Collaboration Environment server and run this command: **`deploy_service -l`**

Your service should not appear in the list output by this command. When you uninstall and delete a service from System Manager, System Manager propagates commands to all configured Collaboration Environment servers. The time for the operation to complete will vary depending on the number of configured Collaboration Environment servers, network issues, and so on. You may need to wait a minute or two and then run the **`deploy_service -l`** command again.

Why don’t I see my (re)deployed service?

If you initially installed your service through System Manager and then later replaced it using the **deploy_service** command, it might be possible that the version changed.

1. Log in to System Manager. On the System Manager Home page, select Collaboration Environment from the Elements panel.
2. In the Collaboration Environment page, select Service Management from the left menu.
3. Find your service and note the Version.
4. In the workspace where you build your service (assuming that you used the service archetype to start your project in Maven), find the **pom.xml** file at the project level. Here is the directory structure of an example service (with many files and directories removed to make this explanation concise):

```
|-helloservice
|---pom.xml    <- this one
|---helloservice-ear
|-----pom.xml
|---helloservice-svar
|-----pom.xml
|---helloservice-war
|-----pom.xml
```

5. Look for the **serviceVersion** property in this pom file. If this version differs from the version shown on System Manager, you will need to either change the **serviceVersion** property or install your service on System Manager.

Why do attributes defined by my service not appear under Service Globals?

Attributes provided in properties.xml of the “-svar” project must conform to the XSD. XSD is part of archive-xsd.jar in the lib directory of the Collaboration Environment SDK.

If the xml is created as per XSD, it is possible that you are trying to install the same service with a new version and the following should be kept in mind while doing this :

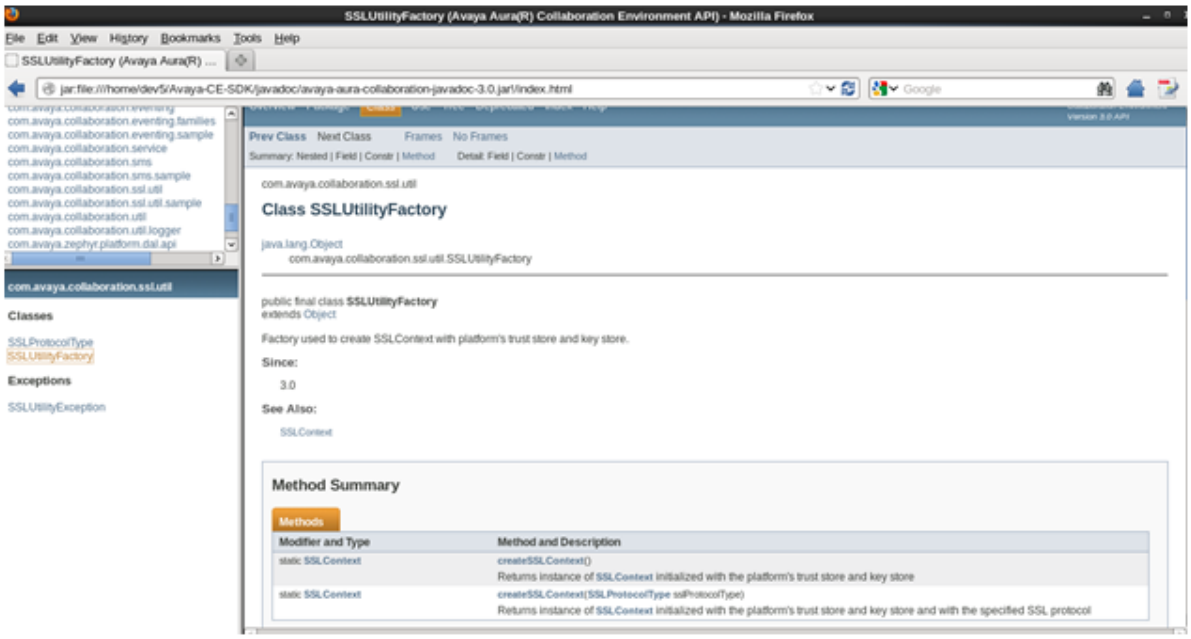
- You cannot modify the attribute values in the new version of the service. Please add an attribute or uninstall the old version of service and delete the old version of the service before installing the new version of the service.
- Make sure to verify the “admin_visible” flag is set to “true” and also “global” flag should be set to “false”.
- You cannot modify the validation type “<validation name="anyString"><type>STRING</type></validation>”

In the above, validation name of “anyString” could be used in other fields of already deployed service. So, if you change the type from “STRING” to “ENCRYPTED_STRING”, the service

would not show the encrypted string. You can add a new attribute to work around this or uninstall the old version of service and delete the old version of the service before installing the new version of the service.

How can I make HTTPS requests using the WebSphere certificate from my service?

Look at SSLUtility APIs provided in the Collaboration Environment Javadoc on how to get the required certificate for making SSL connection from your service.



How can I define alarms in my service?

The alarms are defined in an alarms.xml file and needs to be bundled in the service archive.

```
|-helloservice
 |---pom.xml
 |---helloservice-ear
 |-----pom.xml
 |---helloservice-svar
 |-----src
 |-------main
 |----------resources
 |-------------alarms.xml <- this one
```

```
 |---helloservice-war
 |-----pom.xml
```

Example:

```
<?xml version="1.0" encoding="UTF-8"?>
<p:alarms xmlns:p="http://archiveschemas.aus.avaya.com/snapinalarms"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:schemaLocation="http://archiveschemas.aus.avaya.com/snapinalarms alarms.xsd
">
    <p:alarm>
        <p:eventcode>EVENT_01</p:eventcode>
       <p:alarmDisplayText>Service couldn't connect to the DB.</p:alarmDisplayText>
        <p:severity>major</p:severity>
        <notificationOid>.1.3.6.1.4.1.6889.2.68.10.1</notificationOid>
<clearFlag>true</clearFlag>
<clearNotificationOid>.1.3.6.1.4.1.6889.2.68.10.2</clearNotificationOid>
    </p:alarm>
</p:alarms>
```

How do I define notification OID?

The SNMP Notification OID is the unique identifier for a trap/alarm. Alarms usually have two notifications OIDs - one for setting the alarm and one for resolving the alarm. This is an optional field but if you are defining it make sure it's uniquely defined under the product. It is a combination of dots and digits. `Ex- .1.3.6.1.4.1.6889.2.68.10.1.`

How do I define enterprise product id and product identifier?

The SNMP Notification OID contains the enterprise id and product id for a trap/alarm.

'enterpriseToProductId' – Id starting from enterprise to product.

'enterpriseToProductIdentifier' – Identifier starting from enterprise to product. Ex.-

```
<enterpriseToProductId>6889.2.63</enterpriseToProductId>
<enterpriseToProductIdentifier >avaya.products.ce</enterpriseToProductIdentifier>
<notificationOid>.1.3.6.1.4.1.6889.2.63.10.1</notificationOid>
```

How do I define event codes?

The event code specified in the alarms.xml file must be >= 2 characters and <= 12 characters. Collaboration Environment automatically prefixes the event code with the snap-in name and uses the CLR_ prefix to denote a clear event. The total string formed by combining the snap-in name and the CLR_ prefix must not exceed 32 characters.

How do I enable clearing of a defined alarm?

Alarms usually have two notifications OIDs - one for setting the alarm and one for resolving the alarm. If you are setting clear flag to true, make sure clear notification oid is uniquely defined in alarms.xml. Ex-

```
<clearFlag>true</clearFlag>
<clearNotificationOid>.1.3.6.1.4.1.6889.2.68.10.2</clearNotificationOid>
```

How do I raise defined alarms?

Get the instance of Collaboration Logger (com.avaya.collaboration.util.logger.Logger). Log the event using logger's logEvent API. `Ex. logger.logEvent("EVENT_CODE");.`

How do I convert my Collaboration Environment 2.0 service to a 3.0 service?

Before you begin

- 3.0 SDK is successfully installed
- JDK 1.7 is installed.

Procedure

1. • In the pom.xml present in war project replace dependency below :

   ```
   <dependency>
                    <groupId>com.avaya.collaboration.api</groupId>
                   <artifactId>avaya-aura-collaboration-api-2.0</artifactId>
                   <version>2.0.2.1.202105</version>
                       <scope>provided</scope>
       </dependency>
   ```

   With

   ```
   With
   <dependency>
                    <groupId>com.avaya.collaboration.api</groupId>
                   <artifactId>avaya-aura-collaboration-api-3.0</artifactId>
                   <version>3.0.0.0.300010version>
                   <scope>provided</scope>
   </dependency>
   ```

2. Replace the JRE system Library 1.6 dependency with 1.7 JRE System libraries.

3. In the pom.xml for service svar project , make changes below. Replace the highlighted part:

```
<artifactItem>
        <groupId>com.avaya.collaboration.api</groupId>
        <artifactId>avaya-aura-collaboration-api-2.0</artifactId>
        <version>2.0.2.1.202105</version>
        <type>jar</type>
        <overWrite>true</overWrite>
        <outputDirectory>${project.build.directory}/tmp</outputDirectory>
        <destFileName>sdk.properties</destFileName>
                  <includes>META-INF/MANIFEST.MF</includes></artifactItem>
with
             <artifactId>avaya-aura-collaboration-api-3.0</
artifactId>
             <version>3.0.0.0.300010</version>
```

4. Replace the highlighted part:

```
<dependency>
            <groupId>com.avaya.zephyr.zephyrCommon.xsds</groupId>
            <artifactId>archive-xsd</artifactId>
            <version>2.0.2.1.202105</version>
            <exclusions>
            <exclusion>
            <artifactId>smc-api</artifactId>
            <groupId>com.avaya.zephyr.services.smc</groupId>
            </exclusion>
            </exclusions>
    </dependency>

with
<version>3.0.0.0.300010</version>
```

5. Also need to add snippet below to the pom.xml for service svar project:

```
<validationSet>
            <dir>${basedir}</dir>
            <systemId>${basedir}/target/dependency/alarms.xsd</systemId>
            <includes>
            <include>src/main/resources/alarms.xml</include>
            </includes>
</validationSet>
```

   You can now add an alarms.xml under path: `\<workspace>\<servicename>\<servicename>-svar\sra\main\resources\alarms.xml` to avail of the dynamic snap in alarming capability in 3.0

6. Add snippet below to dist.xml located at

```
/<svcname>-svar/src/main/assembly/dist.xml
    <file>
    <source>src/main/resources/alarms.xml</source>
    <outputDirectory>/</outputDirectory>
    <filtered>true</filtered>
    </file>
```

7. CARRules.xml should contain TerminatingServiceRule similar to the following

```
<TerminatingServiceRule desc="for callback to CMATestService">
        <FeatureURI>CMATestService</FeatureURI>
    </TerminatingServiceRule>
```

CMATestService is an example and it should be replaced with your service name

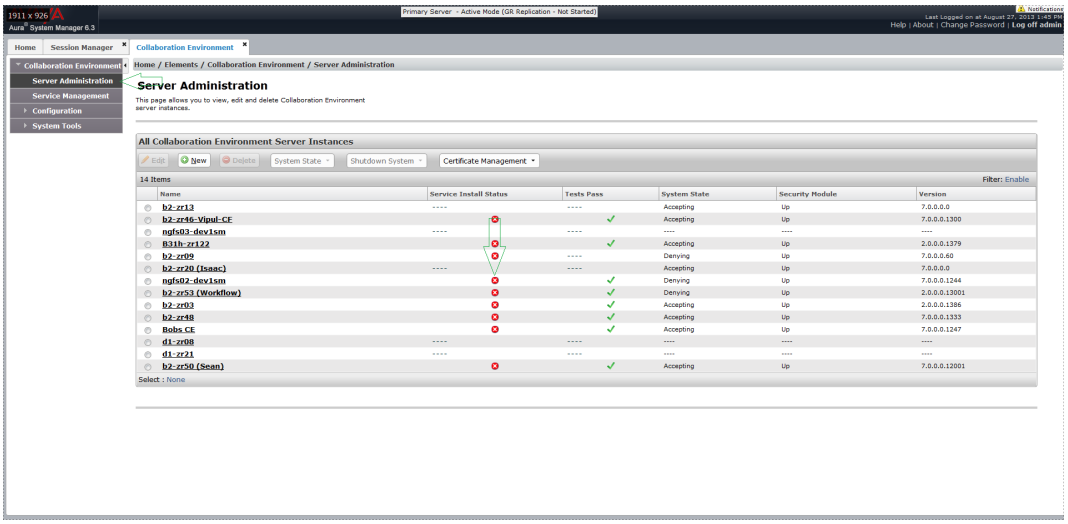
8. Refer the service archetype for Collaboration Environment 3.0 for appropriate sip.xml and web.xml. There are new servlets which are required for services to function correctly. You can also refer to Multichannel Broadcast sample application to get a feel of how the sip.xml and web.xml look from 3.0 onwards.

Chapter 4: Service debugging

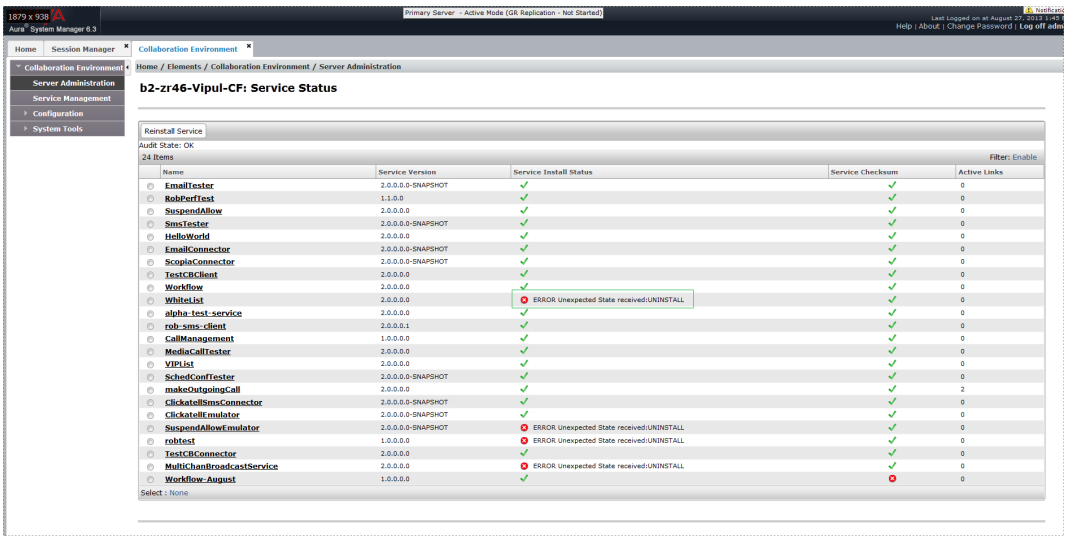
How can I tell if my service got installed properly?

A service is installed from System Manager. System Manager will propagate the service to all configured Collaboration Environment servers, and an audit running on a Collaboration Environment server will deploy the service. On each Collaboration Environment server, the service is deployed. Note that some time is required for System Manager to propagate the service to all Collaboration Environment servers, so this operation takes longer in larger systems with more Collaboration Environment servers. But, in general, the service should deploy within a couple minutes of installing it on System Manager.

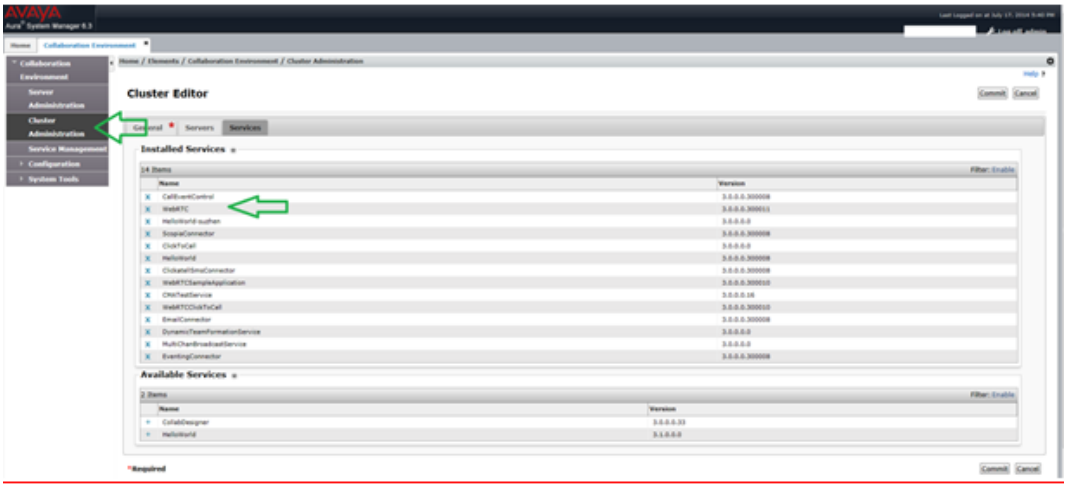
1. Log in to System Manager. On the System Manager Home page, select Collaboration Environment from the Elements panel.
2. In the Collaboration Environment page, select Server Administration from the left menu. You will see a red cross in the Service Install Status column.



3. Click on the red cross and you will be taken to the Service status page for the server. Service install status column on this screen will provide information on which service failed to install on Collaboration Environment.



4. If you are using cluster administration to install the services as in screen below –



Make sure that your service is shown in the installed services list and appropriate servers are selected in the servers tab.

5. If your service is in the Installed state according to System Manager, log in to your Collaboration Environment server and run this command: **`deploy_service -lv`**

   The output will look something like this:

```
load: + deploy: + run: + car: N  HelloWorld-0.0.0.1.3
load: + deploy: + run: + car: +  basicSipCall-1.0.0.0.0
```

   In this example, the service basicSipCall shows a plus (+) for each step of the deployment, so is installed and ready to go. The service HelloWorld deployed but failed the “car” step of the deployment (“N” indicates the step failed), which means that the service will not be invoked.

   You can get an explanation of the **`deploy_service`** command by asking for help: **`deploy_service -h`**

6. If your service is in the list produced by the **`deploy_service -lv`** command, but has an “N” for any step, try manually redeploying the service – the output of that

may provide a clue as to why the service failed a step: **deploy_service <SVAR file>**

When you installed the service from System Manager, the service archive (SVAR) was propagated to all configured Collaboration Environment servers. The copy of the SVAR resides in /var/avaya/aus_svars. You can find your SVAR by running this command: **ls /var/avaya/aus_svars/**

The output will be something like this:

```
basicSipCall-svar-1.0.0.0.0-SNAPSHOT.svar  helloservice-svar-0.0.1-
SNAPSHOT.svar
```

Why if my service is not installed even after I created a cluster and added my service in the cluster?

Select the cluster that was created and click edit. Go to servers tab and make sure that Collaboration Environment servers are added in cluster.

Why do I have to manually select my service for installation while creating a cluster especially since I can see some services which are already selected for installation?

Services that appear to be added automatically are all mandatory services defined for that particular cluster. There is no mechanism to define your service as a mandatory service so it needs to be added manually to get it installed.

How can I find why my service installation failed?

Logon to your Collaboration Environment server and try manually redeploying the service through command line – the output of that may provide a clue as to why the service failed a step: **deploy_service <SVAR file>**

When you installed the service from System Manager, the service archive (SVAR) was propagated to all configured Collaboration Environment servers. The copy of the SVAR resides in /var/avaya/aus_svars. You can find your SVAR by running this command: **ls /var/avaya/aus_svars/**

The output will be something like this: `basicSipCall-svar-1.0.0.0.0-SNAPSHOT.svar helloservice-svar-0.0.1-SNAPSHOT.svar`

In-order to redeploy a command similar to this can be executed: `deploy_service /var/avaya/aus_svars/helloservice-svar-0.0.1-SNAPSHOT.svar`

If the re-deploying of service from command line does not provide obvious hints on deploy failure, run this command: **`ce dlogv`**

This will open the debug log using the vi editor. Search for your service name. This should provide you with information if there were any exceptions that lead to failure in installing your service.

If you added logging statements in your service code, you should be able to spot your log statements in the debug log. If you used the Logger class in the Collaboration Environment API (highly recommended), you can view the service log using the command **`ce dlogv <servicename>`**.For example, to view the logs for the basicSipCall service, run the following command - **`ce dlogv basicSipCall`**.

The logs for different versions of a service all go to the same log file, and each log message contains the service name and version.

When running in normal mode, logs at level INFO and higher priority (FATAL, ERROR, WARN) will be added to the log file; logs at lower priority (FINE, FINER, FINEST) will not be added to the log. The log level can be changed on the fly to provide more verbose logging. Enter this command while logged in to your Collaboration Environment server to enable all log levels: **`ce dlogon <servicename>`**

When you are finished debugging, you should return logging to the default using this command: **`ce dlogoff <servicename>`**

Default logging (sysout or syserr) goes to a different log. If you are using Java “system.out.println (not recommended)”, the text will be written to the default log. Enter this command to view the log in the vi editor: **`ce alogv`**

How can I tell if my service is running or gets invoked?

1. Log in to your Collaboration Environment server and run this command: **`deploy_service -lv`**

   The output will look something like this:

   ```
   load: + deploy: + run: + car: N  HelloWorld-0.0.0.1.3
   load: + deploy: + run: + car: +  basicSipCall-1.0.0.0.0
   ```

   If your service is in the list and shows a plus (+) for each step, it is installed and ready to go. You can run the command **`deploy_service -h`** to get more detail about the output of the **`deploy_service -lv`** command.

   If the car is “N” for a call related service, make sure that the service name and service version is exactly same in sip.xml as well as web.xml. Also, make sure the sip servlets match in both sip.xml and web.xml.

2. The framework will log that it has deployed, registered, and initialized your service. On your Collaboration Environment server, run this command: **`ce dlogv`**

This will open the debug log using the vi editor. Search for the string “initialization complete for”. The full log string will include your service name and version, so as an alternative you could search for your service name.

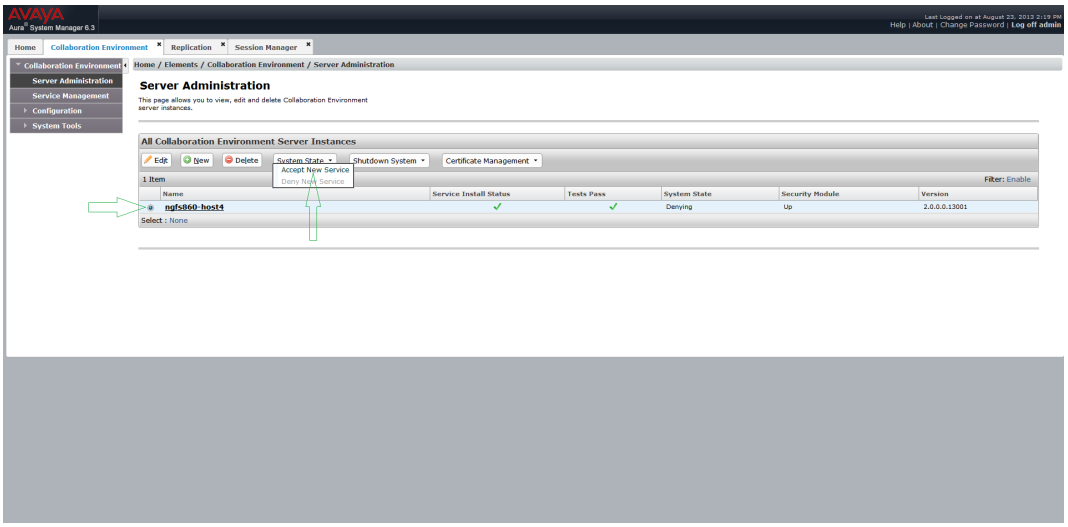
If you added logging statements in your service code, you should be able to spot your log statements in the debug log. If you used the Logger class in the Collaboration Environment API (highly recommended), you can tail/watch the service log using the command **`ce dlogw <servicename>`**

For example, to tail the logs for the basicSipCall service, run the following command- **`ce dlogw basicSipCall`**.

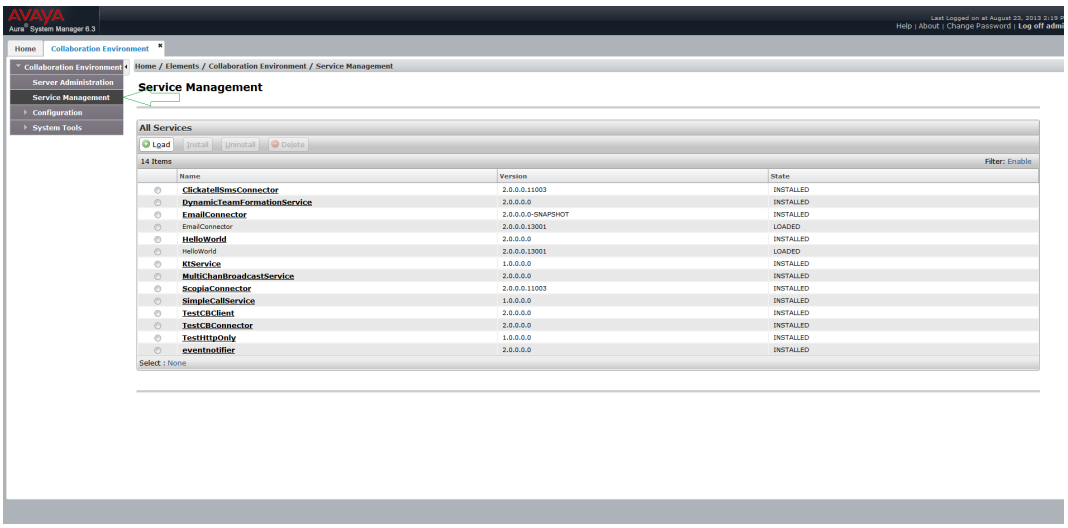
Why is my service not getting invoked?

Your Collaboration Environment server might not be in the state to accept service.

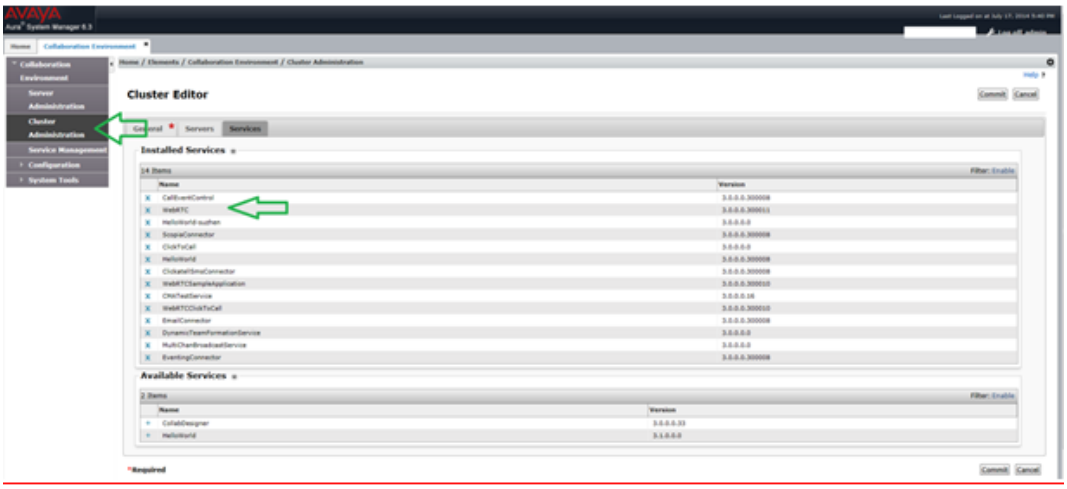
1. On the System Manager Home page, select Collaboration Environment from the Elements panel.
2. Pick the radio button next to your Collaboration Environment server. Click the System State button to pull down the menu, and then select “Accept New Service” from the dropdown menu.



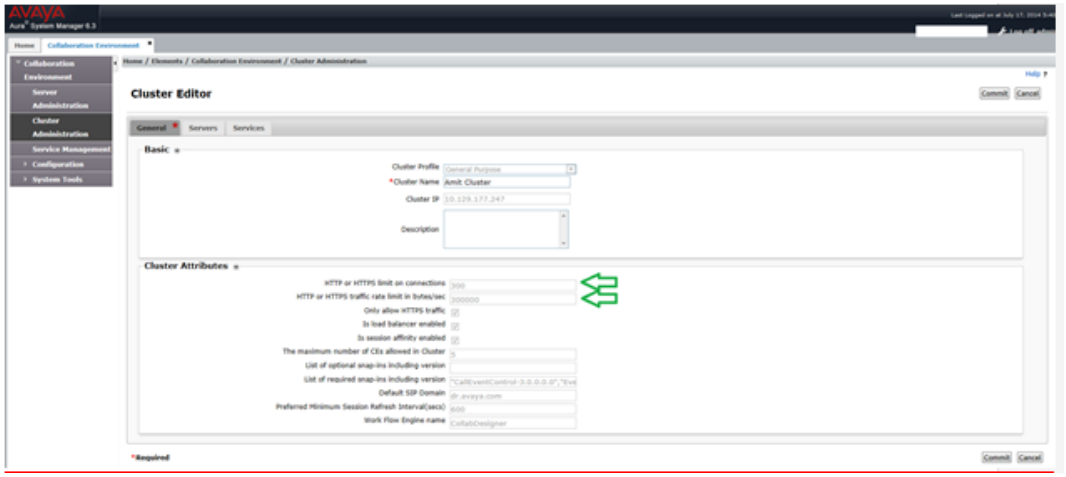
3. Verify that your service is installed. Log in to System Manager. On the System Manager Home page, select Collaboration Environment from the Elements panel.
4. In the Collaboration Environment page, select Service Management from the left menu.



5. Find your service. If your service does not appear in the table, press the Load button to load the SVAR you created when building your service. If the State for your service is Loaded, select the checkbox and press the Install button.
6. If you are using cluster administration to install the services as in screen below –



Make sure that your service is shown in the installed services list and appropriate servers are selected in the servers tab. Also, verify that the preferred version is set for your cluster and the service is installed in the cluster. Verify that the Cluster attributes for HTTP or HTTPs limits on connections and HTTP or HTTPs traffic limit are set up appropriately to handle load for your service.



7. If your service is in the Installed state, log in to your Collaboration Environment server and run this command: **`deploy_service -lv`**

   The output will look something like this:

   ```
   load: + deploy: + run: + car: N  HelloWorld-0.0.0.1.3
   load: + deploy: + run: + car: +  basicSipCall-1.0.0.0.0
   ```

   Service with call related features:

   If your service is in this list and shows a plus (+) for each step, it is installed and ready to go.

   If the car is “N” for a call related service, make sure that the service name and service version is exactly same in sip.xml as well as web.xml. Also, make sure the sip servlets match in both sip.xml and web.xml.

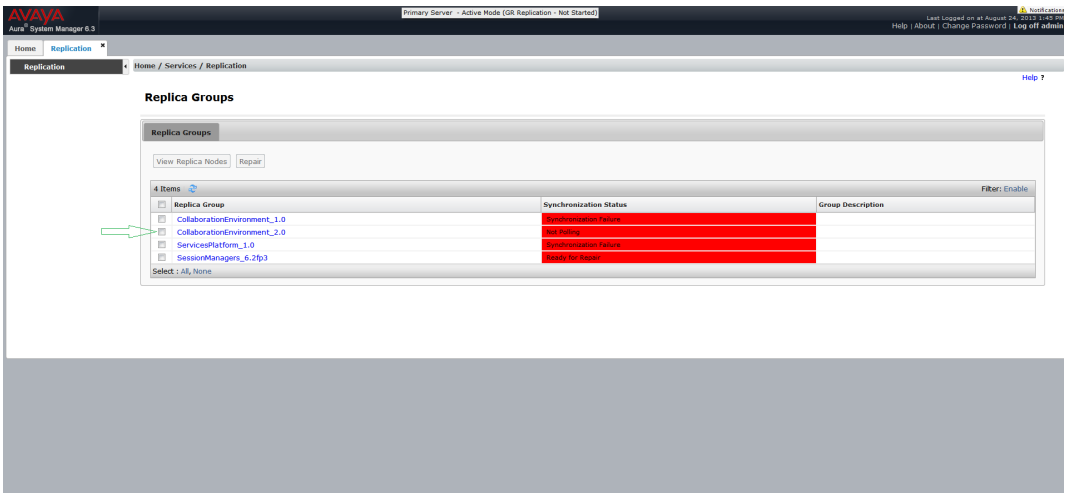
   HTTP service with no call features:

   If your service is in this list and shows a plus (+) for load, deploy and run step, it is installed and ready to go.

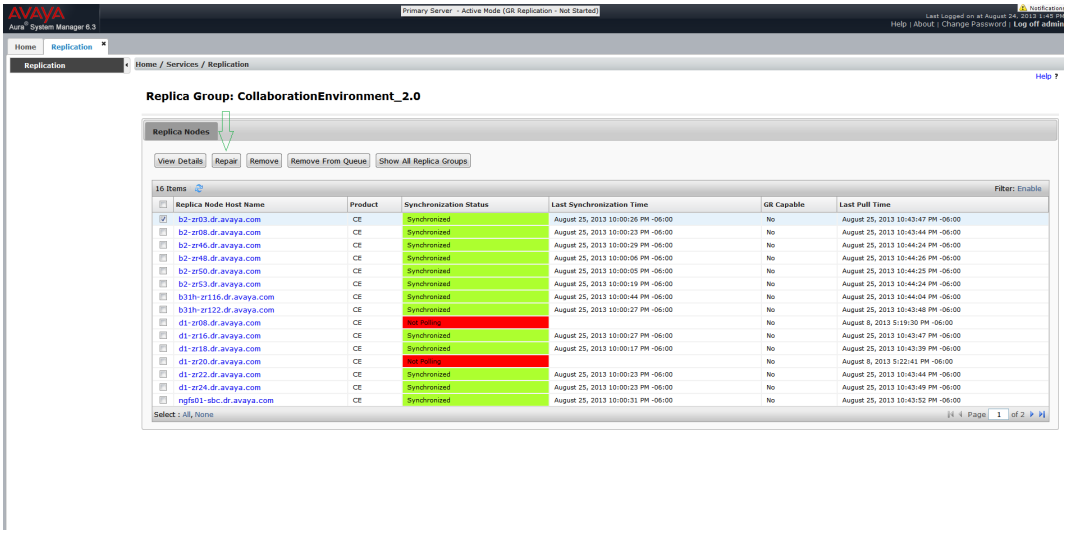
8. Verify that replication is working. From the Home tab on System Manager, select Replication from the Services panel – this will open the Replication tab.

   Look for a Collaboration Environment Replica group in the table. If the Synchronization Status is not “Synchronized”, check the checkbox for the row and press the View Replica Nodes button. It takes about 1 minute a maximum to replicate data once it is changed in System Manager.

   Note that you could press the Repair button here to attempt to repair all servers in the replica group, rather than go to the next page to repair a specific server. Be aware that the Repair operation is service affecting, which means that calls will fail until the operation completes (usually just a few minutes).

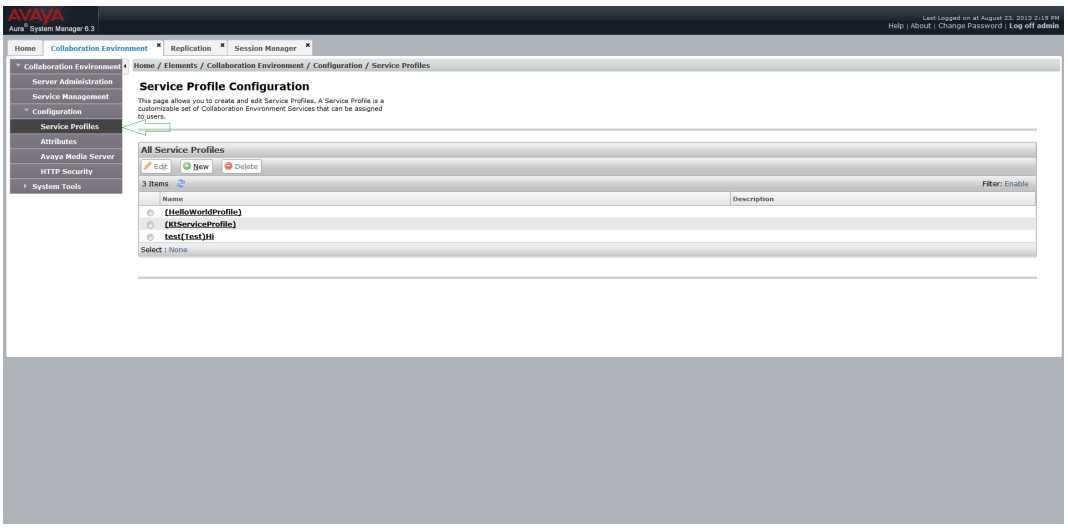


9. If you pressed the View Replica Groups button, you will be in a page that lists the nodes within that Replica Group. Check the checkbox next to your server and press the Repair button. Be aware that the Repair operation is service affecting, which means that calls will fail until the operation completes (usually just a few minutes).

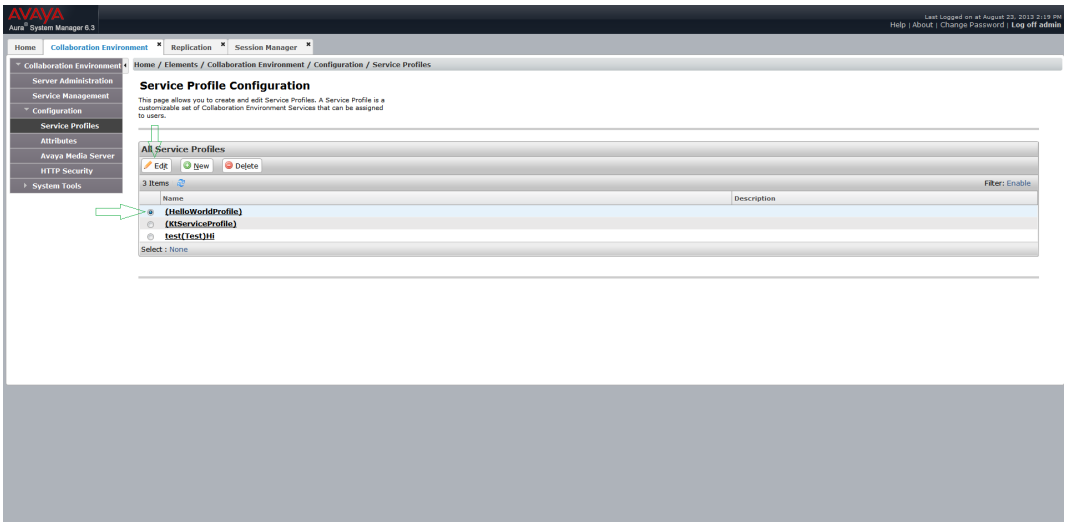


For a Service with Call related features check the following as well:

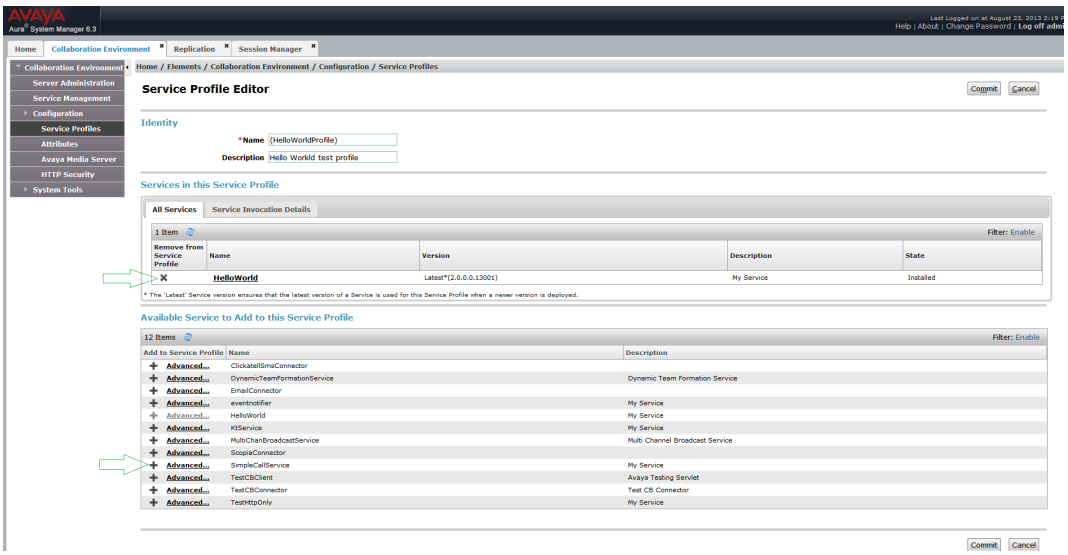
1. Check the service profile definition. From the Home page on System Manager, select Collaboration Environment from the Elements panel.
2. On the Collaboration Environment page, select Configuration, and then Service Profiles.



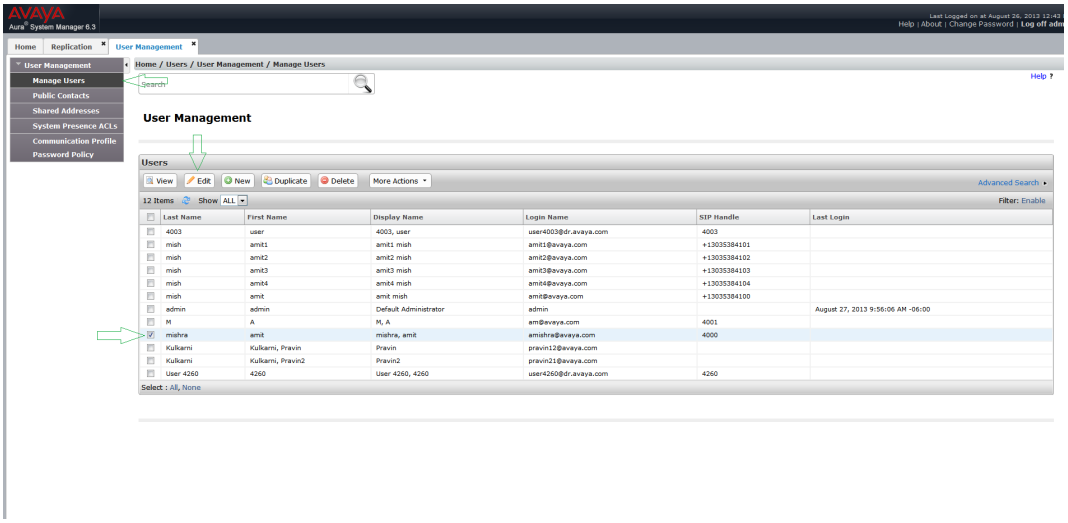
3. If you do not see the profile for your service, press the add button to define a new service. Otherwise, select the radio button next to your service and press the Edit button (as a shortcut, you could simply click the link for your service profile). If you do not recognize any of the service profile names, you may need to view each service profile until you find one that includes your service. You might also find it easier to look at a target user's service profile configuration first (instructions just below here) to find the service profile name.



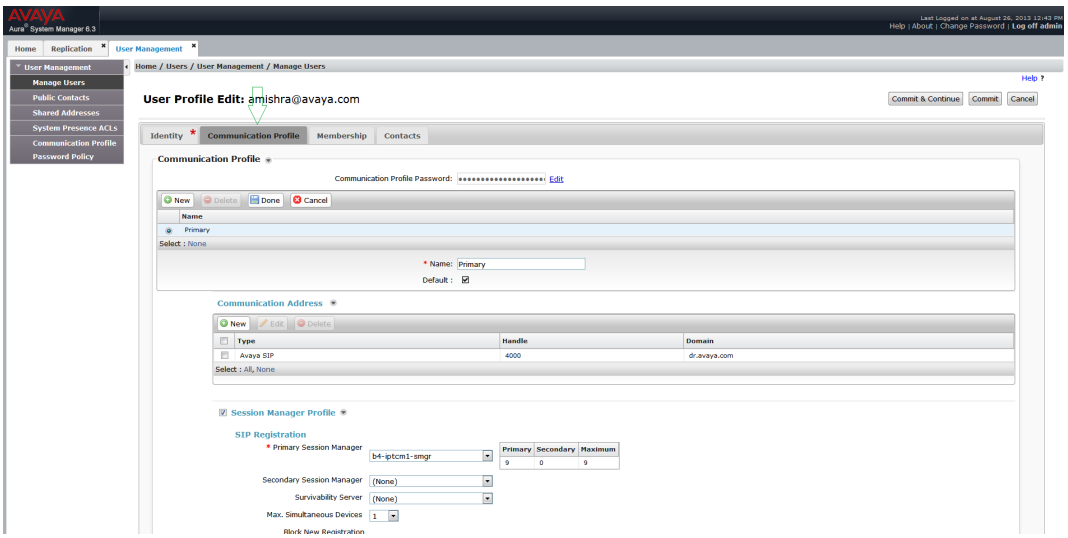
4. The Service Profile Editor page shows the profile at the top, and the list of installed services at the bottom. To add a service to the profile, click the plus (+) icon next to the service in the bottom table; you should then see the service in the profile in the top table. To remove a service, click the X icon next to the service in the top table.



1. Check the user's application sequence and service profile configuration. From the System Manager Home page, select User Management from the Users panel.
2. In the User Management page, select Manage Users from the left menu.
3. Each user to which the service is to be applied must have the service profile assigned to it. Find a user from the list, click the checkbox for that user, then press the Edit button.



4. In the User Profile Edit page, select the Communication Profile tab.



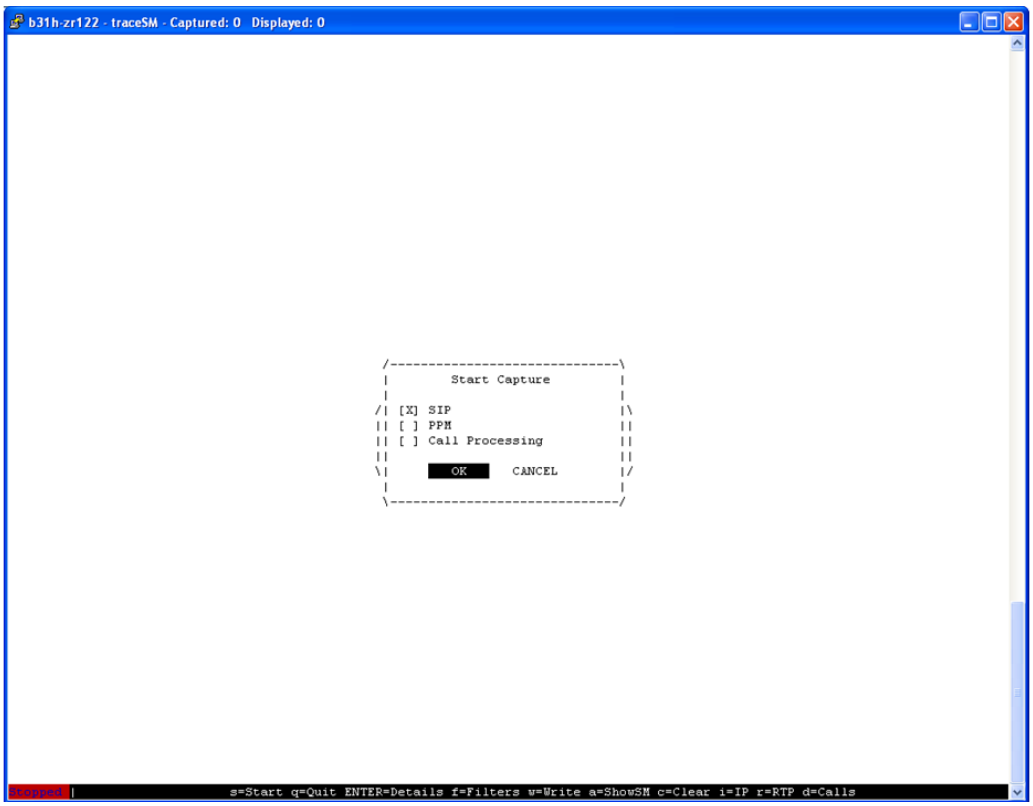
If you made changes to this user that you want to keep, press the Commit button at the top or bottom of the page. Otherwise, press the Cancel button at the top or bottom of the page.

Make sure sequencing is administered on the Session Manager Implicit User Rule Editor page for an individual or group of users. As a best practice, select both sequences, in which case Collaboration Environment is invoked when the user either makes a call or receives a call. At a minimum, if you have a Calling Party service, ensure that there is an Origination Application Sequence that includes Collaboration Environment. If you have a Called Party service, ensure that there is a Termination Application Sequence that includes Collaboration Environment.

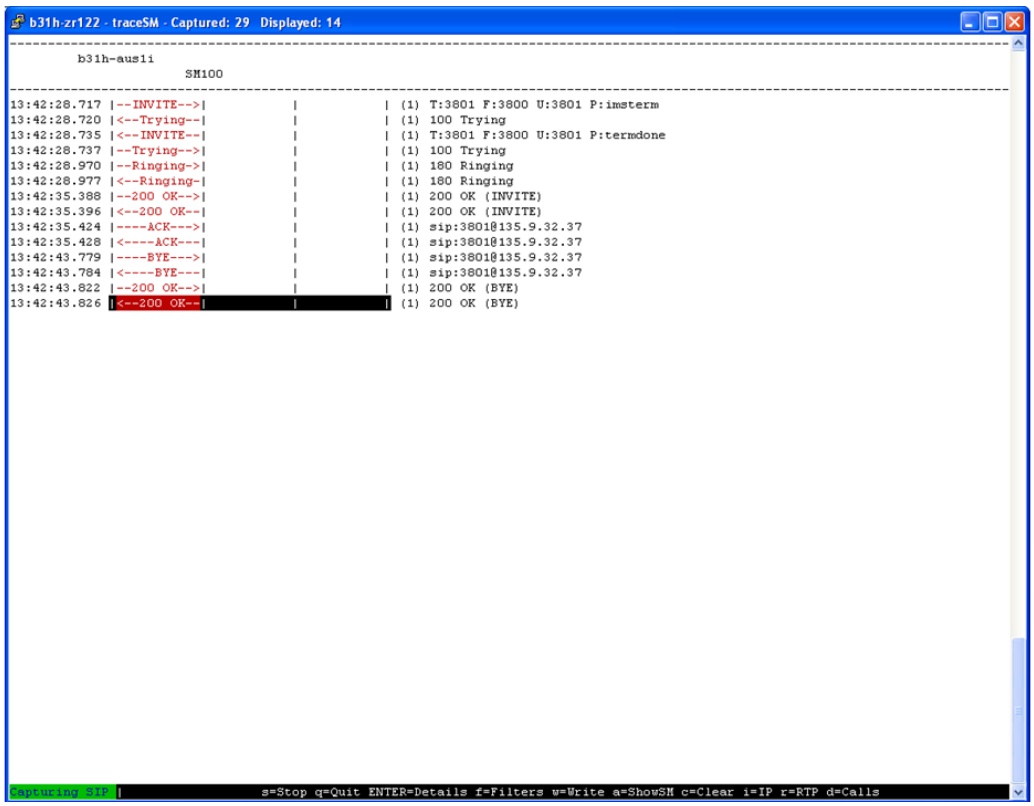
- Navigate to **Home** > **Elements** > **Session Manager** > **Application Configuration** > **Implicit Users**, and click the user or group you want to set up to receive Collaboration Environment services.
- Select your Collaboration Environment Application Sequence for both **Origination Application Sequence** and **Termination Application Sequence**.
- To save your changes, click **Commit**.

For more details, Refer to the document *How to Install, Configure, and Test a Collaboration Environment Service.*

1. Verify that traffic is hitting your Collaboration Environment server. For this edition of Collaboration Environment 2.0, the supported underlying communication protocol is the Session Initiation Protocol (SIP), as defined in IETF RFC 3261 and associated RFCs. Collaboration Environment has a tool to capture and analyze traffic. Run this command: **`traceCE`**
2. When prompted, press “s” to start the capture. You will then be prompted to select the type of traffic – just hit return at this point to take the default.

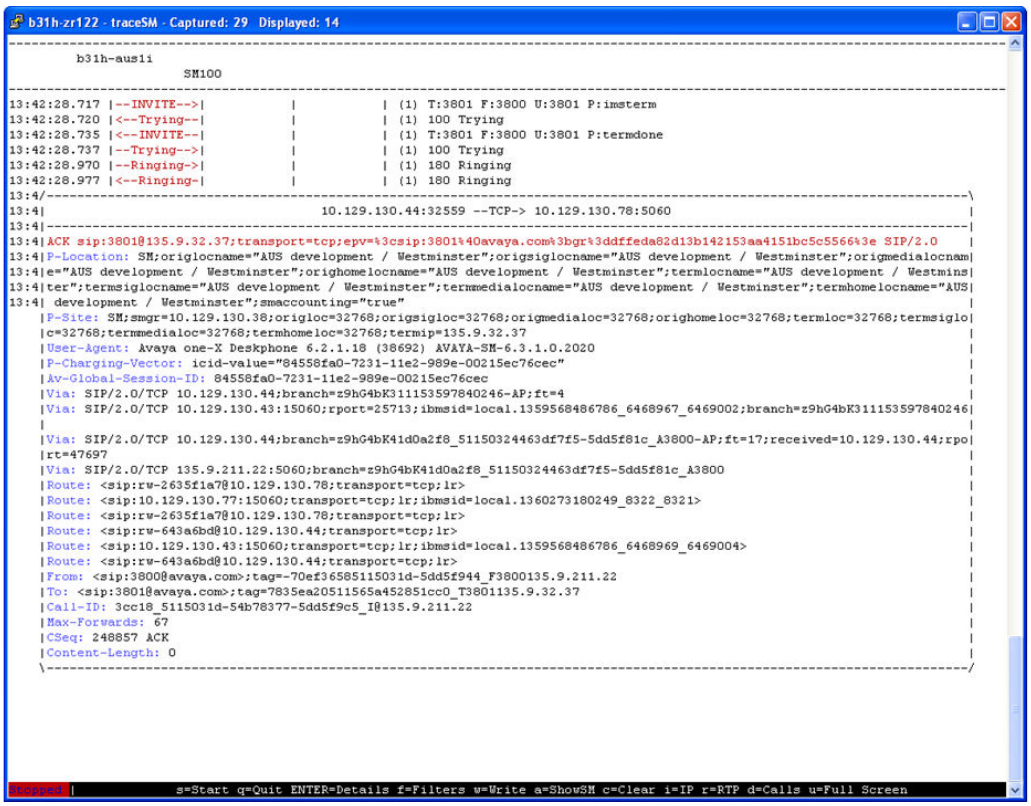


3. Place a call. You should see a fencepost diagram develop that shows the messages exchanged between the Collaboration Environment server and other entities. Here is an example of a simple call, where user A called user B, user B answered and then user A dropped:



Note the menu along the bottom – you may want to stop the trace collection after you have made your test call to limit the amount of data captured.

This trace shows 2 entities involved in the signaling: b31h-ausl1 is a Session Manager, and SM100 represents the Collaboration Environment server. Signaling is usually routed from Session Manager to Collaboration Environment and then back to Session Manager. You can use the arrow keys to move up and down through the message sequence. If you press the enter key, the highlighted message will be expanded.

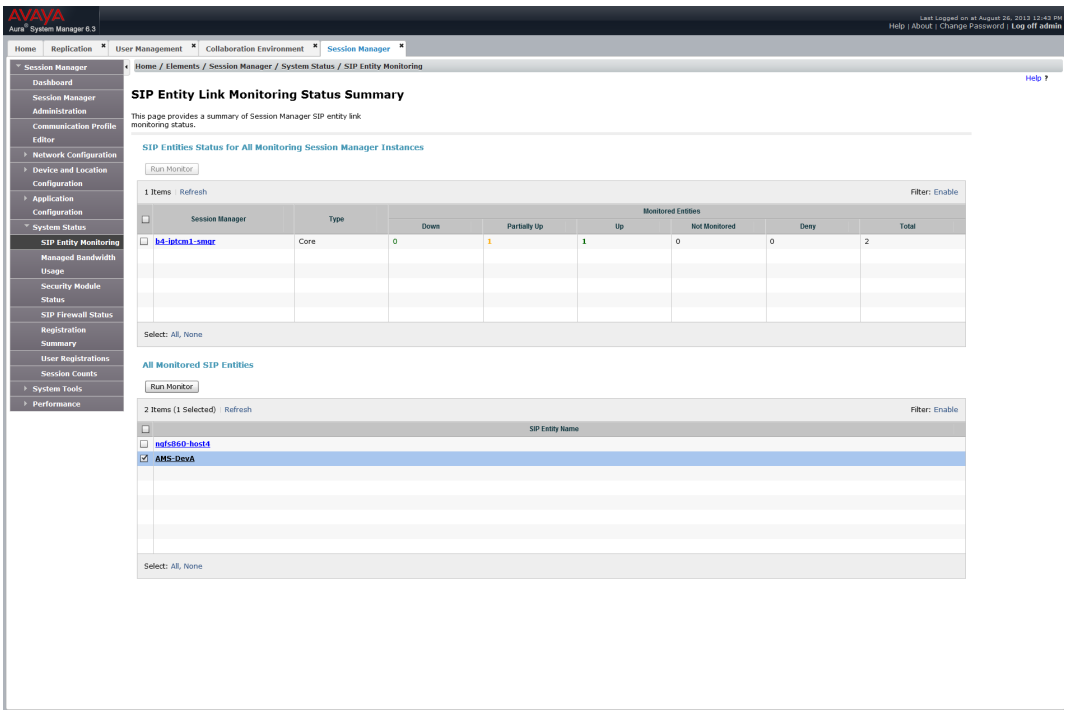


If you start traceCE and do not see any traffic from your test call, and the Origination/ Termination Sequences that we checked above are correct and you placed your Collaboration Environment server into the “Accept New Service” state as described above, then you will need to check other configuration on System Manager. Refer to the document *How to Install, Configure, and Test a Collaboration Environment Service*.

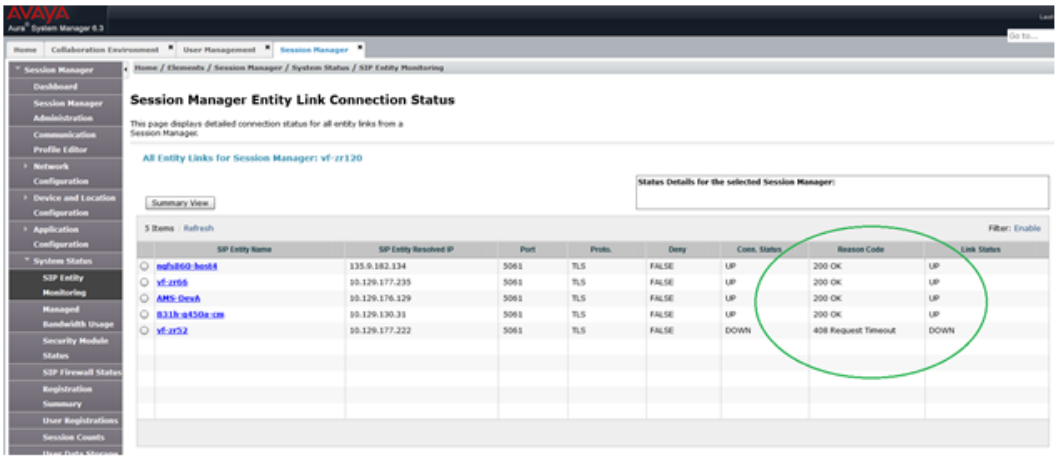
If the trace shows traffic, then check that the users in the call have the correct service profile, as described above.

Both Collaboration Environment and Session Manager monitor the signaling pipes between each other.

1. To check the link monitoring status, on the System Manager Home page select Session Manager from the Elements panel.
2. The Session Manager page appears with the Session Manager Dashboard. Select System Status from the left menu, then SIP Entity Monitoring from the expanded submenu.



If the Session Manager in your environment shows any entities in the Down or Deny state, further inspection is warranted. In this example, the Session Manager named vf-zr120shows an entity link as down. Clicking the link on the name opens a page that shows the status of the individual entity links.



The Reason Code column provides the clue. Refer to Session Manager Documentation for resolution instructions. A tip: a Reason Code of 503 on the link to the Collaboration Environment server probably indicates that the Collaboration Environment server is not in the state to Accept New Service.

More troubleshooting for a Service invoked through HTTP –

- Verify that the URL used to access the service is of the form :

```
http://[CE_FQDN_or_IP]/services/servicename/[Service_Specified_Pattern] or
https://[CE_FQDN_or_IP]/services/servicename/[Service_Specified_Pattern] or
http://[CE_FQDN_or_IP]/services/servicename-serviceversion/
```

```
[Service_Specified_Pattern] or
https://[CE_FQDN_or_IP]/services/servicename-serviceversion/
[Service_Specified_Pattern]
```

- CE FQDN or IP is the Security Module hostname or IP address as configured on System Manager. e.g If you have Multichannel Broadcast Sample service installed on CE, it can be invoked as –

```
http://10.10.10.100/services/MultiChanBroadcastService /
MultiChanBroadcastServlet    or
https://10.10.10.100/services/MultiChanBroadcastService /
MultiChanBroadcastServlet
```

- MultiChanBroadcastService – this is the service name
- The above URLs for Multi-Channel broadcast will invoke the preferred version of the service installed on Collaboration Environment. For example, if you have 2.0.0.0.13002 installed and also 2.0.0.0.13003 on CE, the above example URL will invoke 2.0.0.0.13003 version of the application.
- If you wish to invoke 2.0.0.0.13002 specific version of the example service above, the URL would look like:

```
http://10.10.10.100/services/MultiChanBroadcastService-2.0.0.0.13002/
MultiChanBroadcastServlet     or
https://10.10.10.100/services/MultiChanBroadcastService-2.0.0.0.13002/
MultiChanBroadcastServlet
```

- The preferred version of service can also be invoked using the format of servicename-serviceversion in the URL. i.e

```
http://10.10.10.100/services/MultiChanBroadcastService-2.0.0.0.13003/
MultiChanBroadcastServlet      or
https://10.10.10.100/services/MultiChanBroadcastService-2.0.0.0.13003/
MultiChanBroadcastServlet
```

1. Login to CE server and run this command on command line: **`ce dlogv`**

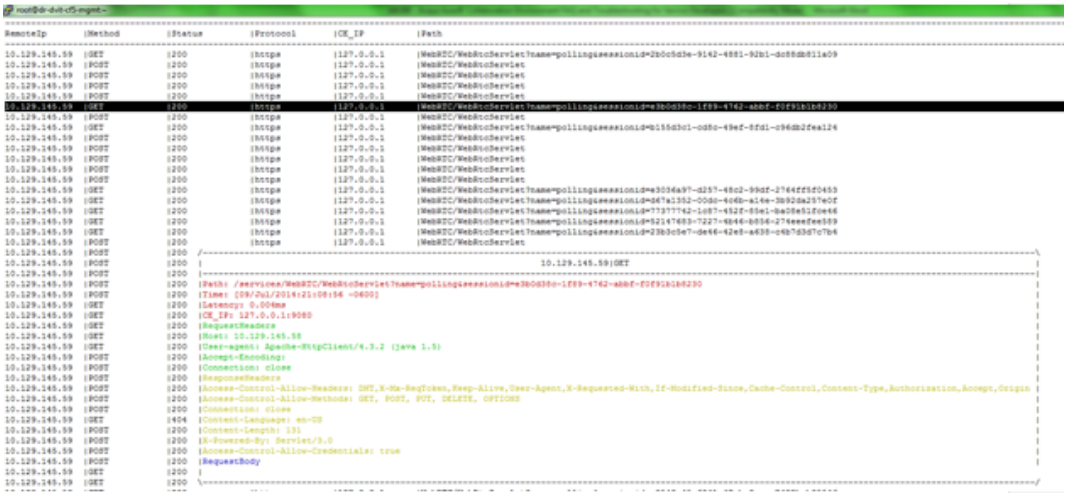
   This will open the debug log using the vi editor. Search for your service name. This should provide you with information if there were any exceptions that lead to issues with invoking the service.

2. If you added logging statements in your service code, and if you used the Logger class in the Collaboration Environment API (highly recommended), you can view the service log using the command **`ce dlogv <serviceName>`**. For example, to view the logs for the basicSipCall service, run the following command - **`ce dlogv basicSipCall`**. The logs for different versions of a service all go to the same log file, and each log message contains the service name and version.

3. When running in normal mode, logs at level INFO and higher priority (FATAL, ERROR, WARN) will be added to the log file; logs at lower priority (FINE, FINER, FINEST) will not be added to the log. The log level can be changed on the fly to provide more verbose logging. Enter this command while logged in to your Collaboration Environment server to enable all log levels: **`ce dlogon`**. To enable all log levels for a specific service use – **`ce dlogon <serviceName>`**

4. When you are finished debugging, you should return logging to the default using this command: **`ce dlogoff`** or **`ce dlogoff <serviceName>`** for a specific service.

5. Default logging (sysout or syserr) goes to a different log. If you are using Java “system.out.println (not recommended)”, the text will be written to the default log. Enter this command to view the log in the vi editor: **`ce alogv`**
6. Look at logs under /var/log/nginx/, if the above errors do not provide sufficient information.
7. You can also run traceHTTP to look at the HTTP requests to your service and responses from your service. Type traceHTTP -> Once you hit enter on command line, you will be able to see the http trace messages. By default the trace is started and it you want to stop press s. Help is available in the tool. You can also see details of a message by selecting the message in the trace.



8. Make sure you are passing all the parameters required to execute your service when invoked through HTTP. In case of Multi-Channel Broadcast Sample Service bundled in SDK, you can look at – MultiChanBroadcastServlet.java class.

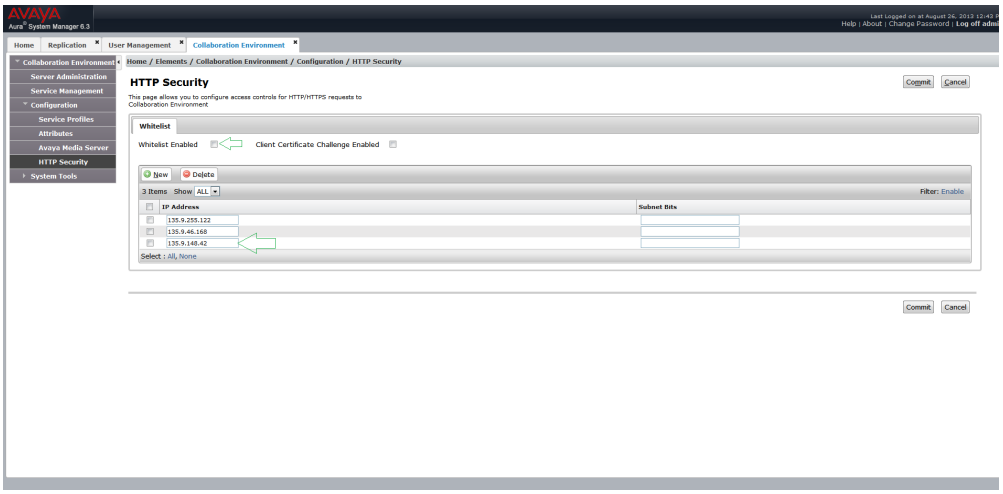
   MultichanbroadcastCollection directory in the sample service provides the Postman collections which is a skeleton for Postman requests to be sent to Multi-Channel Broadcast Sample Service.

   Also look at - “How do I import collections bundled with Sample Services to Postman client in Chrome?”, if you want to know how to import the collections and use them.
9. Make sure Collaboration Environment shows “Accepting” in System State as per the steps mentioned at the start of this question.
10. Check that the HTTP Security either :
    - Does not have Whitelist Enabled.
    - Has Whitelist Enabled and the IP address of the connecting client is in the Whitelist. If the checkbox “Whitelist Enabled” is checked, please make sure the IP address or subnet from which you are trying to connect is in the Whitelist.

      Verify the following if your application depends on CORS (e.g WebRTC does need CORS).

Check that the IP address of the service which serves your JavaScript and tries to make HTTP call to Collaboration Environment is administered in the HTTP CORS page.

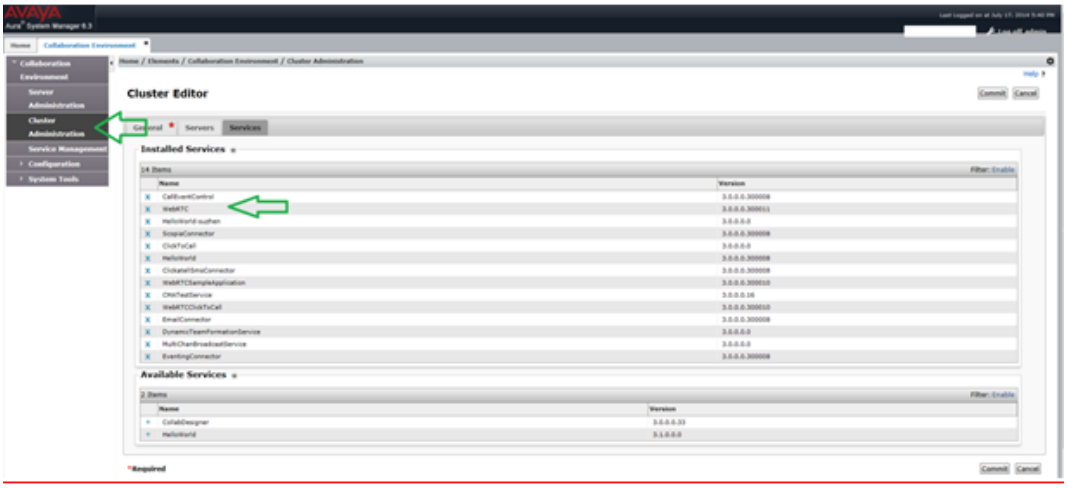
o Alternatively use the checkbox in the screen. But this can have security issues, so please consider those before using this option.



Why are the changes from System Manager to service attributes of my service not reflected during execution or runtime?

Procedure

1. Verify that your service is installed. Log in to System Manager. On the System Manager Home page, select Collaboration Environment from the Elements panel.
2. In the Collaboration Environment page, select Service Management from the left menu.
3. Find your service. If your service does not appear in the table, press the Load button to load the SVAR you created when building your service. If the State for your service is Loaded, select the checkbox and press the Install button. If you are using cluster administration to install the services as in screen below:



4. Make sure that your service is shown in the installed services list and appropriate servers are selected in the servers tab. Also, verify that the preferred version is set for your cluster and the service is installed in the cluster.
5. If your service is in the Installed state, log in to your Collaboration Environment server and run this command: **`deploy_service -lv`**. The output will look something like this:

```
load: + deploy: + run: + car: N  HelloWorld-0.0.0.1.3
load: + deploy: + run: + car: +  basicSipCall-1.0.0.0.0
```

Service with call related features:

If your service is in this list and shows a plus (+) for each step, it is installed and ready to go.

If the car is “N” for a call related service, make sure that the service name and service version is exactly same in sip.xml as well as web.xml. Also, make sure the sip servlets match in both sip.xml and web.xml.

HTTP service with no call features:

If your service is in this list and shows a plus (+) for load, deploy and run step, it is installed and ready to go.

6. Verify that replication is working. From the Home tab on System Manager, select Replication from the Services panel – this will open the Replication tab.

   Look for a Collaboration Environment Replica group in the table. If the Synchronization Status is not “Synchronized”, check the checkbox for the row and press the View Replica Nodes button. It takes about 1 minute a maximum to replicate data once it is changed in System Manager.

   Note that you could press the Repair button here to attempt to repair all servers in the replica group, rather than go to the next page to repair a specific server. Be

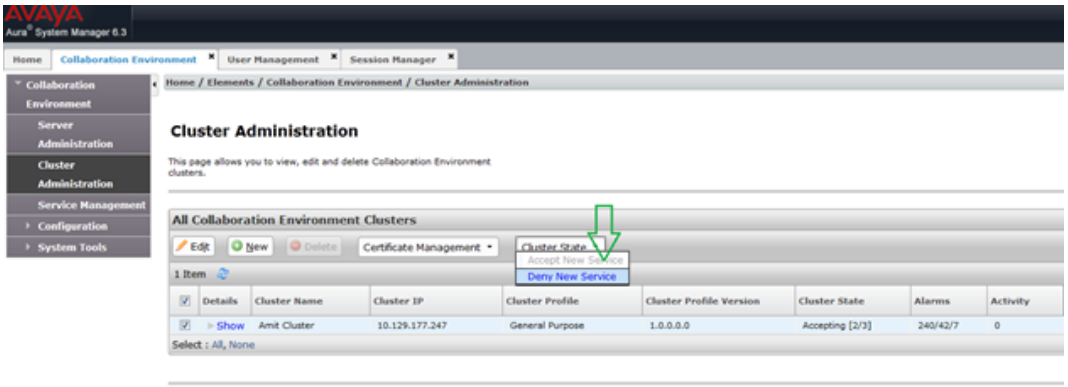
aware that the Repair operation is service affecting, which means that calls will fail until the operation completes (usually just a few minutes).

Why is my service not accessible over http but accessible over https?

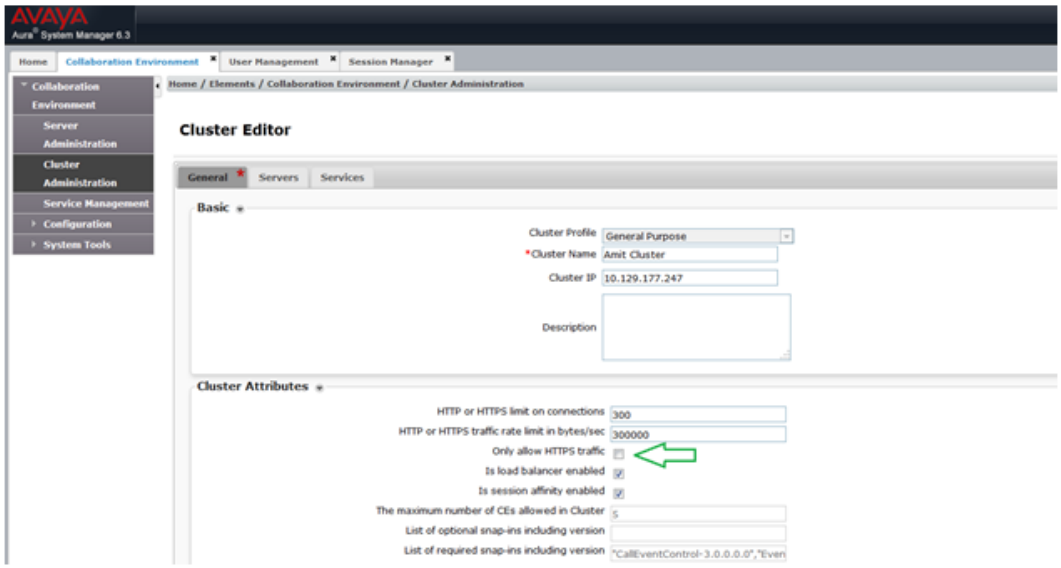
Collaboration Environment has http access disabled by default. Only https is allowed by default. If you wish to have http access as well as https access, from the Home page on System Manager, select Collaboration Environment from the Elements panel.

Procedure

1. On the Collaboration Environment page, select Cluster Administration from left navigation.
2. Put your cluster under Deny New service.



3. And then edit the cluster, uncheck the checkbox “Only Allow HTTPS traffic” and click commit.



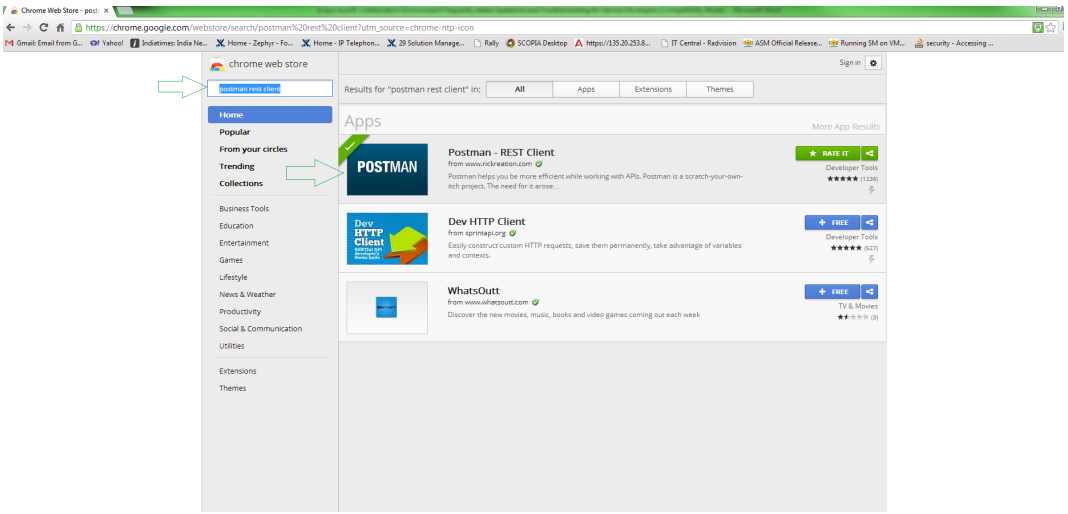
4. Put your cluster back in service. Now you will be able to have both http and https to your service.

How to use postman for https invocation to my service?

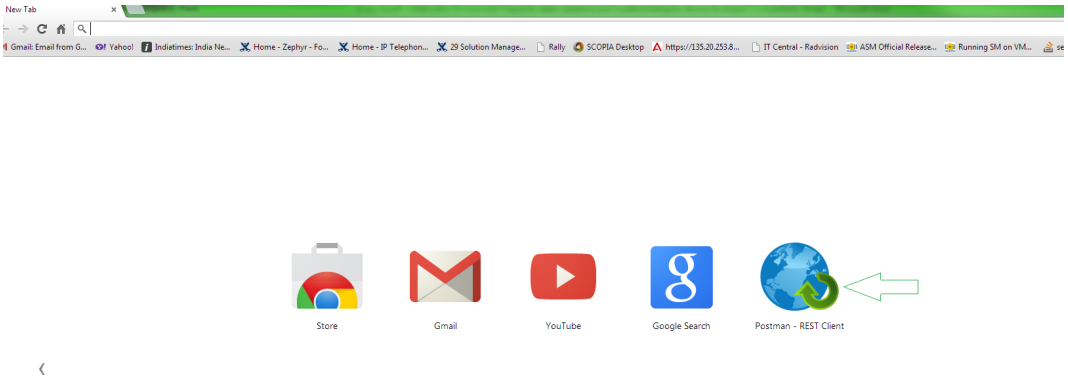
In order to use postman to do https invocation to your service, you need to do accept the certificate of your Collaboration Environment server in Chrome. E.g if you want to execute https to 10.10.10.10 Collaboration Environment server and you are using default certificates – Launch Chrome and execute https://10.10.10.10 as the URL. You will be prompted to accept the certificate. Once accepted, you can use postman to point to your service.

How do I import collections bundled with Sample Services to Postman client in Chrome?

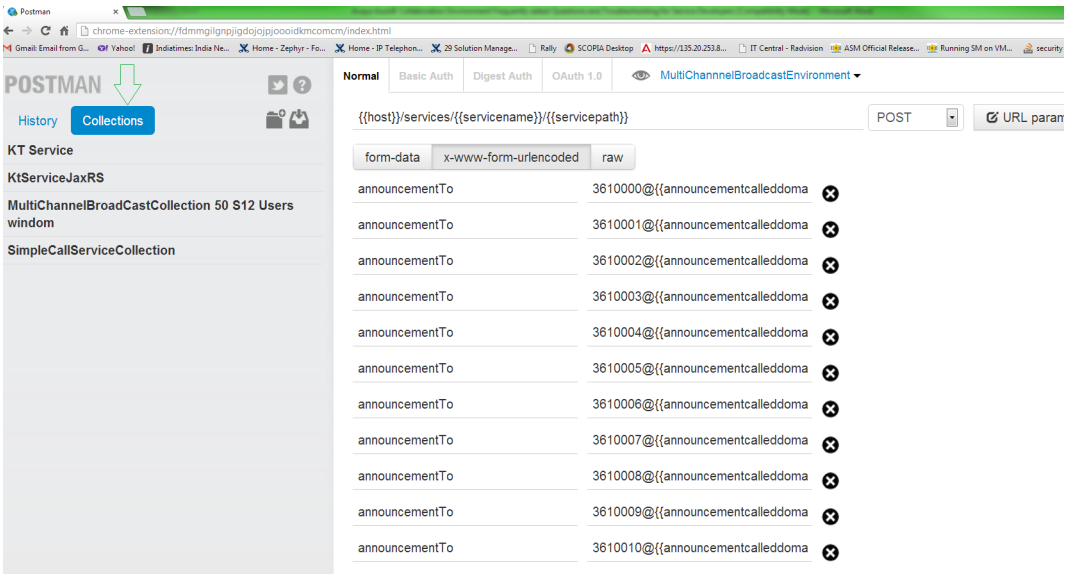
1. First step is to make sure you have Postman Client installed in Chrome. Go to Chrome Web Store and Search for “postman rest client”



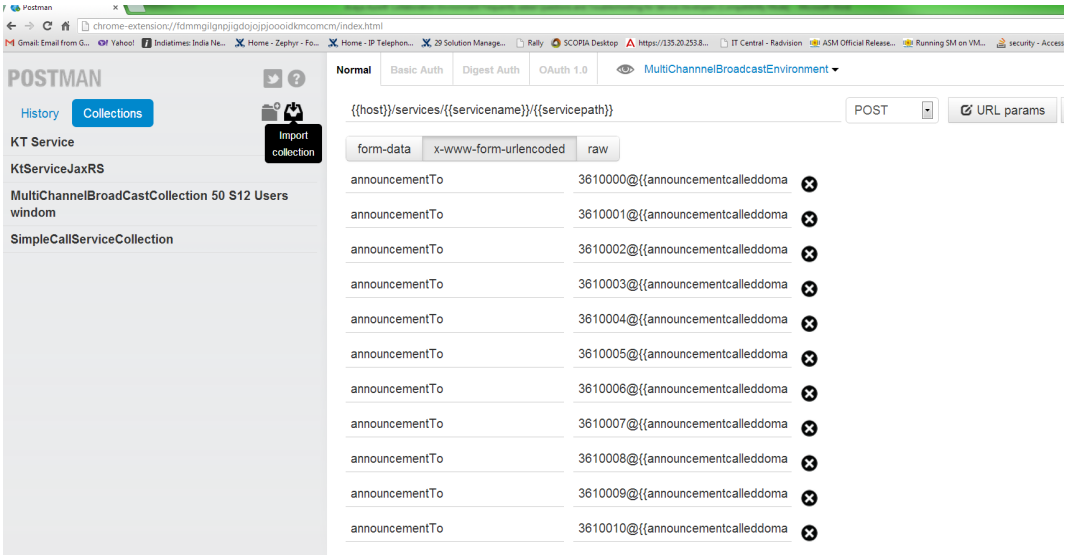
2. Install the client and you should be able to see the client when you launch chrome.



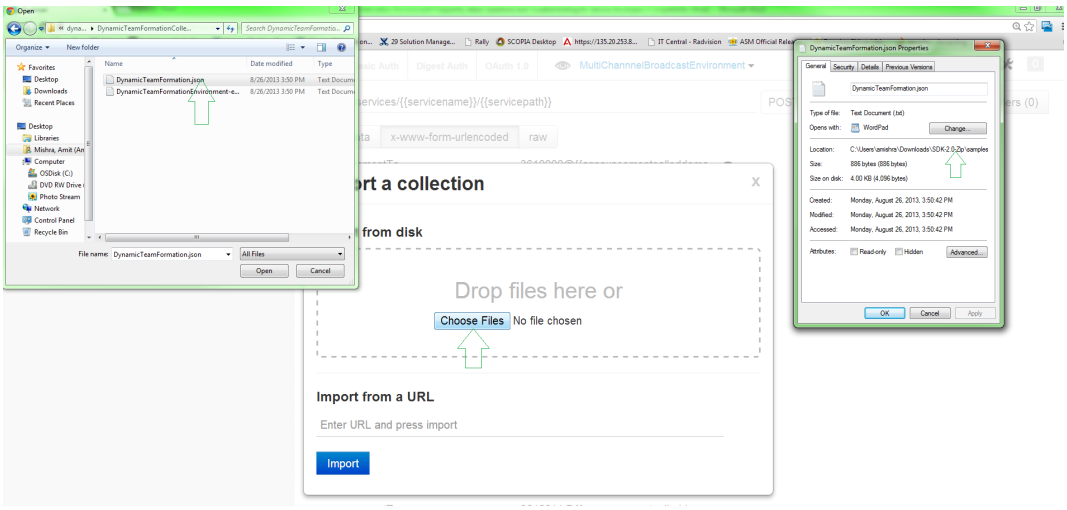
3. Launch the Postman Client and Click on Collections



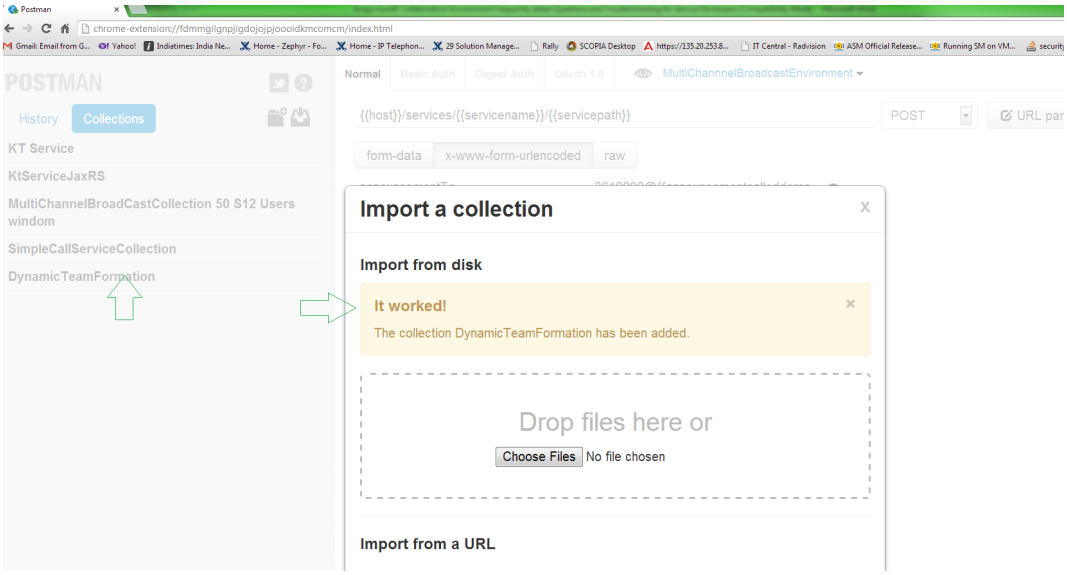
4. Click on Import Collection Button



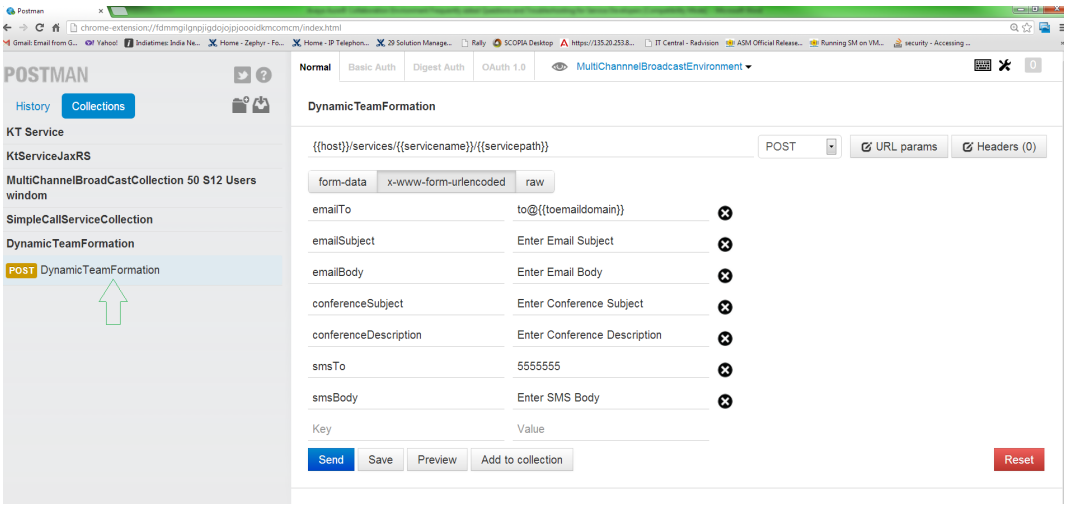
5. Click on Choose files and select the json file from the SDK zip sample service directory. We are taking an example of Dynamic Team Formation Collection here. The collection can be found in following location in the SDK – <SDK-2.0 Folder> \samples\dynamicteamformation\DynamicTeamFormationCollection



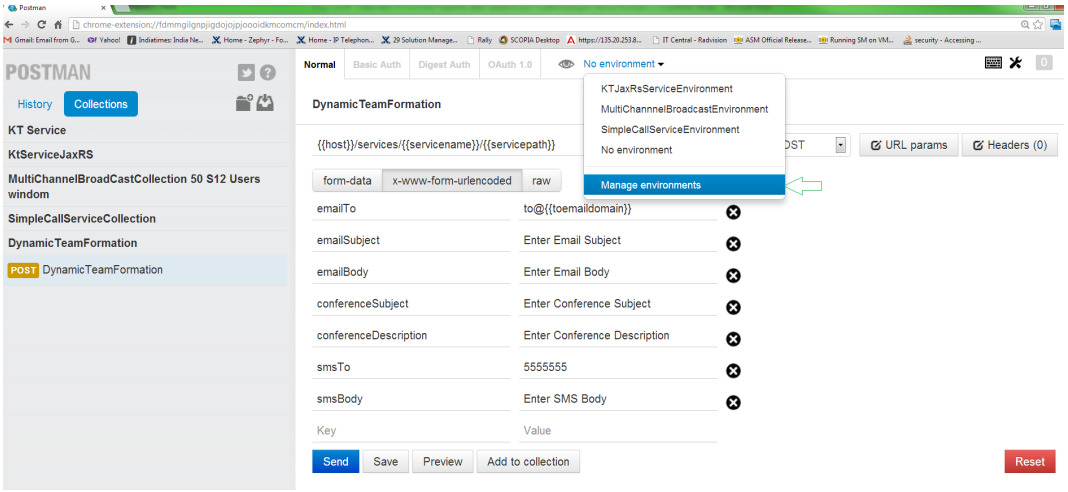
6. You should be able to see a successful import of the collection now.



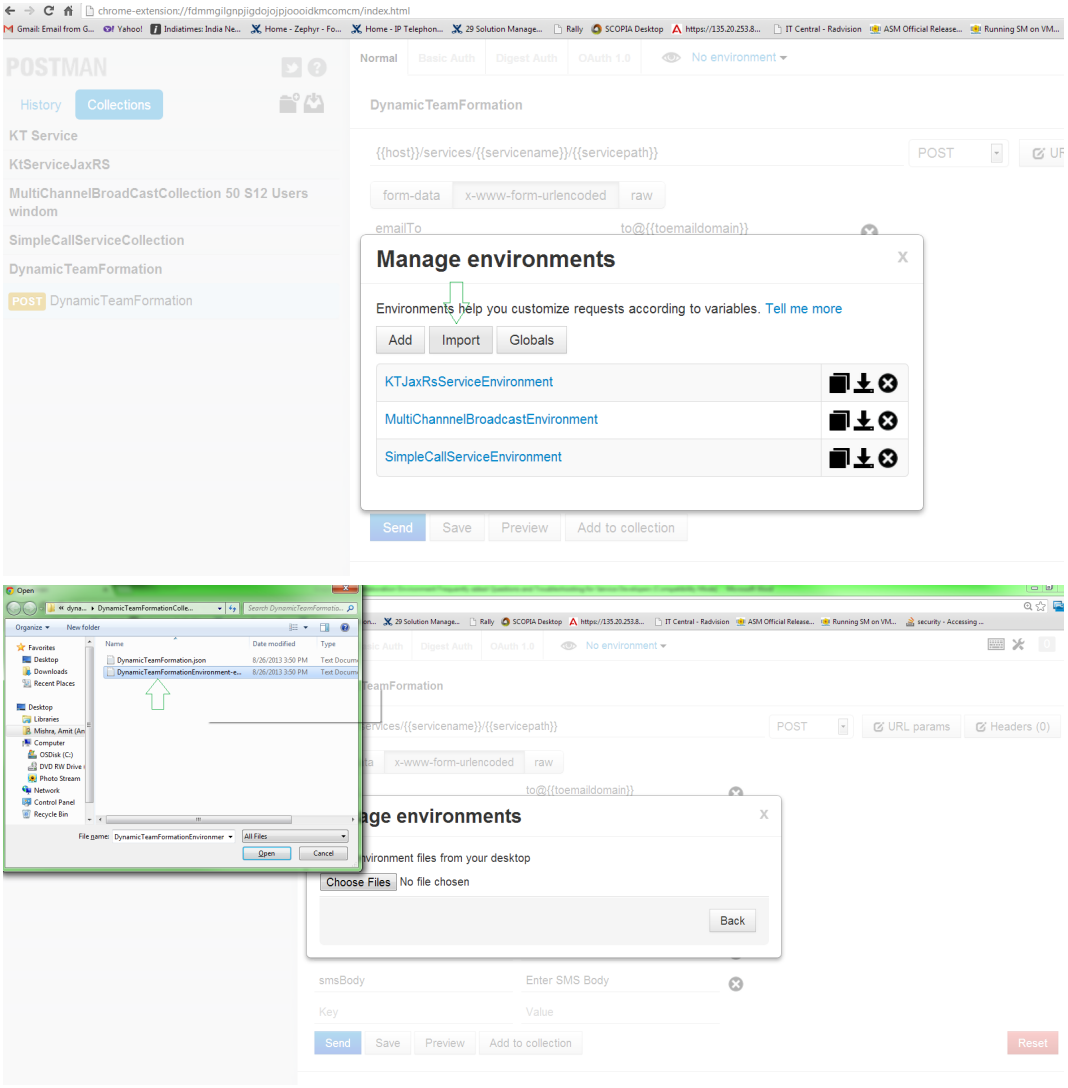
7. Check the collection by selecting the left navigation –



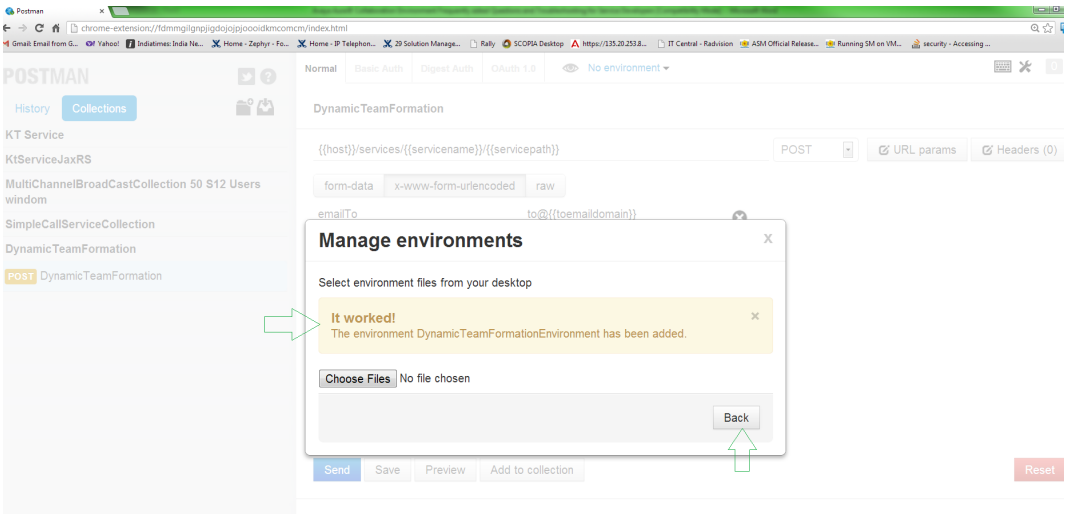
8. Next step is to import the Environment Collection. Select Manage Environment as shown below:



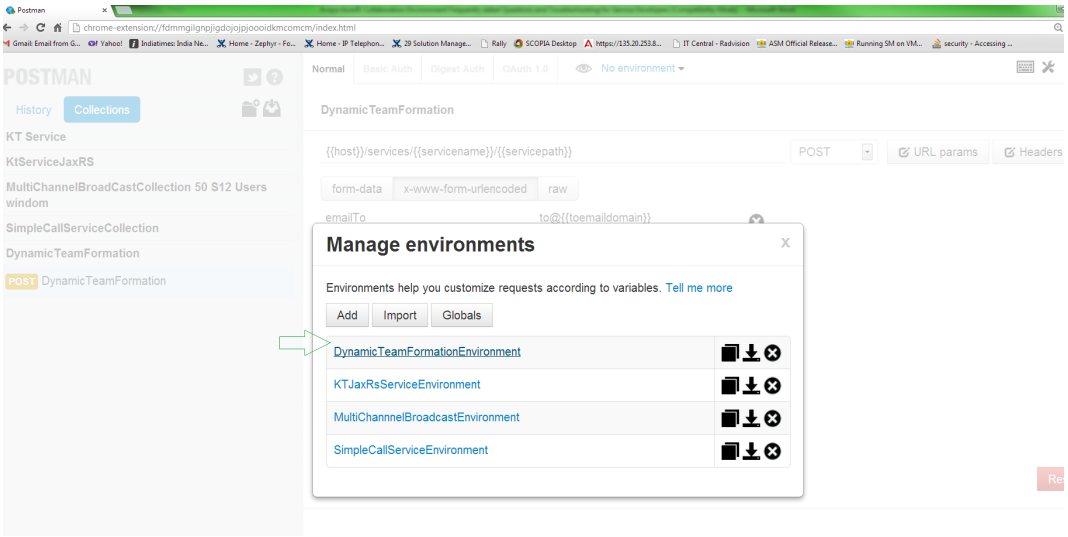
9. Click on import and then click on “Choose Files”:

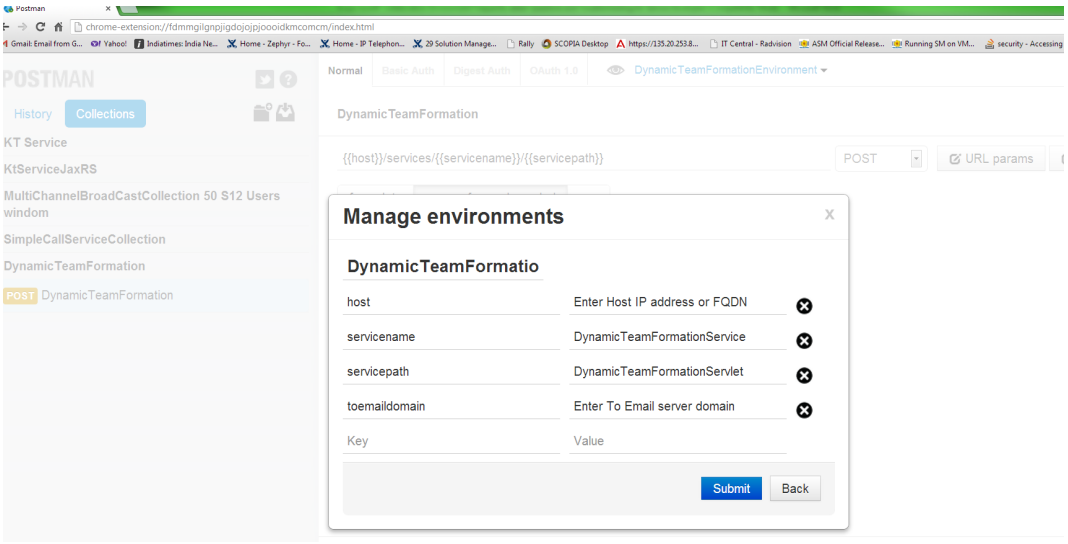


10. You should see it worked -> Click on Back

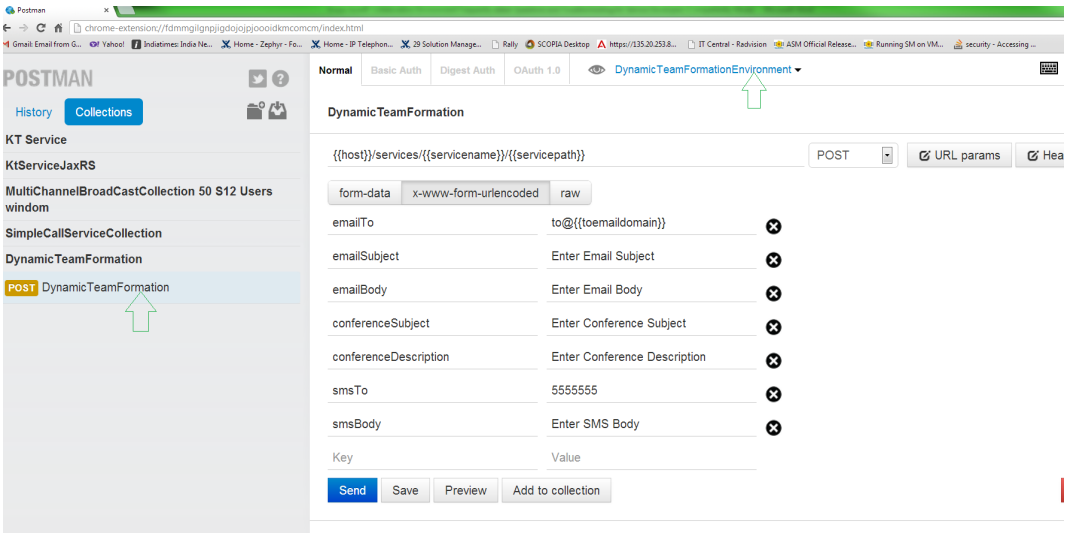


11. Now you can edit the skeleton data as needed to invoke your service by selecting the specific environment –





12. Submit this and select the Environment to be used by your collection



13. Edit the fields in the screen and you are ready to test the Dynamic Team Formation Service assuming that the service is already installed on CE! Click on Send button to see your request sent to the CE.

How do I invoke Multi-Channel Broadcast service and Dynamic Team Formation Service for my testing?

You can write an external application which can send GET and POST HTTP requests to Multi-Channel Broadcast or Dynamic Team formation service to test your service. URL to invoke the service would look like :

Multichannel Broadcast Sample service installed on CE, it can be invoked as –

- `http://10.10.10.100/services/MultiChanBroadcastService/MultiChanBroadcastServlet`

  OR

- `https://10.10.10.100/services/MultiChanBroadcastService/MultiChanBroadcastServlet`

MultiChanBroadcastService – this is the service name

The above URLs for Multi-Channel broadcast will invoke the preferred version of the service installed on Collaboration Environment. For example, if you have 2.0.0.0.13002 installed and also 2.0.0.0.13003 on CE, the above example URL will invoke 2.0.0.0.13003 version of the application.

If you wish to invoke 2.0.0.0.13002 specific version of the example service above, the URL would look like:

- `http://10.10.10.100/services/MultiChanBroadcastService-2.0.0.0.13002/MultiChanBroadcastServlet`

  OR

- `https://10.10.10.100/services/MultiChanBroadcastService-2.0.0.0.13002/MultiChanBroadcastServlet`

The preferred version of service can also be invoked using the format of servicename-serviceversion in the URL. i.e

- `http://10.10.10.100/services/MultiChanBroadcastService-2.0.0.0.13003/MultiChanBroadcastServlet`

  OR

- `https://10.10.10.100/services/MultiChanBroadcastService-2.0.0.0.13003/MultiChanBroadcastServlet`

Dynamic Team formation Service on Collaboration Environment can be invoked as –

- `http://10.10.10.100/services/DynamicTeamFormationService/DynamicTeamFormationServlet`

  OR

- `https://10.10.10.100/services/DynamicTeamFormationService/DynamicTeamFormationServlet`

Alternatively, you can use Postman client and use the Collections provided with the sample services to send the requests to the services. Details on how to import the Collections can be found in the previous question “How do I import collections bundled with Sample Services to Postman client in Chrome?”

More details about the sample services can be found in documents *Multi-Channel Broadcast Sample Service* and *Dynamic Team Formation Sample Service*.

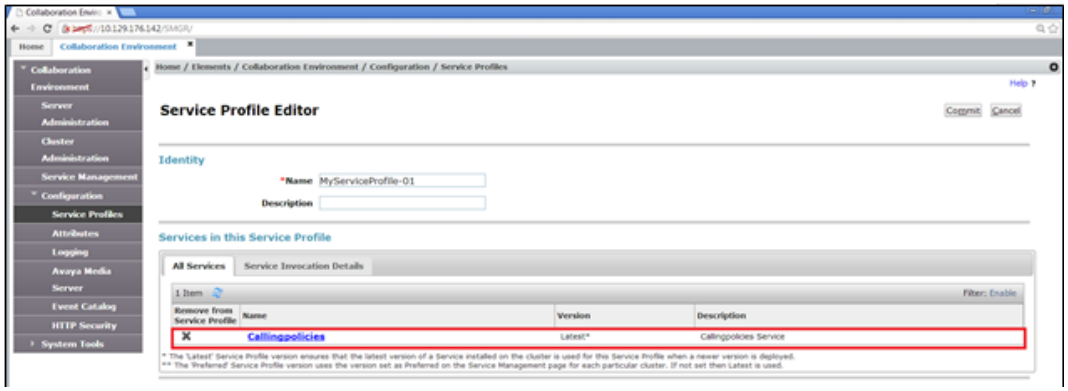
How do I invoke ClickToCall Sample service for my testing?

Open a well-known browser and type: http://<CE IP>/services/ClickToCall/index.jsp. For example: http://10.129.177.214/services/ClickToCall/index.jsp.

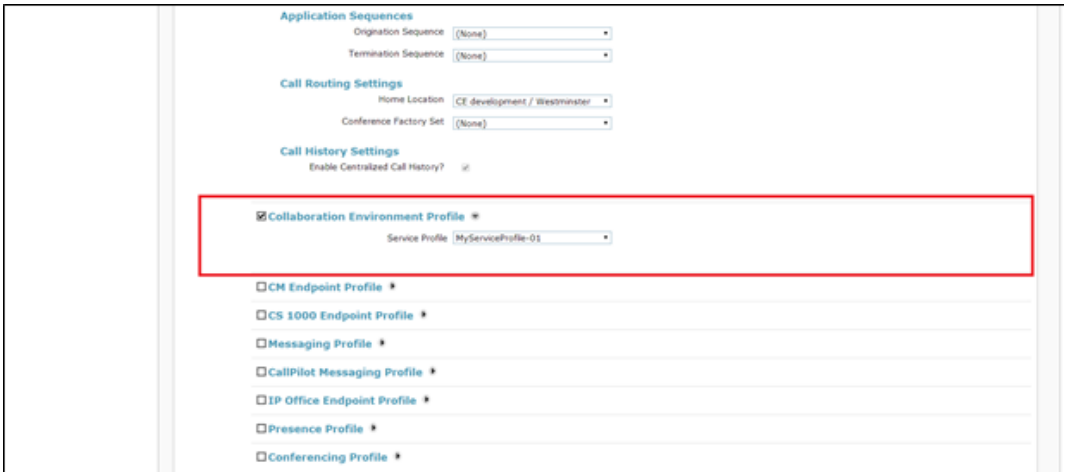
What are the pre-requisites for the Callingpolicies Sample service?

Procedure

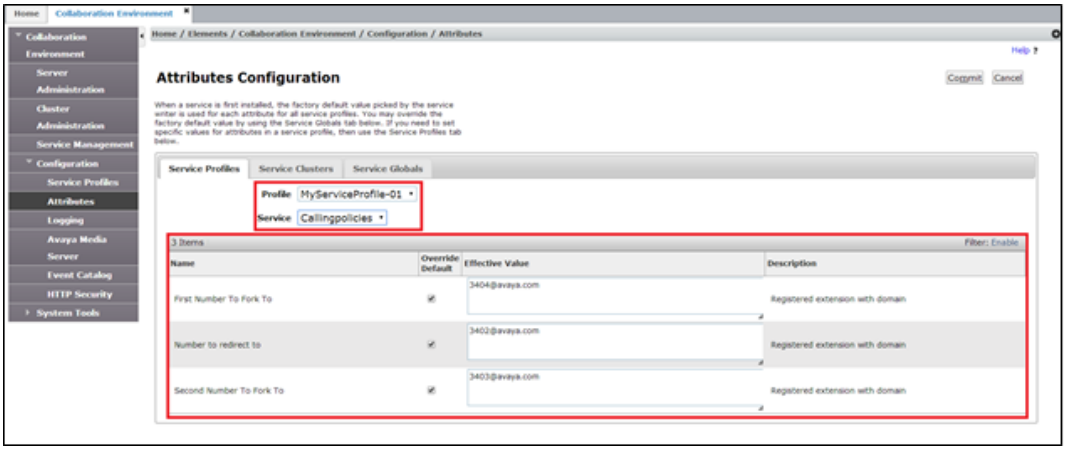
1. Configure the service profile and add the Callingpolicies preferred version to the Service profile:



2. Add the service profile to the selected extension:



3. Under Configuration Attributes, please set the Service profile attributes by selecting the Profile and the Service:



How do I test the Callingpolicies Sample service?

- Since this is a term side called service, call a number example: 3401@avaya.com from another extension. This would be the called extension.
- Called extension rings.
- The announcement is played.
- Based on the dtmf input , call is either forked or, redirected or, allowed to continue or dropped.

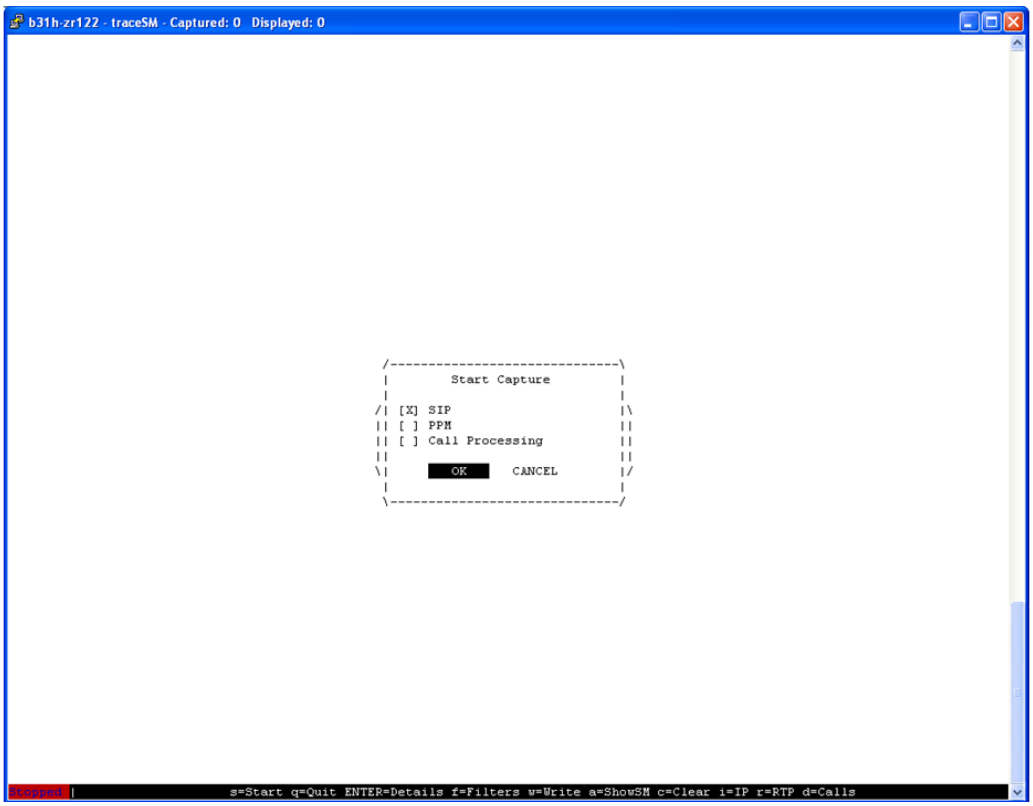
- For 1= Fork, it forks to the extensions provided in the service attributes. Please refer the screen shot above.
- For 2 = redirect, it redirects to the extension provided in the service attribute. Please refer the screen shot above.
- For 3= allow, the called extension (in the above example 3401@avaya.com) rings.
- For 4=drop, the call is dropped and the respective tone is heard on the extension.

My service is installed and I see logs from my service, but things are not behaving like I expected. What should I check?

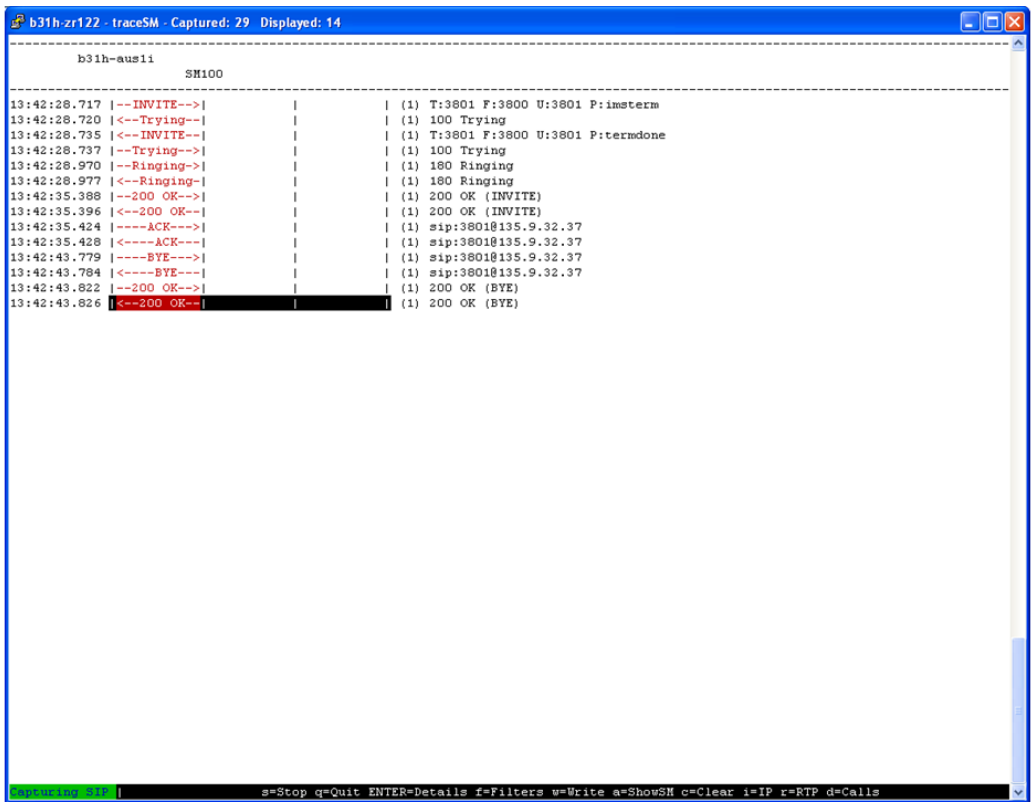
If your application makes use of collaboration bus and think that messages are not passed correctly, you can run a traceBus command to check the messages to and from the collaboration bus.

For this edition of Collaboration Environment, the supported underlying communication protocol is the Session Initiation Protocol (SIP) for call intercept/call control services, as defined in IETF RFC 3261 and associated RFCs. Collaboration Environment has a tool to capture and analyze traffic.

1. Run this command on the Collaboration Environment server: **`traceCE`**
2. When prompted, press “s” to start the capture. You will then be prompted to select the type of traffic – just hit return at this point to take the default.

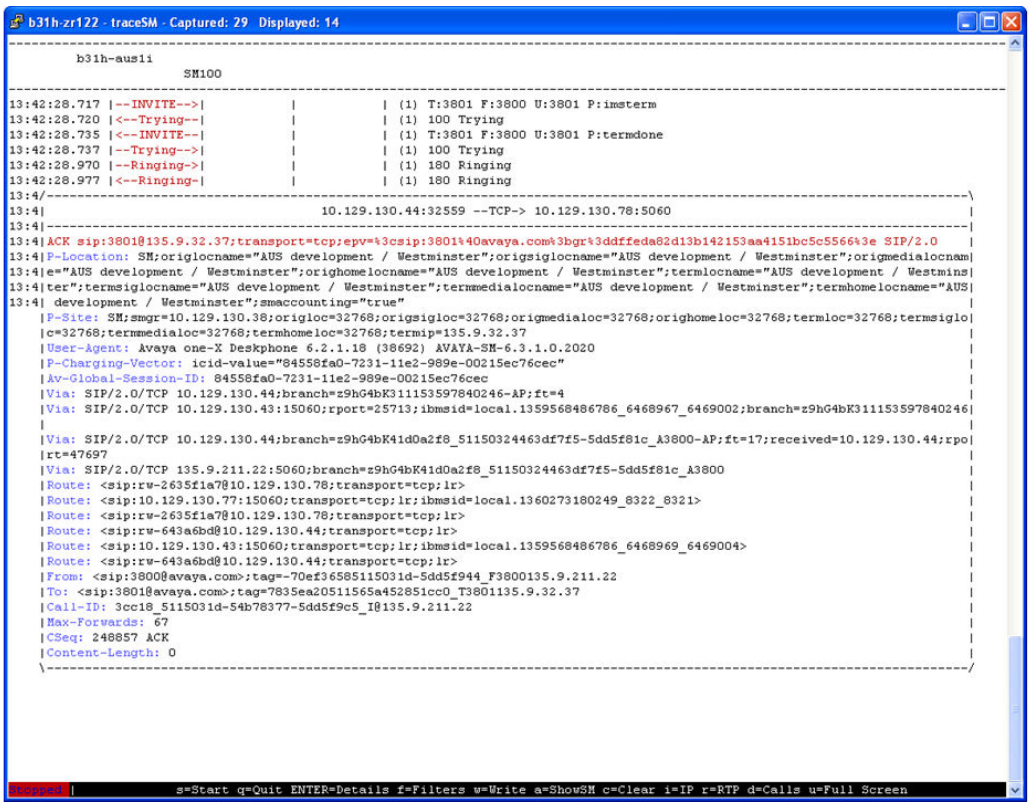


3. Place a call. You should see a fencepost diagram develop that shows the messages exchanged between the Collaboration Environment server and other entities. Here is an example of a simple call, where user A called user B, user B answered and then user A dropped:



Note the menu along the bottom – you may want to stop the trace collection after you have made your test call to limit the amount of data captured.

4. This trace shows 2 entities involved in the signaling: b31h-ausl1 is a Session Manager, and SM100 represents the Collaboration Environment server. Signaling is generally routed from Session Manager to Collaboration Environment and then back to Session Manager. You can use the arrow keys to move up and down through the message sequence. If you press the enter key, the highlighted message will be expanded.



You can inspect the various messages to see if your service has any effect.

- • You can get a view of application status on Collaboration Environment -> Server Administration Screen by looking at the activity column.
- You can get a view of application status by running this command on your Collaboration Environment server: **`metget =sipapp`**

  The output will look something like this:

```
0
SipContainerModule.ApplicationModule.SessionModule.WhiteList-1.0.0.0.0.Active_
SIP_Sessions
1
SipContainerModule.ApplicationModule.SessionModule.basicSipCall-1.0.0.0.0.Acti
ve_SIP_Sessions
2 SipContainerModule.ApplicationModule.SessionModule.pfa.Active_SIP_Sessions
0
SipContainerModule.ApplicationModule.SessionModule.testingLogger-0.0.0.0.1.Act
ive_SIP_Sessions
3 SipContainerModule.Active_SIP_Sessions
```

  The clue here is that the basicSipCall service shows 1 active session.

- Sometimes debugging can be a bit easier by starting with simpler configurations or simpler scenarios. For example, you might want to start with all signaling going over TCP rather than TLS. At a minimum, this would allow you to use network sniffers like wireshark

to analyze traffic at different points in your infrastructure. After resolving your problem, you could try switching back to TCP.

- Make sure you use same protocol across all your systems -> Collaboration Environment, Session Manager, Avaya Media Server and Communication manager. If you use TCP, use TCP on all these servers. If TLS, then configure TLS on all of these servers.

For Service invoked through HTTP –

- Verify that the URL used to access the service is of the form :

```
http://[CE_FQDN_or_IP ]/services/servicename/[Service_Specified_Pattern] or
https://[CE_FQDN_or_IP]/services/servicename/[Service_Specified_Pattern] or
http://[CE_FQDN_or_IP ]/services/servicename-serviceversion/
[Service_Specified_Pattern] or
https://[CE_FQDN_or_IP]/services/servicename-serviceversion/
[Service_Specified_Pattern]
```

CE FQDN or IP is the Security Module hostname or IP address as configured on System Manager. e.g If you have Multichannel Broadcast Sample service installed on CE, it can be invoked as –

http://10.10.10.100/services/MultiChanBroadcastService/MultiChanBroadcastServlet or https://10.10.10.100/services/MultiChanBroadcastService/MultiChanBroadcastServlet

MultiChanBroadcastService – this is the service name

The above URLs for Multi-Channel broadcast will invoke the preferred version of the service installed on CE. For example, if you have 3.0.0.0.13002 installed and also 3.0.0.0.13003 on Collaboration Environment, the above example URL will invoke 3.0.0.0.13003 version of the application.

If you wish to invoke 3.0.0.0.13002 specific version of the example service above, the URL would look like: http://10.10.10.100/services/MultiChanBroadcastService-3.0.0.0.13002/MultiChanBroadcastServlet or https://10.10.10.100/services/MultiChanBroadcastService-3.0.0.0.13002/MultiChanBroadcastServlet

The latest version of service can also be invoked using the format of servicename-serviceversion in the URL. i.e http://10.10.10.100/services/MultiChanBroadcastService-3.0.0.0.13003/MultiChanBroadcastServlet or https://10.10.10.100/services/MultiChanBroadcastService-3.0.0.0.13003/MultiChanBroadcastServlet

- Login to CE server and run this command on command line: **`ce dlogv`**

This will open the debug log using the vi editor. Search for your service name. This should provide you with information if there were any exceptions that lead to issues with invoking the service.

If you added logging statements in your service code, and if you used the Logger class in the Collaboration Environment API (highly recommended), you can view the service log using the command **`ce dlogv <serviceName>`**. For example, to view the logs for the basicSipCall service, run the following command - **`ce dlogv basicSipCall`**.

The logs for different versions of a service all go to the same log file, and each log message contains the service name and version.

When running in normal mode, logs at level INFO and higher priority (FATAL, ERROR, WARN) will be added to the log file; logs at lower priority (FINE, FINER, FINEST) will not

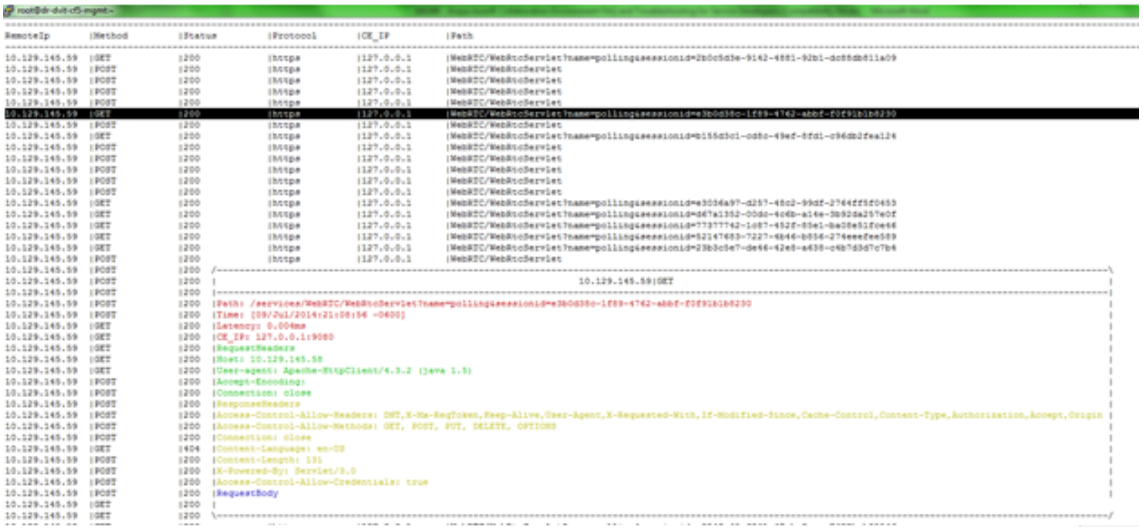
be added to the log. The log level can be changed on the fly to provide more verbose logging. Enter this command while logged in to your Collaboration Environment server to enable all log levels: **`ce dlogon`**

To enable all log levels for a specific service use run **`ce dlogon <serviceName>`**. When you are finished debugging, you should return logging to the default using this command: **`ce dlogoff`** or **`ce dlogoff <serviceName>`** for a specific service.

When you are finished debugging, you should return logging to the default using this command: **`ce dlogoff`**

Default logging (sysout or syserr) goes to a different log. If you are using Java “system.out.println (not recommended)”, the text will be written to the default log. Enter this command to view the log in the vi editor: **`ce alogv`**

- Look at logs under /var/log/nginx/, if the above errors do not provide sufficient information.
- You can also run traceHTTP to look at the HTTP requests to your service and responses from your service. Type traceHTTP -> Once you hit enter on command line, you will be able to see the http trace messages. By default the trace is started and it you want to stop press s. Help is available in the tool. You can also see details of a message by selecting the message in the trace.



- Make sure you are passing all the parameters required to execute your service when invoked through HTTP. In case of Multi-Channel Broadcast Sample Service bundled in SDK, you can look at – MultiChanBroadcastServlet.java class.

  MultichanbroadcastCollection directory in the sample service provides the Postman collections which is a skeleton for Postman requests to be sent to Multi-Channel Broadcast Sample Service.

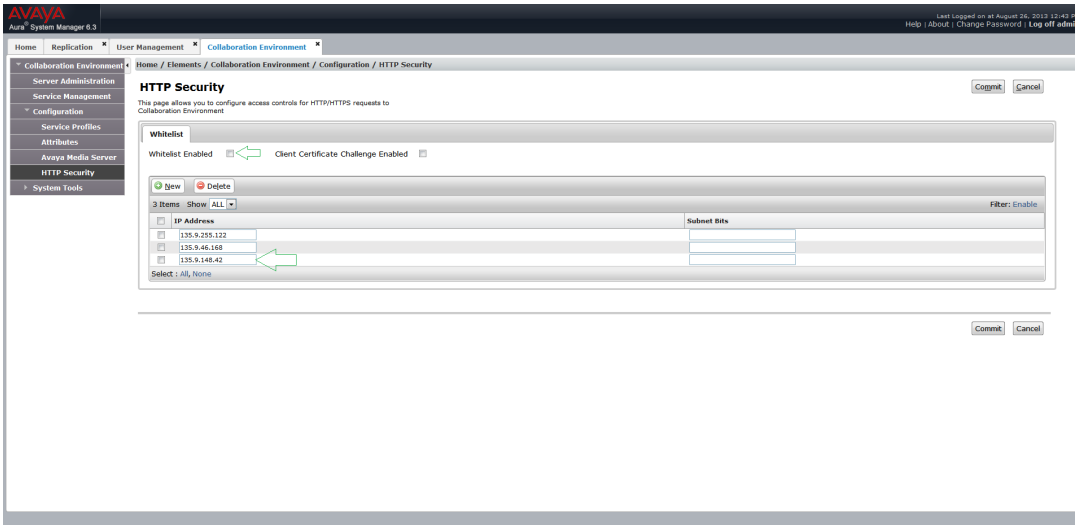
  Also look at - “How do I import collections bundled with Sample Services to Postman client in Chrome?”, if you want to know how to import the collections and use them.

- Make sure Collaboration Environment show “Accepting” in System State as per the steps mentioned at the start of this question.
- Check that the HTTP Security either :
  - Does not have Whitelist Enabled.
  - Has Whitelist Enabled and the IP address of the connecting client is in the Whitelist. If the checkbox “Whitelist Enabled” is checked, please make sure the IP address or subnet from which you are trying to connect is in the Whitelist.

    Verify the following if your application depends on CORS (e.g WebRTC does need CORS).

    Check that the IP address of the service which serves your JavaScript and tries to make HTTP call to Collaboration Environment is administered in the HTTP CORS page.

    o Alternatively use the checkbox in the screen. But this can have security issues, so please consider those before using this option.



Where do I look for the debug logs I added to my service?

Assuming you added logging statement in your service code using the Logger class in the Collaboration Environment API (highly recommended), the debug logs are text files. The default location for log files for a particular service on Collaboration Environment is /var/log/Avaya/services/<servicenameThe logs for different versions of a service all go to the same log file, and each log message contains the service name and version.>.

Log in to the Collaboration Environment server and run this command: **`ce dlogv <serviceName>`**. This opens the service log using the vi editor. Take a look at the Javadoc for the Logger class in the Collaboration Environment API. When running in normal mode, logs at level INFO and higher priority (FATAL, ERROR, WARN) will be added to the log file; logs at lower priority (FINE, FINER, FINEST) will not be added to the log.

The log level can be changed on the fly to provide more verbose logging. Enter this command while logged in to your Collaboration Environment server to enable all log levels: **`ce dlogon <serviceName>`**.

When you are finished debugging, you should return logging to the default using this command: **`ce dlogoff <serviceName>`**.

Default logging (sysout or syserr) goes to a different log. If you are using Java "System.out.println (not recommended)", the text will be written to the default log. Enter this command to view the log in the vi editor: **`ce alogv`**.

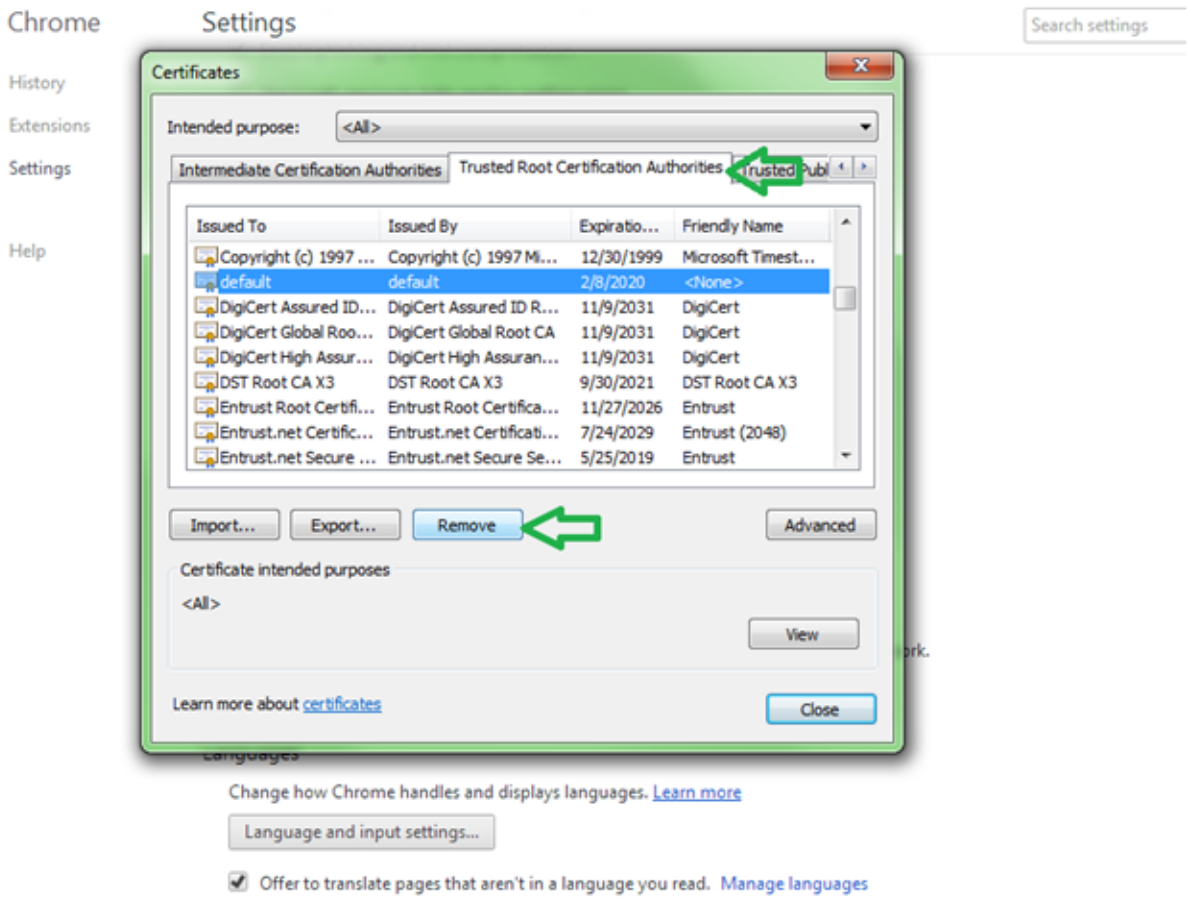
Why does my outgoing HTTP request from my service not use the security module IP?

By default the outgoing HTTP requests from Collaboration environment use the management IP address. In-order to make sure security module IP address is used in outgoing HTTP requests, please look at sample code from Javadoc. `Package: com.avaya.collaboration.bus.sample Class: SampleLotteryHttpClient`

If you would like to make https requests from you snap-in look at sample code in Javadoc for SSLUtility APIs. Please note that this does not use proxy configuration by default. Write the service such that you can use a proxy in your snap-in while making the http request.

Why does HTTPS access to my services give error like "Cannot connect to real <IP address>"?

This happens if you are using default certificates on your Collaboration Environment. The default certificate would not be valid. You can clear the default cert and then access your application over https again. In order to remove the incorrect certificate, launch chrome as "administrator". Go to Settings -> Show Advanced Settings -> Manage Certificates -> trusted Root Certification Authorities. Now select the Default certificate and remove.



Close the browser and launch again. Now access your service using HTTPS. You would be able to accept the certificate and proceed with testing your snap-in.

How do I debug media-related issues?

Make sure the administration is done as per the *Deploying the Avaya Aura® Collaboration Environment* document. Look specifically at the chapter with details about Avaya Media Server Deployment

Do you see an exception in the logs – “mediaService play throws illegalStateException”? This indicates that an unsupported scenario is being executed using the Avaya Media Server. The scenarios in which the operations are supported are:

- An Outbound Calling service that plays announcements for notification purposes
- A Call Intercept service that 1) Plays a message then 2) Drops the call.

Unsupported scenarios include:

- A Call Intercept service that invokes Call.allowCall() then attempts to invoke MediaService methods.
- A Call Intercept service that invokes MediaService methods then attempts to invoke Call.allowCall() or Call.divertTo().

Refer to the details about developing applications with announcements from the *Collaboration Environment Service Development Guide*

Refer to the Sample Service – Multi-Channel Broadcast sample service (which is part of the SDK) on how announcements can be played for an outgoing call.

Login to CE server and run this command on command line: **`ce dlogv <serviceName>`**. This will open the debug log using the vi editor. Search for your service name. This should provide you with information if there were any exceptions that lead to issues with invoking the service.

If you added logging statements in your service code, and if you used the Logger class in the Collaboration Environment API (highly recommended), you can view the service log using the command **`ce dlogv <serviceName>`**. For example, to view the logs for the basicSipCall service, run the following command - **`ce dlogv basicSipCall`**.

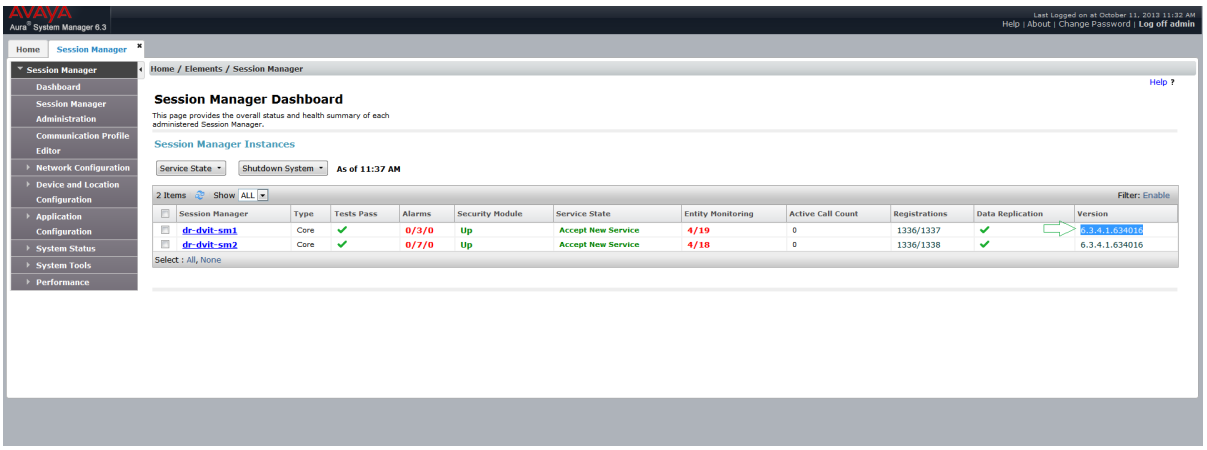
. When running in normal mode, logs at level INFO and higher priority (FATAL, ERROR, WARN) will be added to the log file; logs at lower priority (FINE, FINER, FINEST) will not be added to the log. The log level can be changed on the fly to provide more verbose logging. Enter this command while logged in to your Collaboration Environment server to enable all log levels: **`ce dlogon <serviceName>`**. When you are finished debugging, you should return logging to the default using this command: **`ce dlogoff <serviceName>`**. Default logging (sysout or syserr) goes to a different log. If you are using Java “system.out.println (not recommended)”, the text will be written to the default log. Enter this command to view the log in the vi editor: **`ce alogv`**.

In order to debug and look for ce platform logs, enable logging by: **`ce dlogon`**. Then look at server by running : **`ce dlogv`**. Once you are done disable logging by executing : **`ce dlogoff`**.

I've set up an implicit sequence pattern so that SM will route calls to the Collaboration Environment, but calls are not being sent to the Collaboration Environment. What might cause this?

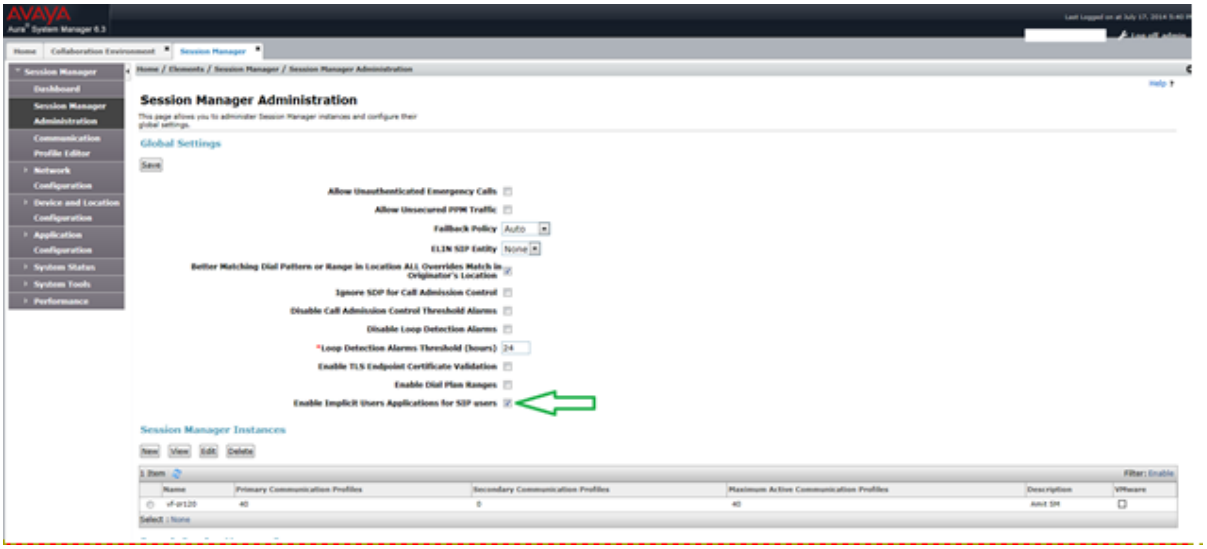
While using Collaboration Environment 3.0 load with Avaya Aura 6.2 Feature Pack 3 GA load, you must install Session Manager Patch for “Implicit/Explicit sequencing” (asm-patch-6.3.4.1.634016).

Login to System Manager and go to Home -> Session Manager, the version column for your session manager should show as 6.3.4.1.634016.



We recommend using Avaya Aura Feature Pack 4 GA load to avoid this confusion.

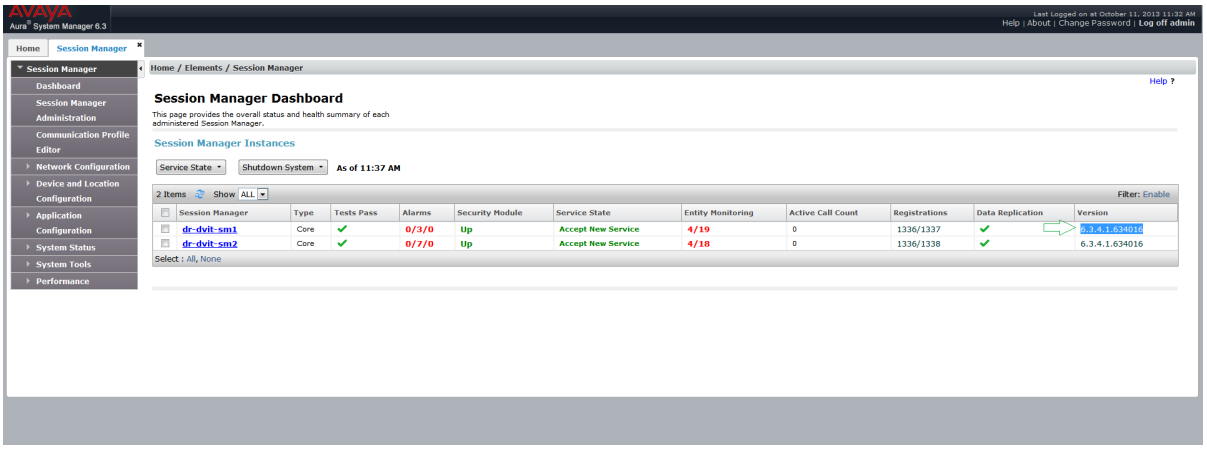
Verify that the “Enable Implicit Users Applications for SIP users” configuration is selected in session manager administration screen.



I've set up an implicit sequence pattern so that SM will route calls to the Collaboration Environment, but some of the calls that should be matching my pattern are not getting sent to the Collaboration Environment. What might cause this?

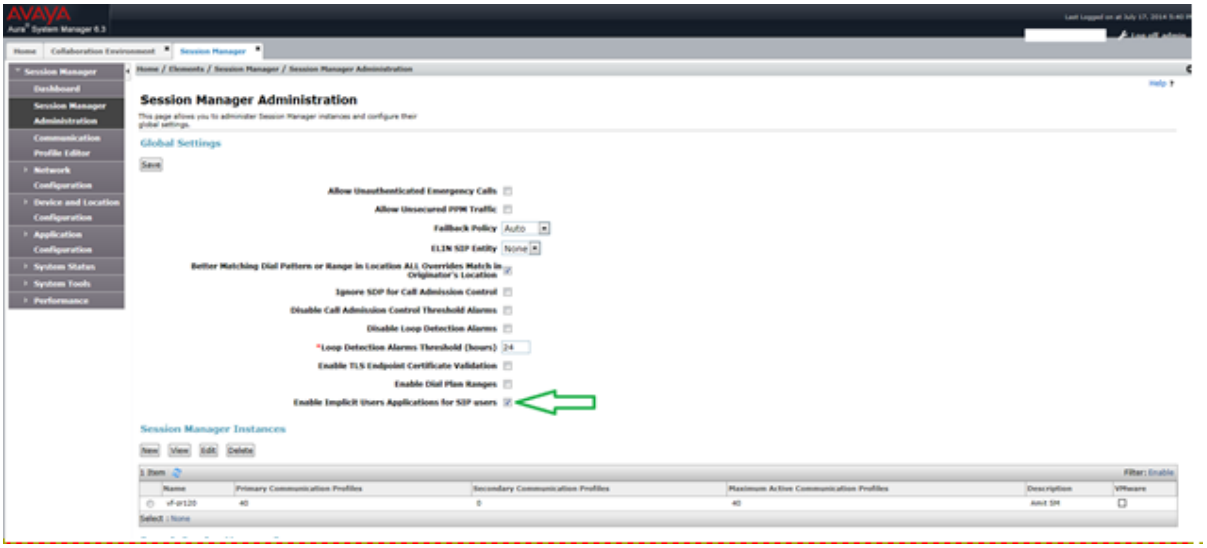
While using Collaboration Environment 3.0 load with Avaya Aura 6.2 Feature Pack 3 GA load, you must install Session Manager Patch for “Implicit/Explicit sequencing” (asm-patch-6.3.4.1.634016).

Login to System Manager and go to Home -> Session Manager, the version column for your session manager should show as 6.3.4.1.634016.



We recommend using Avaya Aura Feature Pack 4 GA load to avoid this confusion

Verify that the “Enable Implicit Users Applications for SIP users” configuration is selected in session manager administration screen.



How can I stop Collaboration Environment from logging me out?

This is actually a security feature that will log you out after a few minutes of inactivity.

While logged in to your Collaboration Environment server, you can disable the timer by entering this command: **`unset TMOUT`**

Or, you can set it to a longer interval with this command: **export TMOUT=<number of seconds>**

Why don't I see my service alarm on System Manager or Data Center?

- Verify Collaboration Environment Serviceability Agent's configuration for User and Target profiles. If not configured properly, reconfigure correctly.
- Verify Serviceability agent is running on Collaboration Environment instance using command > service spiritAgent status. If service is down start the service.
- Verify the event is logged in /var/log/Avaya/services/ event.log file. If event is not logged check server logs for exceptions.

Why is access to my service so slow or sometimes not accessible?

Assuming that you have already looked at other FAQs like – "Why is my service not accessible", "Why is my service not getting invoked", you may want to check following:

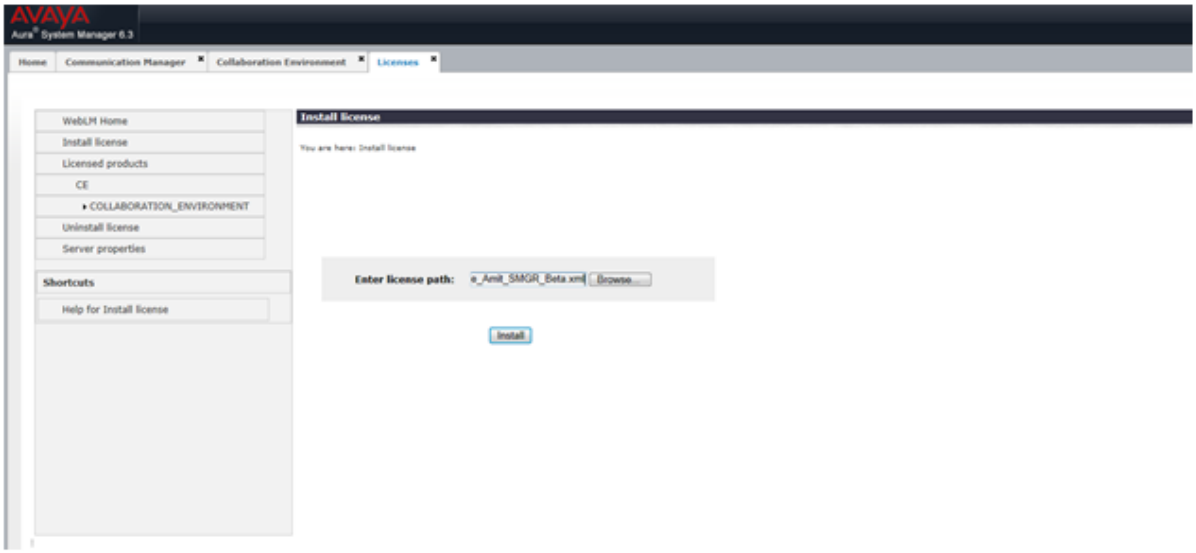
- Is the system under overload? For Collaboration Environment 3.0, system would get into overload state when CPU usage is 80% or beyond that. System would be out of overload when the usage of CPU and memory drops to 60% or below.
- Look at the activity status on Collaboration Environment -> System Administration screen. If there is lot of activity on Collaboration Environment, service access might be slow.

How can I have my service not show "Not Applicable" for my service?

You cannot create a licensed service. Licensed services are only provided by Avaya. Some of these services are – WebRTC, Collaboration Designer etc.

Why does my service show a warning icon with a date?

This means that you have not installed a license file for the licensed service. It is very important that you see a green checkmark for your service. Make sure you install appropriate license for your service.



Once you install the license, it takes a maximum 9 minutes for the License mode to update with a green check mark.

Why do I have call failures when using APIs which use Avaya Media Server?

Make sure the administration is done as per the “Deploying the Avaya Aura® Collaboration Environment” document. Look specifically at the chapter with details about “Avaya Media Server Deployment”

Verify the following configuration on Avaya Media Server Element Manager:

1. System Configuration -> Signaling protocols -> SIP -> General Settings, “Always Use SIP default outbound proxy" is checked.
2. System Configuration -> Signaling protocols -> SIP -> Nodes and Routes : Make sure you have only functions Session Managers mentioned here with default domain i.e. *
3. Avaya Media Server is properly licensed.

4. If using WebRTC:

   Make sure System Configuration -> Server profile -> General Settings have “Firewall NAT Tunneling Media Processor” as checked.

   Also, System Configuration -> Media Processing -> ICE : All settings are correct.

5. After changing any of the above – Restart Avaya Media Server.

Verify the following configuration on System manager:

1. Entity link from Session Manager to Avaya Media Server is set up correctly. If the protocol between Collaboration Environment and Session manager is TLS, make sure protocol between Session Manager and Avaya Media Server is TLS. Or if the protocol between Collaboration Environment and Session manager is TCP, make sure protocol between Session Manager and Avaya Media Server is TCP.
2. Regular Expressions on Routing has pattern of the form ce-msml@<SIP domain>.

My service is invoked and it runs media operations, or indicates that it might run media operations. Why does the call seem to stall?

If your service invokes Call.enableMediaBeforeAnswer() or runs a media operation within the callIntercepted callback in your class that implements CallListener and is annotated with @TheCallListener, then Collaboration Environment will attempt to include the Avaya Media Server in the call. As part of the call setup, Collaboration Environment will instruct the Avaya Media Server to initiate a call to your service. Your service must be configured to accept this call. Check the CARRule.xml file (if you used the service archetype to start writing your service, the CARRule.xml file can be found under yourService-war/src/main/resources) – it should include a TerminatingServiceRule element. Here is an example:

```
        <?xml version="1.0" encoding="UTF-8"?>
<ServiceRules name="${serviceName}" version="${serviceVersion}"
    xmlns:xs="http://www.w3.org/2001/XMLSchema-instance">

   <!-- Following section is needed for call related applications. Refer HelloWorld
        Sample service for more clarifications. If your application is HTTP only,
        please remove the SequencedServiceRule tags. You can refer
DynamicTeamFormation
        Sample service in SDK for help on http only service -->

    <SequencedServiceRule phase="imsterm" order='125'
        desc="Interested in imsterm phase">
    </SequencedServiceRule>

    <SequencedServiceRule phase="imsorig" order='25'
        desc="Interested in imsorig phase">
    </SequencedServiceRule>

    <TerminatingServiceRule desc="for callback to CMATestService">
        <FeatureURI>CMATestService</FeatureURI>
```

```
    </TerminatingServiceRule>

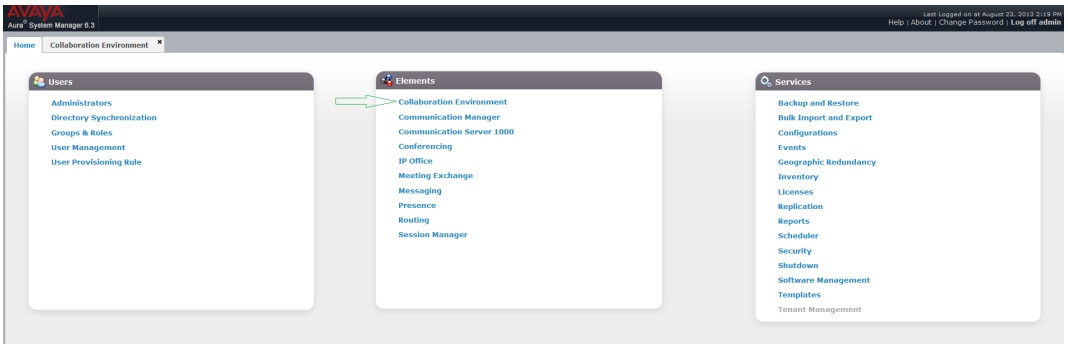
</ServiceRules>
```

The TerminatingServiceRule element contains a child FeatureURI element, the value of which must be the same as the service name.

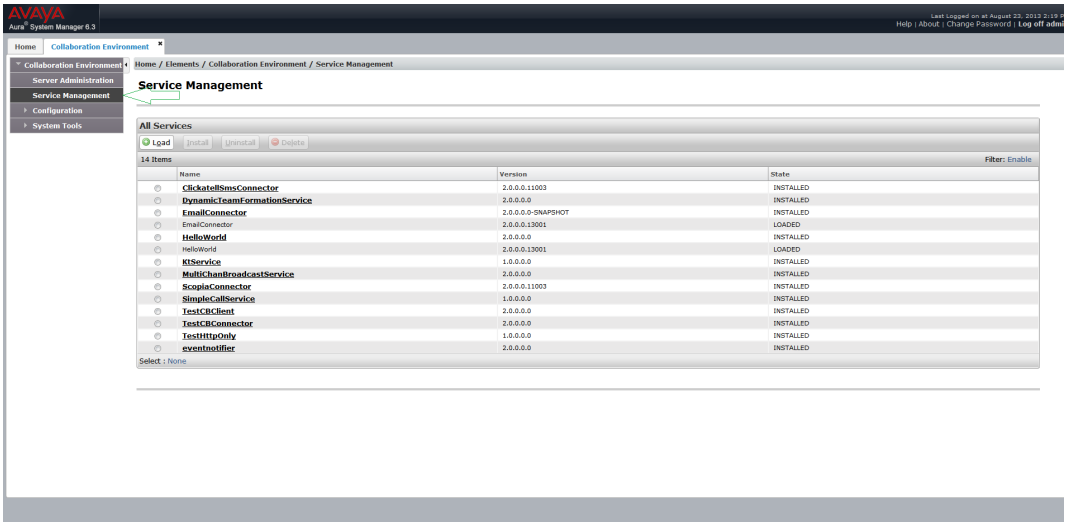
Chapter 5: Connector debugging

Why was email not sent from my service?

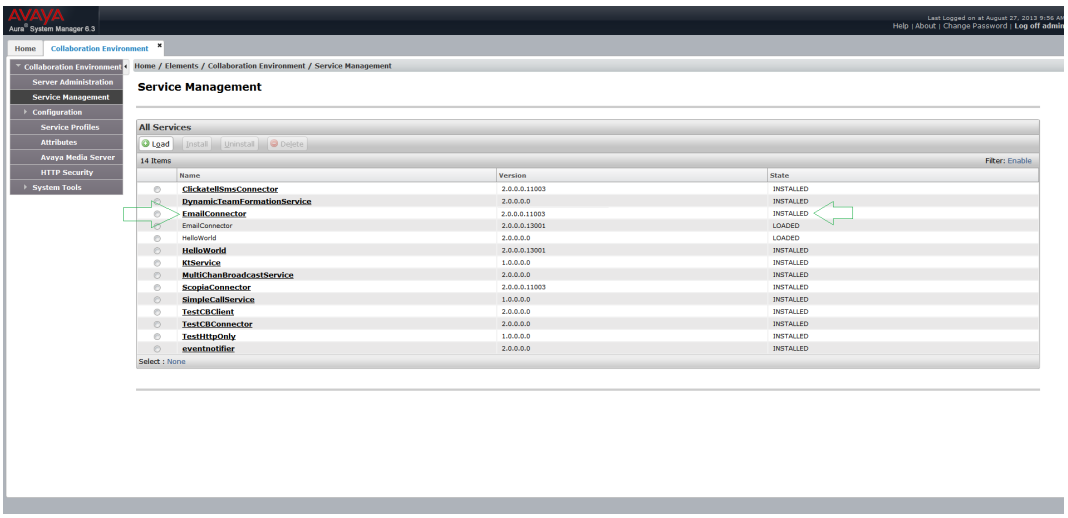
1. Please check that the email connector is installed on the Collaboration Environment Server. By default when you get set up a system, the EmailConnector is only Loaded and not installed. Log in to System Manager. From the home page, select Collaboration Environment from the Elements panel.



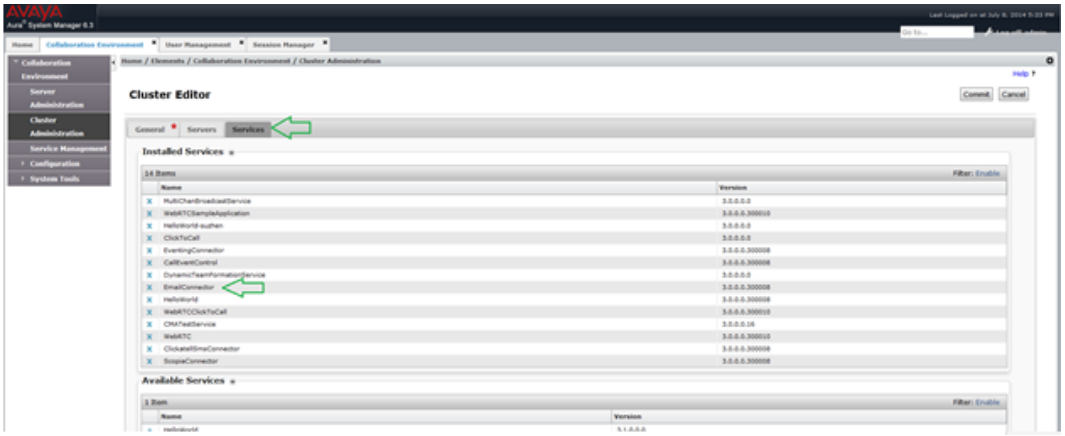
2. Select Service Management from the left navigation



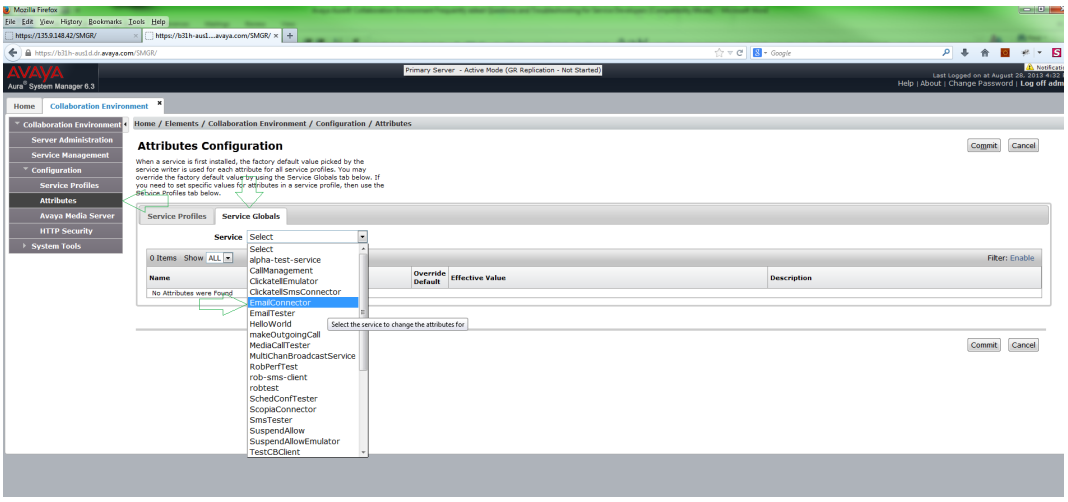
3. Check if the EmailConnector is installed



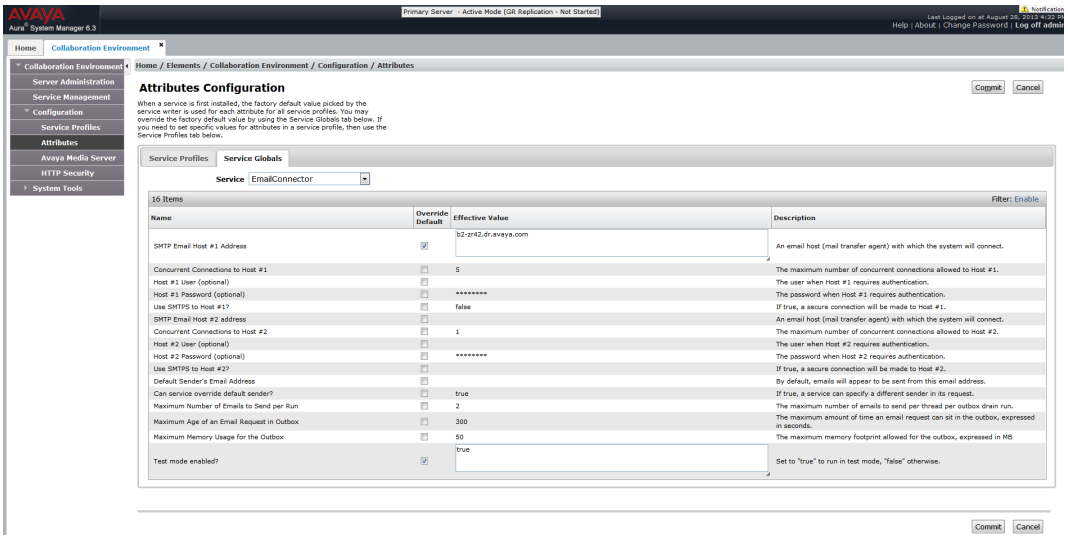
4. Check if Email connector is installed for your cluster. Go to Cluster Administration. Select your cluster and Edit.
5. Go to Services tab to verify that the connector is installed.



6. Go to Configuration -> Attributes in left navigation and select the Service Globals and Service Clusters tab. Then select EmailConnector



7. Make sure the details provided for email server, host are correct.



8. Based on the burst of emails that your server is sending, please configure the setting for “Maximum Number of Emails to Send per Run”, “Maximum Age of an Email Request in Outbox” and “Maximum Memory Usage for the Outbox”. Details about all the parameters are available in *Collaboration Environment Service Development Guide*.
9. Set “Test mode enabled?” to true and check if the request-response is working fine without actually sending an email.
10. Refer to the Multi-channel Broadcast Sample Service which is part of the SDK as it shows how to send an Email using the Collaboration Environment APIs. This can be found in the SDK at – <SDK-3.0 Folder>\samples\multichanbroadcast

    Test using the multi-channel broadcast sample service once. Look at the question “How do I invoke Multi-Channel Broadcast service and Dynamic Team Formation Service for my testing?” if you need details on how to do this.
11. Login to Collaboration Environment server and run this command on command line: **`ce dlogv Email Connector`** This will open the debug log using the vi editor.

Search for your service name. This should provide you with information if there were any exceptions that lead to issues with invoking the service.

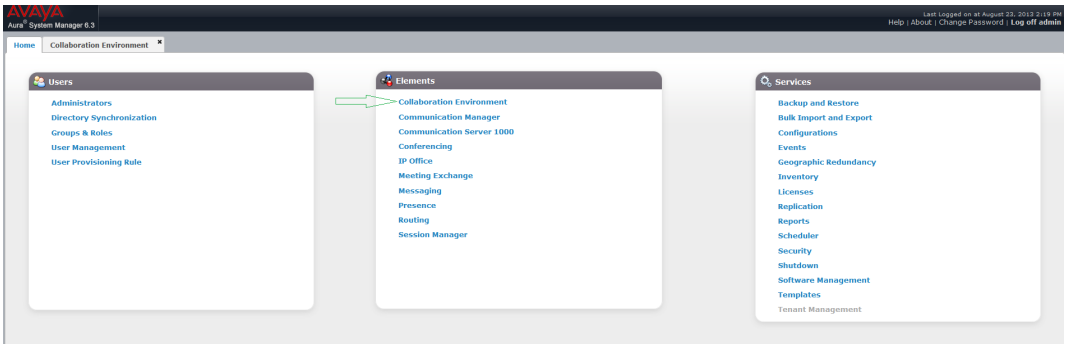
If you added logging statements in your service code, and if you used the Logger class in the Collaboration Environment API (highly recommended), you can view the service log using the command **`ce dlogv <serviceName>`**.For example, to view the logs for the basicSipCall service, run the following command **`ce dlogv basicSipCall`**.

The logs for different versions of a service all go to the same log file, and each log message contains the service name and version.

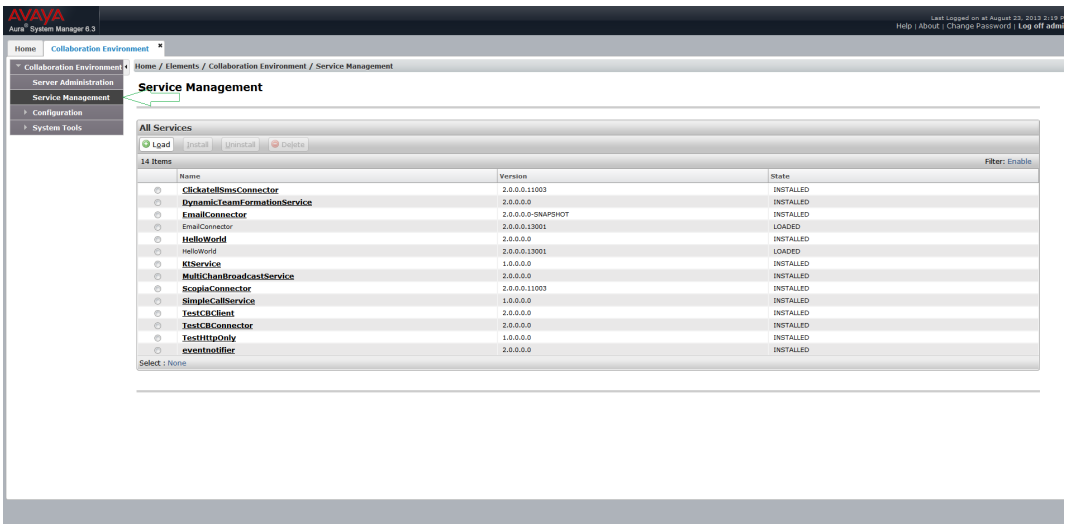
When running in normal mode, logs at level INFO and higher priority (FATAL, ERROR, WARN) will be added to the log file; logs at lower priority (FINE, FINER, FINEST) will not be added to the log. The log level can be changed on the fly to provide more verbose logging. Enter this command while logged in to your Collaboration Environment server to enable all log levels for the Email connector: **`ce dlogon EmailConnector`**When you are finished debugging, you should return logging to the default using this command: **`ce dlogoff EmailConnector`**.

Why was SMS not sent from my service?

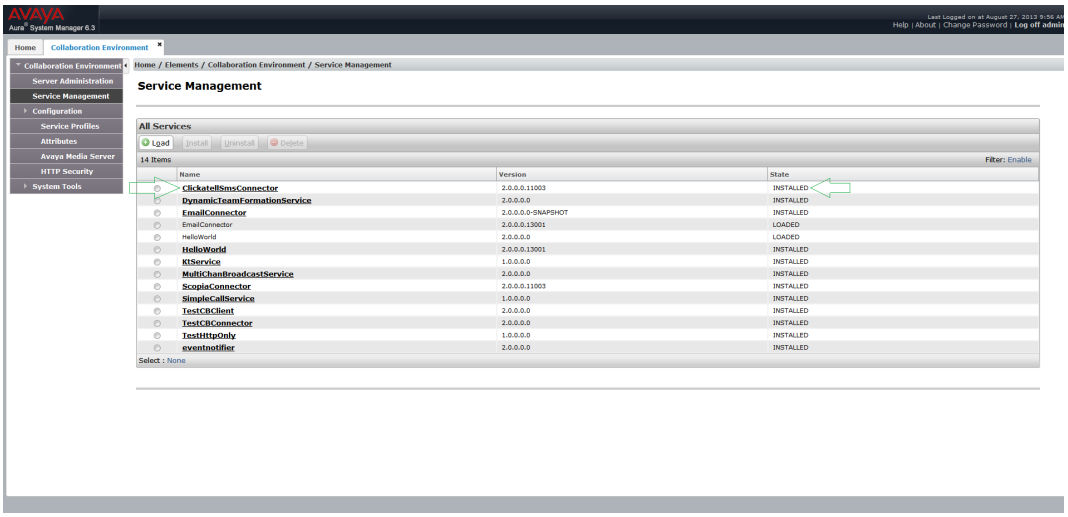
1. Please check that the Clickatell connector is installed on the Collaboration Environment Server. By default when you get set up a system, the ClickatellSmsConnector is only Loaded and not installed. Log in to System Manager. From the home page, select Collaboration Environment from the Elements panel.



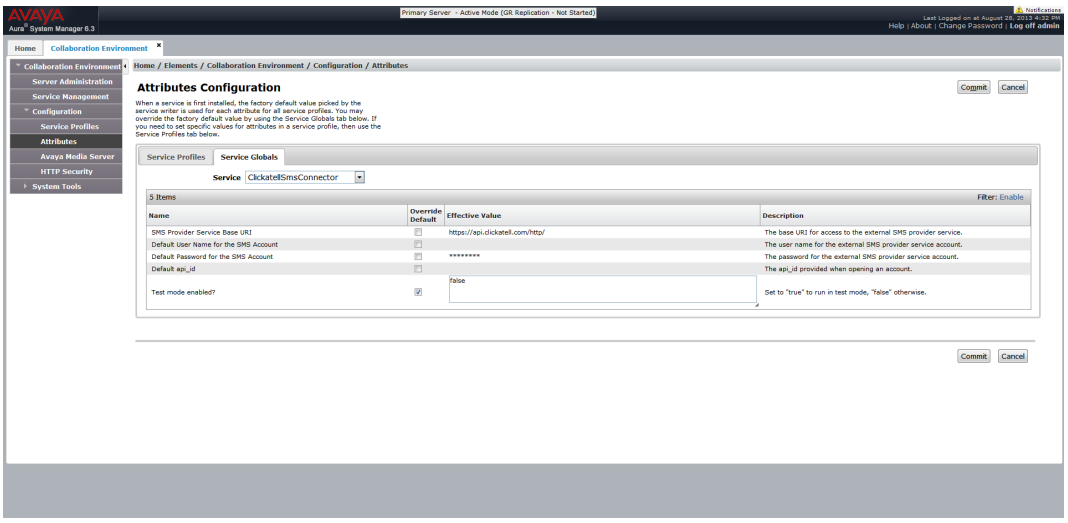
2. Select Service Management from the left navigation



3. Check if the ClickatellSmsConnector is installed



4. Check if ClickatellSmsConnector is installed for your cluster. Go to Cluster Administration, select a cluster and click Edit.
5. Go to Services tab to verify that the connector is installed.
6. Go to Configuration -> Attributes in left navigation and select the Service Globals tab. Then select ClickatellSmsConnector. Make sure the details provided to access Clickatell server is https://api.clickatell.com/http/, user name is a valid account with Clickatell, api_id is provided correctly and also appropriate password is provided. Alternatively, these values can be passed programmatically via the SMS API.



Note:

Details about all the parameters are available in the *Collaboration Environment Service Development Guide*.

7. Set “Test mode enabled?” to true and check if the request-response is working fine without actually sending an SMS.
8. Refer to the Multi-channel Broadcast Sample Service which is part of the SDK as it shows how to send an SMS using the Collaboration Environment APIs. This can be found in the SDK at – <SDK-3.0 Folder>\samples\multichanbroadcast Test using the multi-channel broadcast sample service once. Look at the question “How do I invoke Multi-Channel Broadcast service and Dynamic Team Formation Service for my testing?” if you need details on how to do this.
9. Login to Collaboration Environment server and run this command on command line: **`ce dlogv ClickatellSmsConnector`**

   This will open the SMS connector log using the vi editor. Search for your service name. This should provide you with information if there were any exceptions that lead to issues with sending the SMS.

   If you added logging statements in your service code, and if you used the Logger class in the Collaboration Environment API (highly recommended), you can view the service log using the command **`ce dlogv <serviceName>`**. For example, to view the logs for the basicSipCall service, run the following command **`ce dlogv basicSipCall`**.

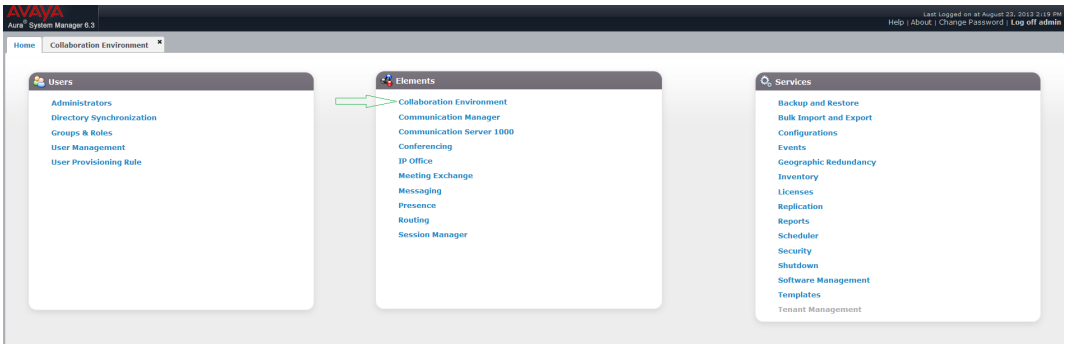
   The logs for different versions of a service all go to the same log file, and each log message contains the service name and version.

   When running in normal mode, logs at level INFO and higher priority (FATAL, ERROR, WARN) will be added to the log file; logs at lower priority (FINE, FINER, FINEST) will not be added to the log. The log level can be changed on the fly to provide more verbose logging. Enter this command while logged in to your Collaboration Environment server to enable all log levels for the SMS connector: **`ce dlogon ClickatellSmsConnector`**

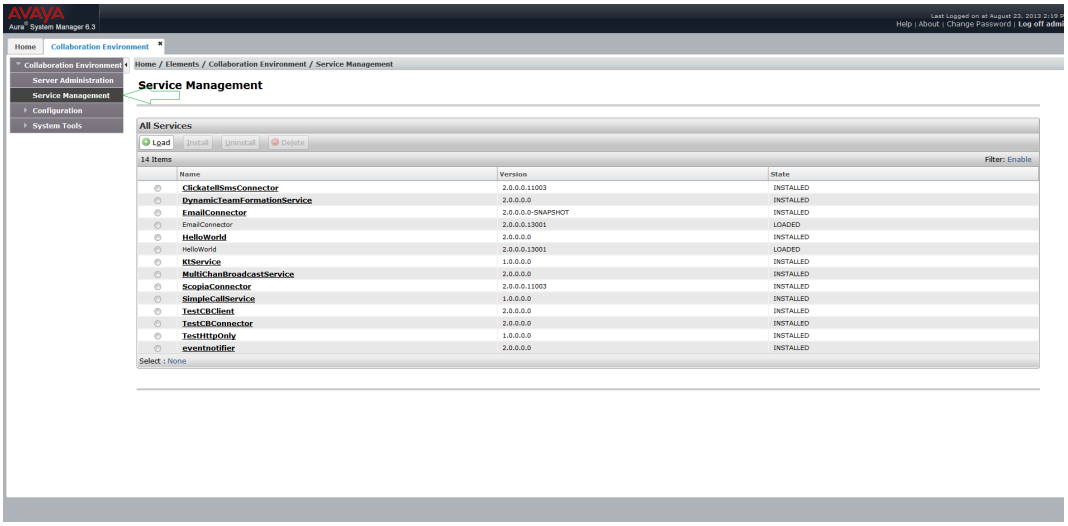
When you are finished debugging, you should return logging to the default using this command: **`ce dlogoff ClickatellSmsConnector`**

Why was Scopia conference not scheduled from my service?

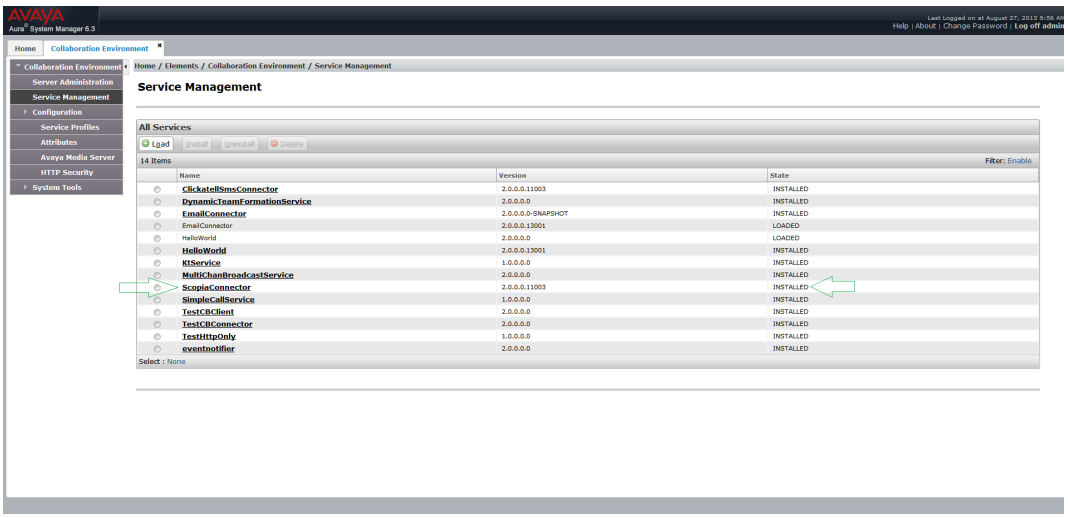
1. Please check that the Scopia connector is installed on the Collaboration Environment Server. By default when you get set up a system, the ScopiaConnector is only Loaded and not installed. Log in to System Manager. From the home page, select Collaboration Environment from the Elements panel.



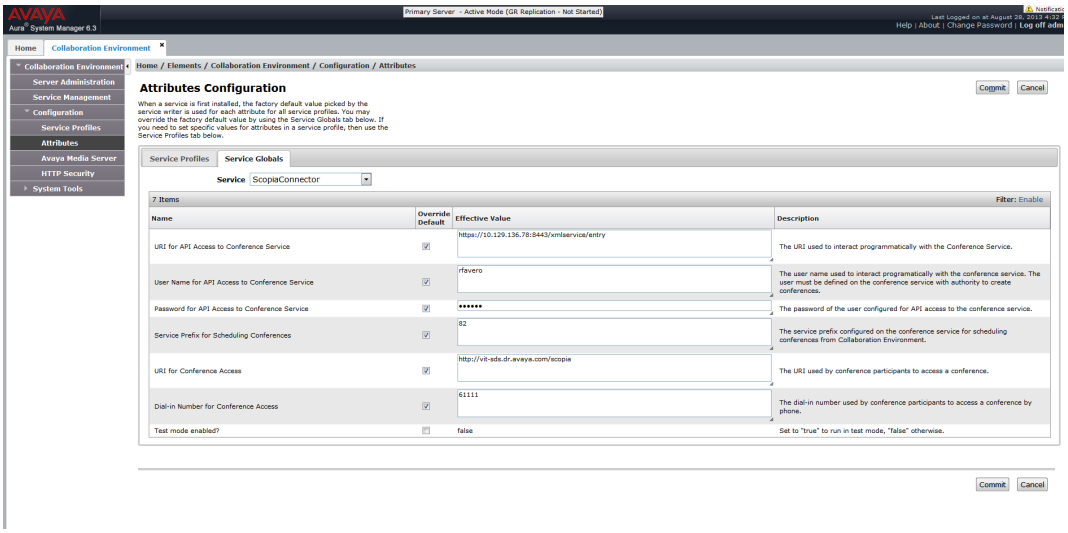
2. Select Service Management from the left navigation



3. Check if the ScopiaConnector is installed



4. Check if ScopiaConnector is installed for your cluster. Go to Cluster Administration, select your cluster and click Edit.
5. Go to Services tab to verify that the connector is installed.
6. Go to Configuration -> Attributes in left navigation and select the Service Globals tab. Then select ScopiaConnector from the drop down for Service. Make sure the details provided for Scopia server address and port is correct. Also, you must have a valid login on the server which you use in User name and password field as shown in screen below.



7. Details about all the parameters are available in *Collaboration Environment Service Development Guide*.
8. Try to launch the URI that you provide in “URI for Conference Access” – You should be able to get to the Scopia Desktop.
9. If you try to call the number specified in “Dial-in Number for Conference Access” – You should hear the attendant announcement.

10. Set “Test mode enabled?” to true and check if the request-response is working fine without actually scheduling a conference on your scopia server.
11. Refer to the Dynamic Team Formation Sample Service which is part of the SDK as it shows how to schedule a Scopia conference using the Collaboration Environment APIs. This can be found in the SDK at – <SDK-3.0 Folder>\samples\ dynamicteamformation

    Test using the dynamic team formation sample service once to verify Scopia server set up is correct. Look at the question “How do I invoke Multi-Channel Broadcast service and Dynamic Team Formation Service for my testing?” if you need details on how to do this.
12. Login to Collaboration Environment server and run this command on command line: **`ce dlogv ScopiaConnector`**

    This will open the debug log using the vi editor. Search for your service name. This should provide you with information if there were any exceptions that lead to issues with scheduling the conference.

    If you added logging statements in your service code, and if you used the Logger class in the Collaboration Environment API (highly recommended), you can view the service log using the command **`ce dlogv <serviceName>`**. For example, to view the logs for the basicSipCall service, run the following command **`ce dlogv basicSipCall`**.

    The logs for different versions of a service all go to the same log file, and each log message contains the service name and version.

    When running in normal mode, logs at level INFO and higher priority (FATAL, ERROR, WARN) will be added to the log file; logs at lower priority (FINE, FINER, FINEST) will not be added to the log. The log level can be changed on the fly to provide more verbose logging. Enter this command while logged in to your Collaboration Environment server to enable all log levels for the SMS connector: **`ce dlogon ScopiaConnector`**

    When you are finished debugging, you should return logging to the default using this command: **`ce dlogoff ScopiaConnector`**

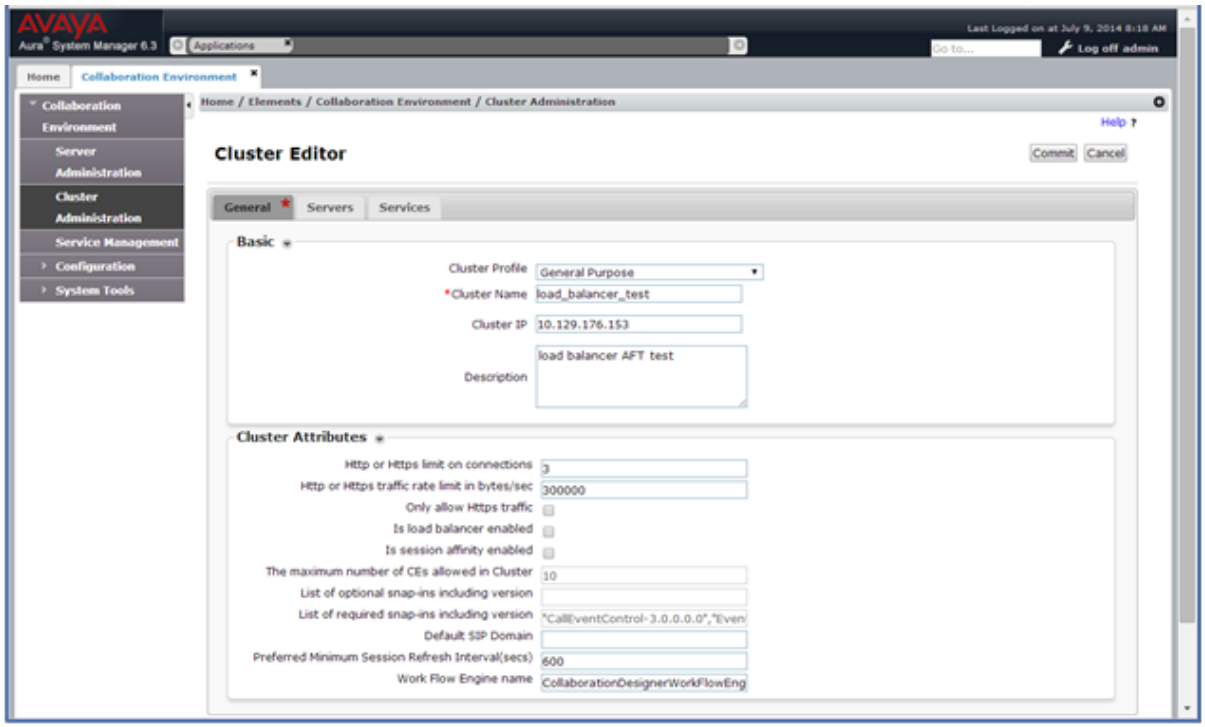
Why do I see the following services – EventingConnector and CallEventControl installed?

These are default services from Collaboration platform and are needed for functioning of Collaboration Environment. These cannot be removed as long as the Collaboration Environment server is part of the cluster and is needed to be functional.

Chapter 6: HTTP Load Balancing and High Availability

How to enable/disable HTTP load balancer in a cluster?

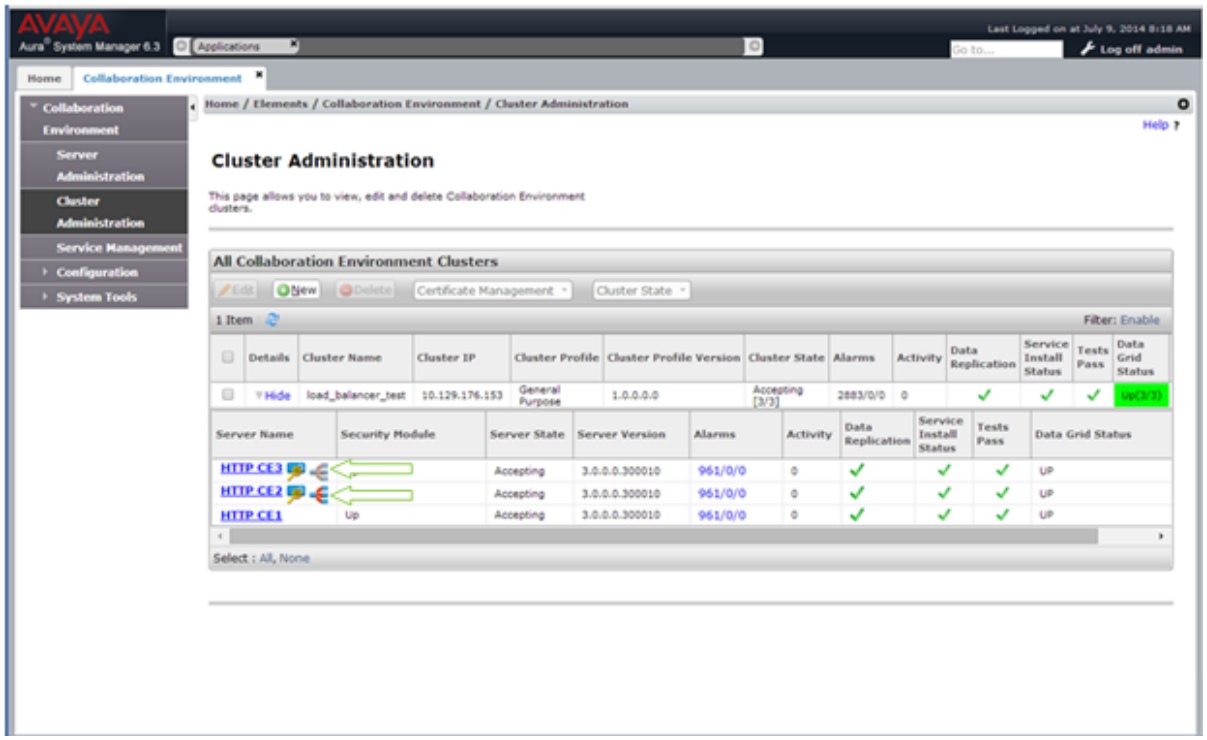
By default, the HTTP load balancer is set to disabled when creating a new cluster. To change the HTTP load balancer state, go to the Cluster Administration page, and click the Is load balancer enabled check box. Click Commit.



How to locate the HTTP load balancer nodes in a cluster?

After you have enabled the HTTP load balancer and added at least two or more servers into the cluster, the System Manager (SMGR) will automatically pick up two Collaboration Environment servers in the cluster to serve the load balancing roles. You can find out the HTTP

load balancer state on Cluster Administration page. The colored icon indicates the active load balancer node and the grayed-out icon indicates the standby load balancer.



You also can verify the load balancer status from the command line on the server. For example:
`$ sohctl -dls curr=ST_HA_ACTIVE new=ST_HA_ACTIVE prev=ST_HA_ACTDISC ACTIVE for=00000:03:08:14:905.`

How is the cluster IP address used on the load balancer?

The cluster IP address will be assigned as an IP alias to eth1:0 interface on active load balancer server. The client should send the HTTP/s request to this IP address and load balancer will distribute the request to one of the upstream servers. The standby load balancer will take over this cluster IP address to become active when failover occurs.

Why hasn't the cluster IP address been assigned to any of the servers?

This condition should not happen unless both active and standby load balancer servers are down or the load balancer was disabled at this time. You can follow these steps for troubleshooting:

1. Make sure load balancer is enabled.
2. Locate the active and standby load balancer servers.
3. Ping management and secure module interfaces on each server.
4. Make sure secmod is running (by running ***statapp*** command).
5. Has cluster administration page been “Commit”ed properly?
6. Make sure the cluster IP address is of the same subnet as the Collaboration Environment servers.

What if planned maintenance needs to be done on active load balancer node?

If you have to manually bring down the server with active load balancer running for any reasons (i.e. software upgrade, hardware replacement,…), you must put this server into Deny New Service state from Server Administration page. That triggers a failover and the standby load balancer will take over the active role within several seconds. By doing this, the server downtime will be minimal with less service affecting.

What if an unplanned failover happened?

In the cases of power outage, network outage, and hardware failure, the standby load balancer will be promoted to active if the failure condition only affected the active load balancer. If both active and standby load balancer servers are affected, the load balancer will be re-established as in previous status once the failure conditions are cleared.

Why do I see both load balancers are in active state?

This condition is called “Split Brain” that is usually caused by the network connection issues. When standby load balancer cannot reach to the active load balancer server thru both management and security module interfaces, it believes itself is the only active server. The standby load balancer will promote itself to become the active load balancer. In the meantime, the previous active load balancer is still up and running, but most likely the server might have lost the network connection.

No action is required in this case. The high availability arbiter will resolve this tie break once the network issue is cleared and both load balancers are able to resume the keepalives between them.

How can I tell when the HTTP Load Balancer failed over?

The easiest way is to check on the load balancing icons on Server Administration or Cluster Administration page. You can also use command line command to list the most recent histories of load balancer on the servers. For example:

```
$ sohctl -lh
!!! Listing only 3 most recent entries for each HealthObject !!!
Resource=replication:
    HealthObject:
    ce_failover                         :
        INFO      2014:07:09-15:06:47:844 O_IVDM52413 "DEMOTED took=001340 ms "
        CLEARED   2012:07:09-10:48:59:000 O_CVDM52606 "reconnected to farend"
        ALARM     2012:07:09-10:48:58:000 O_AVDM52606 "farend NOT reachable SPLIT
BRAIN may have occurred"
Resource=keepalive:
    HealthObject:
    kapath_0                            :
        CLEARED   2012:07:09-10:48:59:000 O_CVDM52502 "path established
10.129.176.147->10.129.176.145"
       ALARM     2012:07:09-10:48:58:000 O_AVDM52502 "path not up ne=10.129.176.147
fe=10.129.176.145"
       ERROR     2012:07:09-10:48:57:000 O_EVDM52502 "path not up ne=10.129.176.147
fe=10.129.176.145"
    kapath_1                            :
        CLEARED   2012:07:09-10:48:59:000 O_CVDM52502 "path established
10.129.176.148->10.129.176.146"
       ALARM     2012:07:09-10:48:58:000 O_AVDM52502 "path not up ne=10.129.176.148
fe=10.129.176.146"
       ERROR     2012:07:09-10:48:57:000 O_EVDM52502 "path not up ne=10.129.176.148
fe=10.129.176.146"
```

How does the load balancer distribute requests in a cluster?

After the load balancer is enabled, all servers in the cluster will be put into the upstream server list, including active and standby load balancers. By default, requests are distributed to all the upstream servers by using a weighted round-robin balancing method.

The load balancer supports failure detection. It will retry the request against the next available upstream server if a server failed to respond to a request with timeout or 5xx error in the response.

How can I tell whether the load balancer distributes the traffic evenly in cluster?

Load balancer uses round-robin to spread traffic evenly between servers in the cluster. For the traffic monitoring, there is no centralized management reporting tool for the traffic and throughput today. You can get details from logs in the following areas if the dive-in investigation is needed:

- Load balancer (/var/log/Avaya/asset/nginx_error.log)
- WebSphere (/var/log/Avaya/sm/TextLog*)
- Snap-in (/var/log/Avaya/services/<snap-in service name>)

How does Session Affinity work on load balancer?

You can enable the Session Affinity on load balancer if session persistence is required. The client's IP address will be used as a hashing key to determine which upstream server should be selected for this request. The method ensures that requests from the same client will always be passed to the same server except when this server is unavailable. In the latter case client requests will be passed to another server in round-robin fashion. Most probably, it will always be the same server as well.

Chapter 7: Eventing Framework

Why do I stop receiving events after 8-9 hours?

Subscriptions have an expiration time associated with them. This expiration time is currently set to 8 hours. Your service will receive an event "EventFamilyEventing.EVENT_SUBSCRIPTION_EXPIRING_QUALIFIED" when its subscription is about to expire. When you receive this event, invoke the "renewSubscription" method on the Eventing Service, passing the payload that was in the received event.

How can I tell which subscriptions were matched when an event is sent?

Turn debug logging on with the command "ce dlogon". If you then watch the logs with "ce dlogw" or view them after events have been sent with "ce dlogv", you will see log statements like the following. These statements tell you which subscriptions were matched.

```
2014-07-14 13:47:04,597
[WorkManager.WebSphere_EJB_Container_AsynchMethods_Internal_WorkManager : 2]
ejb.singleton.PublisherImpl FINEST - sendEvents Sent the following events:
subscriptionId='EventConsumerTe-3.1.0.0.0-5188298e18263886397c1ed620e41d938f3651773
df27ff6f1ceca6a0639cc38'; consumer version to deliver to='3.1.0.0.0'<=>
subscriptionId='EventConsumerTe-
3.0.0.0.0-ec7c68b90748882794a3b6e968a7e06a742b4669f460a555ce7f63af803bbae7';
consumer version to deliver to='3.0.0.0.0'<=>
```

Why am I getting multiple event notifications for a single event?

It's likely that you have multiple subscriptions for the same event family / type. The Eventing Framework has a mechanism to treat subscriptions as being duplicates of each other and ensures that only one event will be sent to a subscriber. This is to handle a case where a service subscribes to events on startup and will therefore subscribe on each server in a cluster. In such cases, it is desirable to treat those multiple subscriptions as a single subscription. However, there are very specific criteria to determine if a subscription is a duplicate and your

subscriptions may not have met those criteria. Please check the Javadoc for the com.avaya.collaboration.eventing package for a detailed description of what qualifies as a duplicate subscription.

One thing in particular to check is to ensure that you do not have server-specific Consumer Private Data in your subscriptions. The Consumer Private Data must match exactly for the subscriptions to be considered duplicates and to therefore only receive a single copy of an event.

Why am I not getting my expected events?

There are several things to consider here. One important point is that a snap-in will not receive any events on a given server until that snap-in has performed at least one subscription. If your snap-in subscribes for events on startup, this is not an issue. If, however, your snap-in subscribes for events based on some stimulus which may arrive at one particular server, you'll want to ensure that all servers are equipped to handle that event (it may be published on a different server). One way to do this would be to always subscribe for the Eventing family of events. These events are sent by default anyway, but this subscription will allow you to receive any other events as a result of a subscription on another server.

Another thing to watch out for is the use of subscriptions based on preferred version. You'll want to ensure that if one version of your snap-in follows this model, they all do.

Why are there log statements saying the Eventing System is still initializing

You may occasionally see log statements such as the following after startup:

```
com.avaya.collaboration.eventing.EventingException
[result=EVENTING_SYSTEM_NOT_READY - The eventing system is still initializing and
cannot yet be used.
```

This is caused by the datagrid component not initializing soon enough after the system starts. The condition should clear itself after a few minutes. If it doesn't, you should log in as root and enter "service was restart" to restart the container.

Why does my subscription not seem to activate?

You may not have defined your event listener properly. Check the log file for your snap-in (ce dlogv <snap-in name>. If everything is OK, you should see the following log statement:

```
2014-07-14 13:27:16,656 [WebContainer : 2] EventConsumerTester INFO -
EventConsumerTester-3.0.0.0.0 - findEventListener Listener class
"com.avaya.zephyr.tester.event.consumer.MyEventListener" found implementing
interface "EventListener" and specified by the properties XML file.
```

The following are error statements you will see if you've done something incorrect:

You haven't specified an event listener in your properties.xml file:

```
2014-07-14 08:45:11,872 [WebContainer : 1] EventConsumerTester ERROR -
EventConsumerTester-3.0.0.0.0 - findEventListener No Event Listener found. See
documentation regarding the annotation "TheEventListener".
```

The class name specified as the event listener in properties.xml doesn't exist:

```
2014-07-14 08:45:11,872 [WebContainer : 1] EventConsumerTester ERROR -
EventConsumerTester-3.0.0.0.0 - findEventListener No Event Listener found. See
documentation regarding the annotation "TheEventListener".
```

You haven't annotated your event listener class with "@Stateless"

```
2014-07-14 13:30:03,496 [WebContainer : 3] EventConsumerTester ERROR -
EventConsumerTester-3.0.0.0.0 - confirmClassDefinition There is an error in the
definition of the class implementing interface "EventListener" and specified by the
properties XML file. Class must be annotated with @Stateless.
```

Your event listener class doesn't implement the EventListener interface

```
2014-07-14 13:33:24,826 [WebContainer : 2] EventConsumerTester ERROR -
EventConsumerTester-3.0.0.0.0 - confirmClassDefinition There is an error in the
definition of the class implementing interface "EventListener" and specified by the
properties XML file. Class does not implement interface.
```