IBM Tealeaf CX
Version 9 Release 0.1
December 4, 2014

Tealeaf Troubleshooting Guide

Note
Before using this information and the product it supports, read the information in "Notices" on page 129.

This edition applies to version 9, release 0, modification 1 of IBM Tealeaf CX and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Tealeaf Troubleshooting Guide

The Tealeaf® Troubleshooting Guide provides techniques for managing common issues with the Tealeaf system. Use the links below to access specific topics in the guide.

Knowing if or when to troubleshoot

Time of occurrence of error messages

Do error messages occur at the same time of day as the daily scheduled restart of Tealeaf services? If so, most likely the error messages are ignorable, as it is expected that access to components of the Tealeaf system would be unavailable while the services are stopping and starting.

Search Server service

These issues are generally ignorable if they are not causing Tealeaf users any problems and the service is being restarted automatically if/when it stops.

TeaLeaf Search Server - authorization error: NetGroupGetUsers() or NetLocalGroupGetMembers()

Applies to versions:
4.x; 5.x; 6.x
Solution Details
If you are receiving the following error: TeaLeaf Search Server Authorization error: NetGroupGetUsers() error (group:): The RPC server is unavailable.
(or the same message but citing the NetLocalGroupGetMembers() function call):

Intermittent messages relating to the NetGroupGetUsers or NetLocalGroupGetMembers function call indicate that the Tealeaf servers cannot contact their Domain Controller(s) to get the Active Directory group information then. Discuss with the NT/AD administrators to investigate why this might be occurring.

If this behavior is intermittent and of short duration, this is more a nuisance than a serious error, as the Search Server keeps a cached copy of the user group membership information it got from the most recent refresh. It should not affect use of the Tealeaf system unless you added a new user to the group at the domain controller, this error occurred, and that new user then attempted to search (or use any part of the system that requires authentication, which would be both the Viewer's search as well as logging in to and using the Portal). But after the next successful refresh the Search Server does of the user group definitions, all should work as expected.

TeaLeaf Search Server Ver: 6.2.0.6241 - main thread proc crashed on cmd/COMMAND. See TLSrchSrv log.

Tealeaf CX problem diagnosis

To properly diagnose an issue with your Tealeaf system, you must first isolate the Tealeaf component that is causing the error. You can use the following steps to run the provided diagnostic tools and check standard issues such as disk space and log files:

To begin diagnosis:

1. **Run Tealeaf Status**. The Tealeaf Status tool runs a series of tests on each Tealeaf server to help identifying problems. Tealeaf Status is a reporting mechanism that is independent of the IBM Tealeaf CX system. It provides a dashboard view into the health of the system.
   - In most installations, Tealeaf Status is configured to run as a scheduled task. You can also trigger it manually. See "Scheduling Tealeaf Jobs" in the *IBM Tealeaf cxImpact Administration Manual*.
2. **Check disk space**. Disk usage by Tealeaf components is also a threshold of the Tealeaf Status report.
3. **View application and event log messages**. Tealeaf generates extensive logs on all components, events, and services. You should review these logs on a periodic basis, as well as to triage urgent issues.
   - Tealeaf stores its log files in the following location: `<Tealeaf_Install_Directory>\Logs`
4. **Run Search Server**. Search Server enables you to check and fix problems in indexes. A check and fix operation enables the Canister to write `TLC` files to the `FilesToIndex` directory, where they are reindexed by the Index Multiprocessor. To perform an index check and fix in Search Server:
   a. Open a browser and enter the following URL: `http://<machinename>:19000`
   b. At the bottom of the Search Server page, click **Run Canister Index Checker**.
   c. The total number of sessions versus the number of indexed sessions is displayed. These indexes are processed in batches of 100 to make index processing more efficient.
      - If the Indexed number is lower, the non-indexed sessions are listed at the bottom of the page.

Troubleshooting strategy: where to start, where to go

Knowing the function of the system's components and the data flow paths among them can help to troubleshoot at or near the likely source of the problem. At least, this knowledge can provide guidance in a likely direction toward the problem's source by eliminating other possible sources in an organized way, by following the data flow paths, either in the natural or the reverse direction.

Data flow paths among the Tealeaf system components

The figure below shows the directions in which data flow through the components of a Tealeaf system. The "pull" label on an arrow indicates that the destination component actively pulls data into it from its source component; otherwise, the destination passively receives data that is pushed to it by the source component.

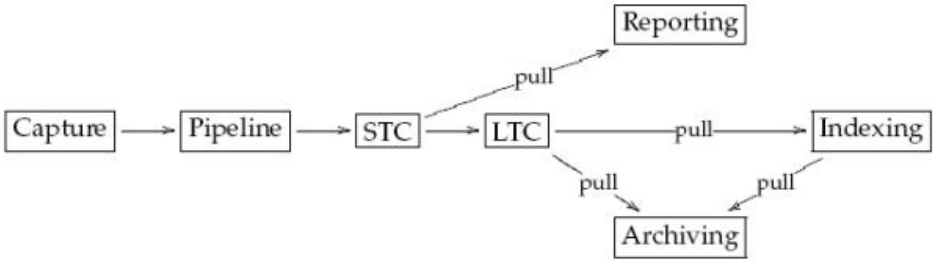

- **Capture:** Tealeaf data capture occurs either via a IBM Tealeaf CX Passive Capture Application server (PCA server) or IIS capture filter.
- **Pipeline:** The pipeline runs on the PCA server or in the Tealeaf Transport Service (or possibly IIS capture, though IIS capture pipelines are minimal).
- **STC:** The Short-Term Canister is an in-memory datastore that performs sessionization of hits, assembly of statistics, and event detection. It also closes sessions when they exceed the configured session idle timeout period or trigger a session-close event.
- **LTC:** The Long-Term Canister provides disk-based storage of sessions that can be replayed by Tealeaf users.
- **Indexing:** To support search of sessions that are stored in the LTC, sessions are indexed by the search engine.
- **Reporting:** Statistics on events and fundamental values, such as hits, pages, sessions, and bytes are collected from the STC and aggregated versus time in a SQL database.
- **Archiving:** The optional Tealeaf archive services provide longer-term storage and faster searching than the LTC, if they are installed on a separate machine from the storage canisters. Archives are also easier to back up using third-party tools than a running Canister.
- **Search:** Searching in the Tealeaf system is performed by the Search Server service, which is not depicted in the diagram. In the figure below, the "search" and "replay" data flows are mediated by these services. Search Server provides search capabilities for the STC and LTC.

The figure below illustrates the data flows between the Tealeaf clients (Portal and IBM Tealeaf CX RealiTea Viewer) and the Canister server-side components. These flows are pulls initiated by the clients; "S" stands for search and "R" stands for session replay.

- Session replay by the Portal might either be performed by Web browser-based replay in IBM Tealeaf cxImpact, IBM Tealeaf cxView 6.0 and later, and IBM Tealeaf cxResults 7.0 and later, or by the IBM Tealeaf CX RealiTea Viewer in any Tealeaf version.

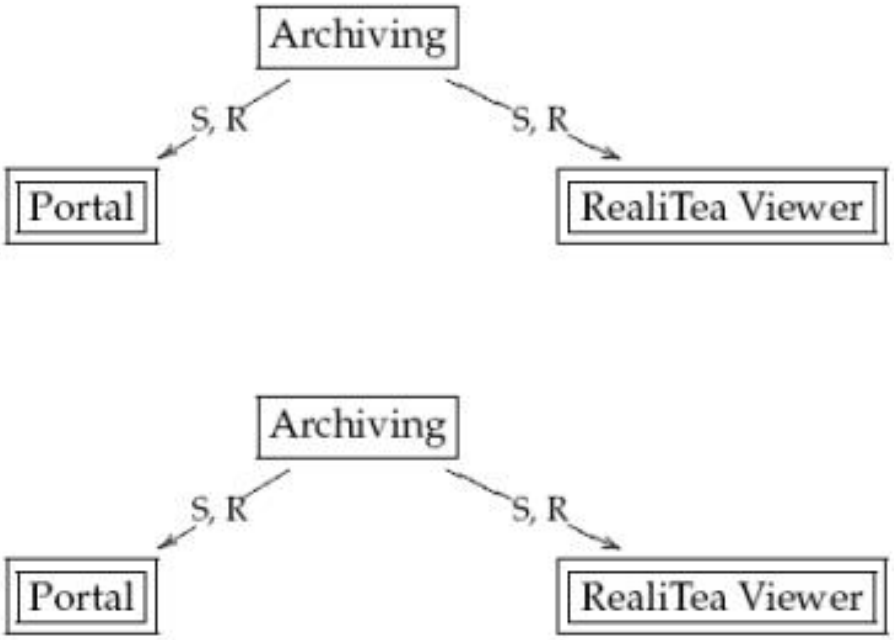

Tools

Diagnostic utilities

- PCA Web UI
  - Summary page
  - SSL Keys page
  - Backups/Logs page
  - Statistics page
- Tealeaf Socket Status
- Tealeaf Pipeline Status

Testing tools

- TealeafArchiveReader
- Event Tester

Configuration points

Configuration files

- `TealeafCaptureSocket.cfg`
- `RTA.ini`
- `Privacy.cfg`
- `RTASplit.ini`
- `HitRouter_PipelineN.cfg`
- `TealeafIIS[6].cfg`

- others referenced by TealeafCaptureSocket.cfg or its contents

GUI configuration utilities

Most configuration of Tealeaf components is now managed through the Tealeaf Management System. See "Tealeaf Management System" in the *IBM Tealeaf cxImpact Administration Manual*.

- For information about troubleshooting TMS, see Chapter 6, “Troubleshooting the Tealeaf Management System,” on page 123.

Events can be created and edited through the Tealeaf Event Manager.

- See "Tealeaf Event Manager" in the *IBM Tealeaf Event Manager Manual*.

Portal configuration

- **Admin tools** (RealiTea 4.6 and earlier)
- **Tools** > **Administration** (IBM Tealeaf cxImpact 5.x and 6.x, IBM Tealeaf cxView 6.x and later)
- **Search** > **Configure Templates** (IBM Tealeaf cxImpact and IBM Tealeaf cxView 6.x)
- **Tealeaf** > **Portal Management** (IBM Tealeaf cxImpact, IBM Tealeaf cxView, IBM Tealeaf cxResults 7.0 and later)
- **Configure** (IBM Tealeaf cxImpact, IBM Tealeaf cxView, IBM Tealeaf cxResults 7.0 and later)

Configuration settings and values

Usernames and passwords

TLADMIN and TLUSER accounts are used to access the SQL Server database(s). The passwords for these two accounts can be changed in SQL Server itself, and the same values would then must be entered in the Report Server configuration in TMS.

- See "Configuring the Report Server" in the *IBM Tealeaf CX Configuration Manual*.

Human users must have Portal accounts that are created within the Portal Web application's UI.

- Even when using the Windows NT authentication integration feature, the Portal must be used to create user accounts, although the Portal can be configured to auto-create accounts for authorized users that are in the NT/AD groups that are specified for authentication through search configuration in TMS.
  - See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

TLADMIN and TLUSER in the Canister datastore, which are shown in Canister configuration and Index configuration in TM, are also only used by the Tealeaf software, not by human users directly. The default passwords are the same as the user names themselves; they never must be changed. However, if needed, they can be changed through one of these two utilities.

- See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.
- See "Configuring CX Indexing" in the *IBM Tealeaf CX Configuration Manual*.

The SQL Server and Canister datastore user names and passwords are unrelated to one another. They happen to have the same names.

Chapter 1. Preserving and rebuilding data

The following sections contain procedures for backing up critical Tealeaf configurations and data so that you can quickly restore or migrate your system in the event of a system failure.

Data backup

Tealeaf recommends that you perform regular backups of all Tealeaf databases using one of the recommended strategies. Depending on the size and daily volumes of captured traffic, one strategy may be more appropriate.

Before you clear data from the Canisters or the SQL Server databases, it is recommended that you perform backups of them.

Note: All tables containing indexes are stored in the Reporting database, which enables accurate "point in time" snapshots to made by performing a backup of this single database.

- See "TLBackup and TLRestore" in the *IBM Tealeaf cxImpact Administration Manual*.
- See "Database Backup Strategy" in the *IBM Tealeaf Databases Guide*.

Preserving environment or Migrate CX settings

In preparation for a complete crash in which the entire machine is lost, you should create copies of the following IBM Tealeaf CX Server items so you can re-create the installation without recustomizing the installation:

- Backup Long Term Canister and indexes
- Tealeaf report definitions (`TL_SYSTEM` database)
- System folder contents
- All configuration files (`*.cfg`)
- Rules Scripts directory (especially if split pipeline)
- Registry hives:
  1. `HKEY_LOCAL_MACHINE\SOFTWARE\TeaLeaf Technology`
  2. `HKEY_CURRENT_USER\Software\TeaLeaf Technology`
  3. `HKEY_CURRENT_USER\Software\TeaLeaf`

Mirroring servers

You can preserve your environment by creating a mirror server that contains the same installation and settings as the one on which you are having problems. To do so, use Norton Ghost or a similar tool to image the server's disks.

Clearing data from the databases

Gradual removal

To perform a gradual clearing of the databases, you can adjust your canister, data collection, and data retention settings so that data is removed without losing any data that was not processed by the Tealeaf system. For example, if you configure

data retention to be a single day and verify that data collection occurs shortly before the daily time for trimming the database, you can remove data more gracefully.

Note: Unless trimmed data was previously archived, it is no longer available for search and retrieval.

- See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

For more information about the data collection and aggregation processes, see "Data Aggregation and Retention" in the *IBM Tealeaf cxImpact Administration Manual*.

- See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.
- See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

Gradual data clearing can be manipulated by changing system parameters.

- **IBM Tealeaf CX databases:** See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.
- **IBM Tealeaf cxResults databases:** See "cxResults Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

The following sections provide more information about how to force clearing of data from the Canister database and the SQL Server databases.

Clearing data from the canister

The canister can be cleared immediately using the CanRebuild utility to rebuild it. While CanRebuild can be used to immediately remove problematic data, by default Tealeaf canisters and databases are trimmed of old data based on configuration settings.

Note: Data that are cleared from the Canister cannot be recovered. Tealeaf recommends that you perform a backup of the Canister for safety. See "TLBackup and TLRestore" in the *IBM Tealeaf cxImpact Administration Manual*.

Rebuilding the Canister

If needed, you can rebuild the Short Term Canister, Long Term Archive, and corresponding dtSearch indexes. The CanRebuild utility also deletes all session data. After deleting all data, it then rebuilds the Short Term Canister tables and reads event definitions.

Note: CanRebuild erases ALL session data from the IBM Tealeaf CX datastore. Use this utility only as a last resort, when the IBM Tealeaf CX Server is not responding and no other attempt to repair the datastore was successful.

Before running the CanRebuild utility:

- Back up session data using the `TLBackup.exe` utility.
- Verify that the capture source stopped sending hits to the Canister. You must either stop the capture filter or queue up hits in the Transport Service pipeline with the Extended Decoupler.

  Note: The CanRebuild utility works only with a local file system. Do not try to remove the files on a remote database, as the files do not exist. For more information about rebuilding a remote SQL database, see “Rebuilding the remote MS SQL database” on page 3.

To rebuild theIBM TealeafCX datastore:

1. Open a command prompt. Navigate to the Tealeaf installation directory.
2. Enter the following command at the command-line prompt:
   `CanRebuild`
   The following options are available, depending on the version of Tealeaf that is installed:

   Note: All existing Canister data is destroyed if you do not select `Preserve Session Data`.

   Option Description

   `Rebuild Full Canister`
   Stops services on the IBM Tealeaf CX Server machine and indexing machine if it is a dual installation. Deletes all session data from the Canister and rebuilds the Canister tables. The Rebuild Full Canister option also deletes the session's corresponding indexes.

   `Delete Spool Files`
   If this option is selected, all files containing spooled hits that are stored on the Canister hard disk are deleted. Use this option only if the spooled data is known to be in a corrupted state or otherwise must be removed.

   `Preserve Session Data`
   Select this option to preserve the session data that is stored in the Long-Term Canister on the server.

   Note: If this option is not enabled, all session data on the server, including index files, is deleted as part of the rebuild.
3. To perform the Canister rebuild, click **Rebuild**.
   - To exit the utility without rebuilding, click **Cancel**.

Clearing data from SQL Server

The SQL Server databases can be cleared by dropping and re-creating them with the following exception:

Note: Unless the Tealeaf system is being reinstalled, do not ever drop and re-create the `TL_SYSTEM` database.

The following IBM Tealeaf CX databases can be dropped and re-created independently:

- `TL_REPORTS`
- `TL_STATISTICS`
- `TL_RSEXTRACTOR`

The IBM Tealeaf cxResults databases are linked and cannot be dropped independently.

Note: If you must drop and re-create one of the IBM Tealeaf cxResults databases, you must drop and re-create both (`TL_VISSTAGE` and `TL_VISREPORT`.

Rebuilding the remote MS SQL database

Dropping and re-creating databases can be managed through the Tealeaf Database Manager.

To rebuild a remote MS SQL database:

Note: This process requires System Administrator privileges on the SQL Server.

1. Back up the existing database to be rebuilt, in case the data must be restored.
2. On the Portal Server, start the Tealeaf Database Manager. Double-click the following executable:

   `\<Tealeaf_install_directory>\Reporting\TealeafDatabaseManager.exe`

   - See "Installing Tealeaf Databases" in the *IBM Tealeaf Databases Guide*.
   - For more information about the Tealeaf Database Manager, see "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.
   - See "Tealeaf Database Manager Reference" in the *IBM Tealeaf Databases Guide*.

Checking the time on Tealeaf systems

Note: It is critical that the time on all Tealeaf servers be closely synchronized. All servers should be synchronized to a common source that is itself synchronized to a reliable master clock.

Tealeaf servers

Tealeaf software typically runs on Windows 2003 or Windows 2008 servers. In the Control Panel of each server, you can use the Date and Time applet to verify that the time zone, date, and time are set properly.

Servers automatically keep their clocks in sync with the clock on the domain controller to which the server connects. If the Tealeaf system does not belong to a domain, open a support case with http://support.tealeaf.com/ asking for assistance in configuring scripted or 3rd-party time synchronization solutions.

Tealeaf timezone

As of Release 8.0, a universal timezone setting is applied to all Tealeaf servers.

All Tealeaf software is aware of the necessary changes for daylight saving time. All time zone selection is done using the Standard time offset. For example, `GMT-5:00 Eastern Time (US and Canada)` should be selected for all systems on the East Coast of the United States, even if daylight saving time is in effect.

- See "Configuring the System Timezone" in the *IBM Tealeaf CX Configuration Manual*.

Passive Capture Application serves (PCA servers)

Since the PCA server may not have access to the enterprise master clock, it should be configured to synchronize its time with the main Tealeaf Transport Service to which it connects.

The PCA server time is the most critical of the time settings. All PCA servers should be within a few seconds of each other. The PCA server creates the `[timestamp]` section of each request. The dates and times that recorded here are in GMT time.

- Note the `Z` appended to the end of each timestamp.

To synchronize the PCA server time:

1. Log in to the PCA server using SSH and the root user ID.

a. You can use the PuTTY program, often installed on the Tealeaf server hosting the main Tealeaf Transport Service in a data center.

2. Run the command `date`, which displays the date, time, and time zone that is configured on the PCA server.
   - See "PCA Web Console - Delivery Tab" in the *IBM Tealeaf Passive Capture Application Manual*.
   - For more information about changing the time zone, see "Configuring the System Timezone" in the *IBM Tealeaf CX Configuration Manual*.

RealiTea Viewer timezone

Individual users of the IBM Tealeaf CX RealiTea Viewer stand-alone application can configure the timezone that is used by RTV, which is typically the time zone of the RTV user.

1. Start RTV.
2. Select **Tools** > **Options** from the menu bar.
3. Select the **Replay** tab in the resulting dialog.
4. Select the correct time zone from the TimeZone drop-down.

Configuration

For more information about configuring the system time, see "Configuring the System Timezone" in the *IBM Tealeaf CX Configuration Manual*.

Before you begin

Before you begin troubleshooting, you should review and verify that you completed the following configuration steps successfully:

- After the PCA is installed, more configuration may be required to effectively capture all required data types and to tune the PCA for capture in your environment. See "Installation" in the *IBM Tealeaf Passive Capture Application Manual*.
- You should also review and verify that your PCA was properly configured when initially installed. See "Initial PCA Configuration" in the *IBM Tealeaf Passive Capture Application Manual*.

Troubleshooting a passive capture application server (PCA server)

You can determine whether a Tealeaf Processing Server is receiving any captured data from a PCA by reviewing the Pipeline Status tab in TMS.

1. Log in to the Tealeaf Portal as an administration.
2. From the menu, select **Tealeaf** > **TMS**.
3. Click the **Pipeline Status** tab. See "TMS Pipeline Status Tab" in the *IBM Tealeaf cxImpact Administration Manual*.
4. Look for a connection that is labeled with the name or IP address as the PCA, showing nonzero page views and bytes of received data. If you can find a connection, then the Processing Server is capturing data from the PCA.

   Otherwise, proceed with the rest of this solution:
5. Check the current statistics on the Summary (home) page of the PCA's Web UI.
   a. When it auto-refreshes to show approximately the past 15 seconds of activity, verify if there is a nonzero number of Hits and Packets.
   b. If there are zero Hits, continue to the next part of this solution.

6. Run the following TCPDump command from the PCA's command line.
   a. Run the command for a sufficient interval to determine that your Web server IP addresses appear both on the left and right sides of the > character in the output lines. If not, tell the network team the side of the > sign on which the Web server IP addresses do not appear.
   b. Assuming your Web servers are operating on the standard IP port number 80, the command is:

      ```
      tcpdump -n -i <NIC> port 80
      ```
7. If there is bidirectional port 80 traffic from your Web servers being seen on one or more of the PCA's network interfaces, then further diagnosis of the traffic is necessary.
   a. Use TCPDump to write some of the network traffic to a file with the following command, which writes binary format data to `output_file_name`:

      ```
      tcpdump -n -i <NIC> -s0 -w output_file_name port 80
      ```
8. `output_file_name`

   should be opened and analyzed for anomalies such as missing packets using the open source Wireshark protocol analyzer (http://wireshark.org/) or an equivalent tool.
   a. Tealeaf can help with this diagnosis if it is not clear what you are seeing, either with a Live Meeting session or by getting the dump file from you, if it does not contain any user's personal information.

/var/log/messages grows very quickly on PCA

The default `syslog.conf` causes /var/log/messages to receive all log notice messages, including the ones from passive capture. To change `syslog.conf` so that the "messages" file no longer receives passive capture messages on Red Hat Linux, edit the /var/log/messages line in the `/etc/syslog.conf` file:

```
*.info;mail.none;authpriv.none;cron.none;local0.none /var/log/messages
```

Include the `local0.none` setting to prevent the passive capture log messages from being written to /var/log/messages.

- These messages continue to be written to the Tealeaf specific `capture.log` file. Ultimately, the presence of many messages from passive capture in either of these log files suggests a problem with the input data coming into the IBM Tealeaf CX Passive Capture Application servers (PCA servers) capture NIC's. The above procedure is only a means of eliminating redundant logging of the passive capture messages; corrective action on the input data stream to the PCA is likely still required.

Large number of `ReqCancelled=Client` hits

Sometimes, a TeaLeaf system records a large number of `ReqCancelled=Client` requests. There may be 5% or more hits that are scattered randomly through the session with a `ReqCancelled=Client`, or there may be a concentrated number of `ReqCancelled=Client` hits occurring in immediate proximity to each other.

In either case, it raises the question as to the effectiveness of the Tealeaf capture. Troubleshooting this issue involves the use of the tcpdump command.
Running tcpdump to a file can result in some large files quickly. To make best use of tcpdump and the http://support.tealeaf.com/ resources to analyze the data, it is

important to be able to reproduce the sessions that have a large number of `ReqCancelled=Client` hits. Before starting to record dump files for analysis, you should be able to effectively produce the problem behavior. So, a first step is some investigation into how to reliably reproduce `ReqCancelled` hits.

Note: This solution applies to Tealeaf version 4.5 SP0 or later, using Standard Indexing.

- If you have the Tealeaf Data Extractor or IBM Tealeaf cxConnect for Data Analysis in place, ask the BI analysis team to run some queries against the Tealeaf data to determine:
  - The top 10 URL's on which `ReqCancelled` is occurring as a percentage of total hits
  - Top ten hit numbers within a session, particularly if it is occurring more often in the beginning of a session
  - Top 3 hours of the day when `ReqCancelled` is occurring.
  - Whether `ReqCancelled` occurs more often during a GET or POST operation.
  - Is there a specific server or data farm on which it occurs most often?
  - Is there a specific proxy or load balancer that causes these (analysis of the HTTP VIA® REQ field)?
- If there is no Data Extractor, the analysis is more involved:
  1. Add the following fields to the `RTA.ini` rule that specifies more fields to index, if not already present in the list.
     - `ReqCancelled`
     - `HTTP VIA`, if it is provided in the REQs captured by Tealeaf. Some networks do not have this field, but networks involving physical load balancers or proxies will.
  2. Run data collection as usual for at least one business day.
  3. Manually use searches to get the same information specified in case 1 above. Once the analysis led to specific conditions most likely to cause sessions with high numbers of `ReqCancelled=Client`, it is time to record a data dump.

Recording a data dump

1. As closely as possible, you should replicate conditions that are known to most likely cause `ReqCancelled=Client` hits. You should verify that you can access any specific pages or send the requests to any specific Web servers or proxies. Complete a few test runs, drive specific test sessions, and use the Tealeaf captured data to verify that your test sessions still have a high number of `ReqCancelled=Client` pages.
2. Prepare a time for making the real tcpdump recording.
   a. If possible, schedule it for a quieter time of the day.
   b. If the tcpdump data contains only a few sessions, that will make it much easier to spot the problem.
   c. You want to run tcpdump for a short period, not more than 5 minutes if possible. By this time, you created a repeatable test case that is guaranteed to cause the behavior quickly, and you will not need to run tcpdump for a long period.
3. Log in to the PCA using SSH and set up the tcpdump command, specifying as many restrictive conditions as you can. For example, if you know the test case goes against a specific Web server, specify to tcpdump to listen only to that IP address.
4. Start the tcpdump command.

5. Run the test.
6. Stop the TCPDump command using `Ctrl-C`.
7. Find the session in the Tealeaf Canister.
   a. Verify there are `ReqCancelled=Client` hits in the sessions.
   b. Save the session as a `.tls` file. Remember to run Get Images first.
8. If the tcpdump can be guaranteed not to contain any user's sensitive information, you can arrange with http://support.tealeaf.com/ to FTP the tcpdump files to Tealeaf for analysis.
   a. If the tcpdump files contain production data and there is a chance that it may contain sensitive or personal information, contact http://support.tealeaf.com/ to arrange a time for remote analysis of the data.
9. After you have the captured session and the corresponding tcpdump of the raw traffic during the same slice of time that the session was captured, Tealeaf engineering should be able to help identify the root cause of the `ReqCancelled=Client` hits.

TL Queue fails to start and capture is disabled

If capture fails to be properly initialized, check the PCA `capture.log` for a line similar to the following:

```
Sep 3 15:33:51 tealeaf-dev reassd[15921]: TL Queue system failed to
   create (-10).
```

If the above line appears in `capture.log`, verify the following settings with the listed commands. Expected answers are listed below the command:

- **net.core.rmem_max:**

  ```
  sysctl -n net.core.rmem_max
  50000000
  ```

- **net.core.rmem_default:**

  ```
  sysctl -n net.core.rmem_default
  50000000
  ```

- **kernel.shmmax:**

  ```
  sysctl -n kernel.shmmax
  209715200
  ```

If the numbers displayed on your screen do not match the expected values, you can reconfigure these settings with the commands listed below:

```
sysctl -w net.core.rmem_max=50000000
sysctl -w net.core.rmem_default=50000000
sysctl -w kernel.shmmax=209715200
touch /usr/local/ctccap/var/startup
chown ctccap:ctccap /usr/local/ctccap/var/startup
chmod 644 /usr/local/ctccap/var/startup
```

When you restart the PCA, the TL Queue should initialize, and PCA capture should commence.

PCA Capture log contains "client_random missing" message

When the above message appears in the `capture.log`, it indicates that the PCA corrupted SSL traffic.

- The source may be many possible issues. There may be duplicate SSL keys that the PCA is using.

- For more information about the keys in use, see "PCA Web Console - SSL Keys Tab" in the *IBM Tealeaf Passive Capture Application Manual*.
- See "SSL Keys" in the *IBM Tealeaf Passive Capture Application Manual*.

PCA Capture log contains "Max SYN/WAIT limit reached"

The PCA logs this message as a result of processing the raw TCP/IP data.

Initiating a TCP connection requires three handshake packets:

- The source host sends a sync packet (SYN)
- The destination host then responds with its own sync packet (SYN+ACK)
- The source host sends an acknowledgement packet (ACK).

The PCA only attempts to reassemble a TCP connection into HTTP hit data if it sees all thee above packets. PCA maintains a finite buffer for caching a TCP connection's packets until all three handshakes are seen. If the PCA doesn't see all three before the buffer becomes full, those packets is flushed from the buffer and the above message is logged.

The three handshakes may not be seen for the following reasons:

- Packets lost on the network
- Packets corrupted
- Packets missing from the data stream that is fed to the PCA
- Packets arriving too far out of order
- Malicious SYN attacks

If you are seeing this message regularly, it could mean that the TCP connection cache is not big enough, and you can increase the cache's size. However, over time, these messages may appear in any Tealeaf system. If the messages occur with only sporadic frequency, the cache may be sized appropriately; the issue may be simple network "noise".

The incomplete TCP connections could was introduced at the network switch's SPAN monitor session without affecting you. It is also possible that you may have broken connections, which the SPAN session would pass on exactly as it was transmitted. From the point of view of a passive observer, there is no easy way to determine which of these occurred. However, visitors were probably not affected as far as browsing was concerned; if the connection was lost (no ACK packet received), then the browser automatically sends another request to the Web server for the wanted data, causing at worst nothing to happen. It should not be the source of an HTTP 500 error.

Other common sources of these errors are requests to an IP address that is not active. For example, suppose the PCA is configured to listen to IP port 8080, and there is a keep-alive service trying to get a page from each of 8 Web servers on port 8080. If only six Web servers are currently installed, the failure of the other two servers to respond may cause the "SYN/WAIT limit reached" message to appear.

When debugging related problems, the PCA's Max SynWait limit can be increased. For sites with many high-traffic spikes, this limit may need to be increased during installation/configuration. After increasing the maximum limit, monitor the capture log.

- If the Syn/Wait buffer fills up rapidly, and there are many dropped connections every minute, the problem is probably in the SPAN port configuration.
- If there are only a few missing Syn/Wait packets each hour, it is probably normal network behavior.

Replacing expired SSL certificate on the PCA

If an SSL certificate was added to the PCA to encrypt its Web Console, that certificate eventually expires. At expiration, a new certificate must be added to the PCA by following the instructions on "How to encrypt the PCA Console" solution at http://support.tealeaf.com/.

The new certificate must also be added to the certificate store on the PortalStatus server under the security context of the user running PortalStatus as a Windows Scheduled Task. If the certificate was not upgraded or is not added to the certificate store, the PortalStatus report emails have the message "Unable to connect to the Passive Capture server, server may be down or certificate is not installed."

PCA capture log error message " reassd: Couldn't find SSL session cache entry . . . "

Every SSL session is unique, and after a fully negotiated session occurs, its session key is cached for reuse during a long session. This message may occur in two cases:

- (More likely) The SSL session cache table size is too small. The default is 10,000. The PCA's statistics page shows the maximum number of entries used. If it reaches the configured maximum size and you are getting this message, then there was probably much SSL activity for a given time period that caused the still-needed session key to be dropped to make room for new entries.
  - To fix the issue, try increasing the maximum SSL session limit to 20,000 on the PCA Web console interface, on the Interface page, under the entry "Max SSL sessions to cache". See "Passive Capture Configuration via Web Console" in the *IBM Tealeaf Passive Capture Application Manual*.

    Note: If the issue isn't resolved after setting the limit to 20,000, then another issue is likely the cause and should be investigated. Do not continually increase the session cache limit beyond the maximum of 20,000.
- (Less likely) A PCA restart occurred, which caused the SSL session entry cache table to be cleared. Any SSL transactions that were in progress then would cause the "not found" message to be generated after the PCA software was reinitialized. Restarts occur when saving Web Console configuration changes or when the health/maintenance script issues a restart instruction based on its assessment of the PCA's health.

PCA could not create reveal object

After upgrading to Build 3324 or later, you may find an error message similar to the following one in the PCA capture log. The PCA may fail to start.

```
Oct 12 12:05:03 sh005 reassd[4763]: Couldn't create reveal object: 1
```

In the above case, the likely problem is that the private keys generated by the old PCA build cannot be validated by the new build. The basic solution is to remove the current PTL files from the appropriate directory, start the PCA to localize the problem to the PTL files, and then regenerate them from their source.

1. Log in to the PCA server.
2. Navigate to the following directory:

   ```
   /usr/local/ctccap/etc/capturekeys
   ```
3. Move the `.ptl` file or files in the directory to a location outside of the PCA installation.
4. Comment out any capture keys that are listed in `ctc-conf.xml`.
   a. Open `/usr/local/ctccap/etc/ctc-conf.xml` in a text editor and comment out all the `<CaptureKey>` nodes and their children. These nodes are children of the CaptureKeys node, which should remain enabled in the file.
      - To comment out a section of the `ctc-conf.xml` file use html style comments (`<!- ->`).

   b. Before example:

      ```
      <CaptureKeys>
                  <CaptureKey>
                      <Label>mykey</Label>
                      <PrivateKeyFile>/usr/local/ctccap/etc/mykey.ptl
                      </PrivateKeyFile>
                  </CaptureKey>
      </CaptureKeys>
      ```

   c. After example:

      ```
      <CaptureKeys>
      <!--            <CaptureKey>
                      <Label>mykey</Label>
                      <PrivateKeyFile>/usr/local/ctccap/etc/mykey.ptl
                      </PrivateKeyFile>
                  </CaptureKey> -->
      </CaptureKeys>
      ```
5. Restart the PCA:

   ```
   tealeaf restart
   ```
6. When the PCA has restarted, verify through the Web Console that all PCA processes are working and that data is being passed to the appropriate targets.
   - See "PCA Web Console - Summary Tab" in the *IBM Tealeaf Passive Capture Application Manual*.
   - See "PCA Web Console - Delivery Tab" in the *IBM Tealeaf Passive Capture Application Manual*.
7. If PCA operations have been verified, then the problem has been localized to the problematic `.ptl` keys.
   - If the PCA still fails to start, the problem may be elsewhere. You should retain the moved `.ptl` files until you can troubleshoot the problem completely. See “Troubleshooting a passive capture application server (PCA server)” on page 5.
8. For more information on regenerating the PEM and PTL files, see "Exporting the SSL Private Key" in the *IBM Tealeaf Passive Capture Application Manual*.
9. After you have regenerated the PTL keys, store them in the directory listed above.
   - If the PTL keys are saved in the same location with the same names as the originals, uncomment the CaptureKey nodes in the `ctc-conf.xml` file.
   - If a new location and/or filename is used, the `.ptl` files can be added through the PCA Web Console or placed in the `capturekeys` directory to be automatically loaded.

– For more information on loading through the console, see "PCA Web Console - SSL Keys Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

10. Restart the PCA.
11. If the PCA is able to decrypt SSL traffic, then the PTL files generated by the old build and moved out of the directory can be deleted.

`ReqCancelled` pages increase significantly after PCA reboot

Note: This solution only applies to TeaLeaf Passive Capture versions before 2058.

After a reboot, the number of `ReqCancelled` pages may increase significantly. The Linux kernel resets two memory parameters to default values after a reboot. Correcting this behavior requires the following steps:

- Change the running values:
  1. SSH to the PCA and login (or su) to root.
  2. Stop the Tealeaf capture software using the following command:
     ```
     service tealeaf-pca stop
     ```
  3. Issue the two following commands:
     ```
     sysctl -w net.core.rmem max=50000000
     sysctl -w net.core.rmem default=50000000
     ```
  4. Start the Tealeaf capture software using the following command:
     ```
     service tealeaf-pca start
     ```
- Edit the `sysctl.conf` file so the correct values are loaded after a reboot. Edit `/etc/sysctl.conf` and add:
  ```
  net.core.rmem max=50000000
  net.core.rmem default=50000000
  ```
- Upon reboot, Linux loads up the new settings in `/etc/sysctl.conf`.

Note: Versions of PCA Capture SW 2058 and later automatically check and configure these values when started. If you are running one of these versions, you do not need apply these workarounds.

You can also make the `sysctl.conf` file change now, and later upgrade to 2058 or higher. You will NOT have to remove the `sysctl.conf` file changes when you upgrade.

How to drop files of a specified type

You can configure the PCA to drop hits that are files of a specific type, which is useful for eliminating large binary file types, such as .PDF, from being captured and processed by Tealeaf. In the following sets of steps, response files with a specific extension are removed from capture, yet the request in which the files are queried remains in the capture stream.

Configure the PCA to capture a non-standard file extension

1. Open the PCA web console.
2. Click the **Pipeline** tab.
3. Add the wanted extension to the Included File Extensions list.
   a. Type the wanted extension into the **Included File Extensions** text field.
   b. Click **Add**.

c. Click **Save Changes**.

Configure the PCA to capture a non-standard MIMEtype

1. Open the PCA web console.
2. Click the **Pipeline** tab.
3. Add the wanted extension to the Capture All Mimetypes list.
   a. Type the wanted MIMEtype into the Capture All Mimetypes text field.
   b. Click **Add**.
   c. Click **Save Changes**.

Configure the PCA to drop the response for hits with a specified extension

1. Open the PCA web console.
2. Click the **Rules** tab.
3. Create a drop response action:
   a. In the Actions section, click **Add**.
   b. Complete a name for the action.
   c. Set the action to `DropResponse`.
   d. Click **Save Changes**.
4. Create a rule that uses the created action:
   a. Click the **Insert Rule 1** button.
   b. Set `ReqField` to `TL_URLEXT`.
   c. Set `ReqOp` to `=`.
   d. Set `ReqVal` to the wanted extension.
   e. To add the drop response action to the rule, select the action and click **Add**.
   f. Click **Save Changes**.

Freeing Up PCA disk space

By default, the IBM Tealeaf CX Passive Capture Application is installed on the `/usr` partition. Through various messages, you might receive indication that the partition is full. You might complete the following steps to verify that the partition is full and take measures to free up disk space.

1. Verify that the `/usr` is running out of space. To verify disk space on all available partitions, run the following command on the Linux server hosting the PCA:

   ```
   df -h
   ```

   - Verify the disk space available on the `/usr` partition.
2. On the server, navigate to `/usr/local/ctccap/bin-debug`. Look for files whose name begins with `core`. These core dump files can grow large and should be deleted or moved to another location for issue resolution.
   - To search for all core dump files, navigate to the `ctccap` directory and run the following command:

     ```
     find /usr/local/ctccap/ -name "core*" -print
     ```

   - Where possible, remove these files to free up disk space.

Troubleshooting failover

The PCA can be configured to failover from a master IBM Tealeaf CX Passive Capture Application to a slave machine. In the event of network outage, system failover, or other interruption on the master machine, the slave machine becomes the active machine and begins capturing traffic.

- See "PCA Web Console - Failover Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

If you configured failover and are experiencing difficulties, you may find troubleshooting tips and steps in this section to help you resolve the issue.

PCA failover mode requires two PCAs, designated Master and Slave. Typically, the **Failover** tab in the PCA Web Console enables defining the Master/Slave PCAs' IP/Port Addresses.

Note: Use their respective physical IP addresses instead of their host names.

Note: Both the Master and Slave PCAs must be running the same PCA build.

For proper identification of which is Master or Slave, the PCA examines the `/etc/hosts` file to find its IP address based on its local hostname. It then matches that IP address to the failover Master/Slave IP entries for assignment.

By running the Linux cmd `hostname`, the host name that is listed must be in the `/etc/hosts` file with its corresponding IP address. For example, executing `hostname` returns the pca01machine, then an entry similar to the following should appear in the `hosts` file.

```
10.10.100.1    pca01machine
```

Note: The wanted hosts file entry should be the first entry in the file.

If no match is found in the hosts file, then Failover mode fails to start. An error message appears in the `capture.log` file:

```
Both MasterAddress and SlaveAddress must be specified in configuration file.
```

Start, stop, and restart of failover

The correct starting order for Master-Slave failover PCAs is as follows:

- Start Master PCA first.
- After the Master PCA successfully started, start the Slave PCA.

  Note: If you are experiencing issues with correct operation of failover, you should always use the command line to manually stop and start it.

Note: The PCA may not correctly indicate the failover active state if the service was enabled or disabled multiple times. This is a known issue affecting all PCA builds.

Restarting Failover Through the web Console

The Failover service can be started and stopped through the **Failover** tab in the CX PCA web console.

1. Open the CX PCA web console on the slave CX PCA server. For more information, see "Open PCA web Console" in the *IBM Tealeaf Passive Capture Application Manual*.

2. Click the **Failover** tab.
3. Click **Restart failover**.
4. Click **Save changes** to apply your changes.
5. Repeat this procedure on the master CX PCA server.

After the failover status is reset, the CX PCA web console displays the new failover status. For example, the master CX PCA reports the status as `Failover is active (master)` and the slave CX PCA reports a status of `Failover is active (slave)`.

At the command line

Older versions of the PCA Web Console do not correctly accept changes to failover. When in doubt, you must manually stop and start the failover service from the command line.

1. Stop the failoverd service:
   ```
   tealeaf stop failoverd
   ```
2. Confirm failoverd was stopped and removed by running the following command:
   ```
   tealeaf ps
   ```
3. Review `capture.log` messages for more information.
4. Make changes as necessary.
5. To start the failoverd service:
   ```
   tealeaf start failoverd
   ```

Determining failover state

Use the following sections to help determine PCA failover issues.

Master PCA stats

Below, you can review the statistics that are published by the PCA for the master failover server in various failover states. These statistics are published in the failover section of the **Statistics** tab.

- When the `Node state` statistic is set to `active`, the server is delivering hits to other Tealeaf servers.
- See "PCA Web Console - Statistics Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

The Node state is active, meaning it is delivering hits:

Value	Statistic
`master`	Node role
`active`	Node state
`running`	Capture state
`yes`	Failover active

The Master wasforced to failover to the slave. Master was stopped:

Value	Statistic
`master`	Node role
`passive`	Node state

stopped
Capture state

yes Failover active

See "PCA Web Console - Statistics Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

Master PCA failover log messages

Below you can review log messages for various failover conditions for the failover master PCA server:

- These log messages appear in `capture.log`.
- Any `capture.log` messages that refer to `peer` is a reference to the other PCA, not the local one.

Failover is disabled:

```
TLAPI: Failover is disabled. Delivery is always enabled.
```

Failover Master is in active delivery state:

```
TLAPI: Failover is enabled. Delivery is currently enabled.
```

Failover service falls over to the PCA slave machine:

```
failoverd: Peer node is down (connection refused).
failoverd: Peer node is alive.
failoverd: Capture has stopped. Initiating failover to peer.
failoverd: Delivery stopped.
```

Failover failed back to the PCA master machine:

```
failoverd: Requesting failback from peer.
failoverd: Delivery started.
```

Slave PCA statistics

Below, you can review the statistics that are published by the PCA for the slave failover server in various failover states. These statistics are published in the failover section of the **Statistics** tab.

- When the `Node state` statistic is set to `active`, the server is delivering hits to other Tealeaf servers.
- See "PCA Web Console - Statistics Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

Master is running in active Node state. Slave is in passive state (no delivery):

Value **Statistic**

slave Node role

passive
Node state

running
Capture state

yes Failover active

Slave PCA is running, and Master PCA is stopped or non-existent (Failed over to Slave PCA):

Value **Statistic**

slave Node role

active Node state

running
Capture state

yes Failover active

When Slave failbacks to Master, the Node state returns to passive again.

See "PCA Web Console - Statistics Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

Slave PCA failover log messages

Below you can review log messages for various failover conditions for the failover master PCA server:

- These log messages appear in `capture.log`.
- Any `capture.log` messages that refer to `peer` is a reference to the other PCA, not the local one.

Failover is disabled:

```
TLAPI: Failover is disabled. Delivery is always enabled.
```

Failover Slave is in passive Node state:

```
TLAPI: Failover is enabled. Delivery is currently disabled.
```

Failover Slave took control:

```
failoverd: Received TakeControl request from peer. Taking control.
failoverd: Delivery started.
```

Slave failed back to Master server:

```
failoverd: Received ReleaseControl request from peer. Releasing control.
failoverd: Delivery stopped.
```

On restart of Web Console, ulimit core file size limit cannot be modified error

When the PCA Web Console is started or restarted, the following error may appear in the PCA log file:

```
tealeaf/bin/tealeaf: line 323: ulimit: core file size: cannit modify limit:
   operation not permitted.
```

The ulimit command is a Linux system utility. In most cases, this utility is run by the root user when the PCA Tealeaf script is first run. Then, the PCA is restarted using the PCA user, which is the `ctccap` user by default.

When the PCA starts using the non-root user, it defines the default system-level settings using ulimit, specifically to allow unlimited core file size. PCA uses this utility to provide core files when a PCA crash occurs for troubleshooting activities. Depending on the security, system policy that is used for the installed Linux OS, as well as platform and version dependencies, ulimit may not work as expected from the non-root user PCA perspective.

You can see what the default setting is when running under the non-root user:

```
su ctccap
ulimit -c
```

The expected line to look at is this:

```
core file size (blocks, -c) unlimited
```

In some Linux environments, the behavior of the ulimit utility is modified from which is expected by the PCA Tealeaf script that is used at startup, and the ability to set the core file size to an unlimited size is blocked.

To fix this issue, complete the following steps:

1. Edit the following configuration file, which is used to set the default ulimit core limits (`-c` option) for specific users.

   Note: The PCA user should be used to perform the file modifications. By default, this user is `ctccap`.

   ```
   /etc/security/limits.conf
   ```

2. To set the user's core setting to unlimited for both hard and soft limits, add the following two lines:

   ```
   ctccap hard core unlimited
   ctccap soft core unlimited
   ```

   Where:

   - `ctccap` corresponds to the PCA user.

3. Verify that there are no lines in the file that set the core setting for all users (`*`).
4. Save the file.
5. Restart the PCA.
6. Check the log files to confirm the issue was resolved.

PCA fails to start after adding a NIC

After you added a network interface card (NIC), the PCA might fail to start, resulting in an error message similar to the following error message.

```
Apr 13 10:27:09 tealeaf2 deliverd[5757]: Ending main loop with 0.
Apr 13 10:27:09 tealeaf2 deliverd[5757]: main(), Exiting with 0
Apr 13 10:27:09 tealeaf2 captured[5740]: Restarting too rapidly (0 seconds).
   Shutting down.
Apr 13 10:28:23 tealeaf2 tealeaf: info: Starting:
   /usr/local/ctccap/bin-debug/failoverd -q
Apr 13 10:28:23 tealeaf2 tealeaf: pem2ptl: error: Please specify the name
   of one or more PEM files to encrypt.
Apr 13 10:28:23 tealeaf2 tealeaf: info: Starting:
   /usr/local/ctccap/bin-debug/captured -P
Apr 13 10:28:23 tealeaf2 captured[6173]: Captured starting:
   revision 1277489920
Apr 13 10:28:23 tealeaf2 reassd[6182]: OpenSSL hw engine(0): None
Apr 13 10:28:23 tealeaf2 reassd[6182]: Couldn't create reveal object: 1
Apr 13 10:28:23 tealeaf2 reassd[6182]: Exiting
Apr 13 10:28:23 tealeaf2 captured[6174]: Caught signal (17). Restarting.
Apr 13 10:28:23 tealeaf2 reassd[6176]: OpenSSL hw engine(0): None
Apr 13 10:28:23 tealeaf2 deliverd[6184]: Ending main loop with 0.
Apr 13 10:28:23 tealeaf2 deliverd[6184]: main(), Exiting with 0
Apr 13 10:28:23 tealeaf2 reassd[6178]: OpenSSL hw engine(0): None
Apr 13 10:28:23 tealeaf2 reassd[6176]: Couldn't create reveal object: 1
Apr 13 10:28:23 tealeaf2 reassd[6176]: Exiting
Apr 13 10:28:23 tealeaf2 reassd[6178]: Couldn't create reveal object: 1
Apr 13 10:28:23 tealeaf2 reassd[6178]: Exiting
Apr 13 10:28:24 tealeaf2 captured[6174]: Restarting too rapidly (0 seconds).
   Shutting down.
```

```
Apr 13 10:34:32 tealeaf2 tealeaf: info: Stopped httpd(5760).
Apr 13 10:34:32 tealeaf2 tealeaf: info: captured is not running.
Apr 13 10:34:32 tealeaf2 tealeaf: info:
   Starting: /usr/local/ctccap/bin-debug/failoverd -q
Apr 13 10:34:32 tealeaf2 tealeaf: pem2ptl: error: Please specify the name
   of one or more PEM files to encrypt.
Apr 13 10:34:32 tealeaf2 tealeaf: info: Starting:
   /usr/local/ctccap/bin-debug/captured -P
Apr 13 10:34:32 tealeaf2 captured[9446]: Captured starting:
   revision 1277489920
Apr 13 10:34:32 tealeaf2 tealeaf: info: Starting: /usr/local/ctccap/bin/httpd
Apr 13 10:34:32 tealeaf2 reassd[9449]: OpenSSL hw engine(0): None
Apr 13 10:34:32 tealeaf2 reassd[9449]: Couldn't create reveal object: 1
Apr 13 10:34:32 tealeaf2 reassd[9449]: Exiting
Apr 13 10:34:32 tealeaf2 captured[9447]: Caught signal (17). Restarting.
Apr 13 10:34:32 tealeaf2 deliverd[9458]: Ending main loop with 0.
Apr 13 10:34:32 tealeaf2 deliverd[9458]: main(), Exiting with 0
Apr 13 10:34:32 tealeaf2 captured[9447]: Restarting too rapidly (0 seconds).
   Shutting down.
```

This issue might be caused by the PTL keys that are installed on the PCA. In some cases, these keys are encrypted using aspects of the addresses of the NIC cards.

The solution is to re-create the PTL keys from the PEM clear-text versions. The PCA can perform this recreation for you automatically.

- See "Encrypted Key Setup" in the *IBM Tealeaf Passive Capture Application Manual*.

Error System.Web.HttpRequestValidationException

```
Symptom:
Error System.Web.HttpRequestValidationException: A potentially dangerous
Request.Cookies value was detected from the client (ErrorId=
Resolution:
```

In IIS Manager, verify that the TeaLeaf IIS capture filter is installed at the global websites level rather than in a specific website.

What does the `PSIO_SIGCAUGHT` error mean?

Whenever you see the `PSIO_SIGCAUGHT` error, you need to look earlier in the capture `logfile` for the true cause. Those `PSIO_SIGCAUGHT` errors are from pipelined trying to do communicate with delivered, but that communication was interrupted by a signal.

In most cases, that signal is from captured telling pipelined to shut down. Captured tells pipelined to shut down for various reasons, such as the following:

- Another child process of captured crashed or exited.
- Captured is shutting down.
- A configuration change in the web console told captured to restart its child processes.

In the following sample snippet from a `capture.log`, you can see a message from reassd saying it failed to create the shared-memory queue for hit processing (TLQueue):

```
Jul 24 17:08:10 abc reassd: TL Queue system failed to create (28).
Jul 24 17:08:10 abc reassd: Exiting
Jul 24 17:08:10 abc captured: Caught signal (17). Restarting.
Jul 24 17:08:11 abc pipelined: Evaluation of tltHeartbeatCallback returned
ERROR(1): tltSocketWrite: Failed to write socket command(3): PSIO_SIGCAUGHT
```

In build 3101, the PSIO_SIGCAUGHT message was changed to minimize exposure to this misleading message.

Why is my messages directory full of crond(pam_unix) session closed for user root messges

Why is my /var/log/messages full of message like the following:

```
crond(pam_unix) ... session closed for user root
```

On IBM Tealeaf CX Passive Capture Application servers running Red Hat Enterprise Linux 4, messages of the following types are common in /var/log/messages:

```
# fgrep crond /var/log/messages | tail -5
Sep  6 16:23:02 ganymede crond(pam_unix)[23922]: session closed for user root
Sep  6 16:24:01 ganymede crond(pam_unix)[23971]: session opened for user root
   by (uid=0)
Sep  6 16:24:03 ganymede crond(pam_unix)[23971]: session closed for user root
Sep  6 16:25:01 ganymede crond(pam_unix)[24020]: session opened for user root
   by (uid=0)
Sep  6 16:25:02 ganymede crond(pam_unix)[24020]: session closed for user root
```

On machine ganymede, all the messages are these crond(pam_unix) messages. Consider the following count of the lines with these messages:

```
# fgrep crond /var/log/messages | wc -l
13233
# wc -l /var/log/messages
13241 /var/log/messages
```

The inclusion of the pam_unix text hints that this message is coming from the PAM configuration for the machine. Looking in /etc/pam.d reveals a crond file whose contents describe it as, "The PAM configuration file for the cron daemon."

Verify if this configuration file is owned by some RPM package:

```
# rpm -qf /etc/pam.d/crond
vixie-cron-4.1-36.EL4
# rpm -ql vixie-cron | fgrep pam
/etc/pam.d/crond
```

This same /etc/pam.d/crond file does not exist in an RHEL 3 Update 5 machine:

```
# rpm -q vixie-cron
vixie-cron-3.0.1-76_EL3
# rpm -ql vixie-cron | fgrep pam
#
```

It seems likely that these crond(pam_unix) messages start appearing in RHEL 4 because of a change that is introduced by the newer vixie-cron package.

The manual page for crond on RHEL 4 also documents this relationship (from "man crond"):

```
PAM Access Control
     On  Red  Hat systems, crond now supports access control with PAM - see
     pam(8).   A  PAM  configuration  file  for  crond  is   installed in
     /etc/pam.d/crond. crond loads the PAM environment from the pam_env
     module, but these can be overriden by settings in the crontab file.
```

An Internet search for crond(pam_unix) messages /etc/pam.d/crond reveals that many users are running into this situation with the PAM configuration in vixie-cron triggering messages in /var/log/messages. Here are two solutions:

Solution 1: suppress the messages

The only point to having pam_unix in the session stack is to provide this level of logging. This logging can be suppressed by commenting out the line in /etc/pam.d/crond that starts the system-auth session stacking and uncommenting the session pam_limits line, since the only pam modules started in the system-auth session stack are pam_unix and pam_limits.

On a RedHat ES4 system, the /etc/pam.d/crond file should contain as follows:

```
#
# The PAM configuration file for the cron daemon
#
#
auth       sufficient pam_rootok.so
auth       required   pam_stack.so service=system-auth
auth       required   pam_env.so
account    required   pam_stack.so service=system-auth
account    required   pam_access.so
#session    required   pam_stack.so service=system-auth
session    required   pam_loginuid.so
# To enable PAM user limits for cron jobs,
# configure /etc/security/limits.conf and
# uncomment this line:
session  required   pam_limits.so
#

Restart crond daemon:
  service crond restart
```

Solution 2: Redirect messages to a different log file

1. Edit /etc/syslog.conf as follows:

   ```
   # Log anything (except mail) of level info or higher.
   # Don't log private authentication messages.
   *.info;mail.none;local5.none;authpriv.none;cron.none;auth.!=info
   /var/log/messages
   ```

   Note: The config line has just auth.!=info appended above.
2. Log cron auth messages in a separate file. Add the following lines:

   ```
   auth.info
   /var/log/cron.auth
   ```
3. After service syslog restart the syslog messages caused by cron appear in /var/log/cron.auth only.

IIS 6 Capture filter for Siebel application does not capture correctly

The error as shown by IBM Tealeaf CX RealiTea Viewer is "Invalid Deflate Data".

The Siebel ISAPI filters are performing compression. For the IIS capture filters to handle this compression, add the following two lines to the [Globals] section of the TeaLeafIIS6.cfg file, on all the Siebel Web servers where the Tealeaf IIS capture filter is installed:

```
IgnoreChunked=False
IgnoreGZip=False
```

After making this change, you must restart IIS on those Web servers.

Information to provide customer support

If you are still experiencing difficulty with the IBM Tealeaf CX Passive Capture Application or its host server, you should find and review the following information before contacting Tealeaf Customer Support:

- What version of the PCA is running?
- What NICs are being listened on?
- Did this problem appear as part of a new install or an upgrade? If this occurred during an upgrade, was this problem also available before the upgrade?
- Run `tcpdump`. Was the memory dump filtering for anything or was it just on the listening NICs?
- Provide root access or personnel who have root access to the Tealeaf PCA.

Additional tcpdump commands

Extended tcpdump Details

To get the expanded details for tcpdump, run `man tcpdump` on the PCA server or get it from the http://www.tcpdump.org/ website. The short form is:

```
tcpdump -ni NIC -s0 -w dumpfile filter_string
```

where

- `NIC` is the sniffer/capture NIC's device name. To see a list of NIC device names, use the `ifconfig` command or the Interface page of the PCA's Web UI, which is located by default at `http://<PCA>:8080/interface.php` or `https://<PCA>:8443/interface`. `<PCA>` is the host name or IP address of the PCA.
- `dumpfile` is the name of the file to which the dump is written.
- `filter string` is the IP/port expression that you want to capture.

If you want to run a test to validate, remove the `-w` option and its argument; output is then printed to the console.

Test capture sniffer ports for traffic

To use tcpdump to determine if capture sniffer ports are receiving traffic:

```
tcpdump -ni eth1 | more
```

To look for specific traffic:

```
tcpdump -ni bge0 host <IP address> and port <IP port number> | more
```

Where `<IP address>` and `<IP port number>` might be 10.10.10.2 and 80, for example.

Manual Capture of tcpdump to a File

The manual method of using tcpdump to capture HTTP(S) traffic to a dump file is the following:

```
tcpdump -ni bge0 -s0 -w dumpfile host ipAddr and port \(80 or 443\)
tcpdump -ni bge0 -s0 -w dumpfile host ipAddr and port '('80 or 443')'
tcpdump -ni bge0 -s0 -w dumpfile net xxx.xxx.xxx.0/24 and port 80
```

Using existing tcpdump file to extract specific traffic to another dump file:

```
tcpdump -nr <dumpfile> -s0 -w <new_file_name> host <ip address> and \
port <port number>
```

Display tcpdump Extra Header

If tcpdump is unable to filter any traffic, whereas running tcpdump unfiltered does show traffic, then the problem may be due to an extra Ethernet header being added, which is typically a VLAN header: 802.1Q VLAN. Use the `tcpdump -e` option to see this extra header information, which should look like the following:

```
. . . ethertype 802.1Q, length 64: vlan 128, p 0, ethertype IPv4,
IP 192.168.128.42.8001 > 192.168.128.90.20700:
```

Port Filtering

Trying to filter using tcpdump fails. An example is to filter on a known port number, such as the following:

```
tcpdump -ni eth2 port 8001
```

If tcpdump is unable to provide a filtered output, then the passive capture software is not able to do so either.

If it is VLAN-type traffic, use the `vlan` expression operator as part of the filter expression:

```
tcpdump -ni eth2 vlan and port 8001
```

Other examples of filtering with VLAN packets:

```
tcpdump -nr tst.dmp 'ether[12:2] = 0x8100'
tcpdump -nr tst.dmp vlan and ip and port 8001
```

To show both types of traffic:

```
tcpdump -nr tst.dmp ip or vlan
```

Display client side traffic only for specific IP address

The following command only displays traffic from the client side, which are requests made from browser to web server from the specified IP address.

```
tcpdump -nieth1 src host ipAddr and port #nmbr
```

Display All Client-Side Traffic

The following command displays all client side traffic that attempts to connection on the specified port number. All requests from all clients on the specified port are reported.

```
tcpdump -nieth1 dst port #nmbr
```

Display only SYN and FIN Packets

The following command only displays the SYN and FIN packets in the `tst.dmp` file. This command can be used to verify that in addition to both directions of traffic being present, the PCA is also seeing the successful start and stop of the tcp connections.

Valid output that is expected for each connection would be a SYN packet in each direction at the start of the connection and a FIN packet in each direction at the end of the connection.

```
tcpdump -nr tst.dmp 'tcp[tcpflags] & (tcp-syn|tcp-fin) != 0'
```

Using tcpdump to determine if traffic is not bidirectional

Execute the following or similar tcpdump command:

```
tcpdump -ni <NIC> -s0 host 192.168.149.201 and port 443
```

Where:

- `-s0` - do not truncate large packets. Without this option, unidirectional traffic recording does not occur.
- `<NIC>` - the device name of one of the PCA's capture NIC's.

The command should generate output similar to the following:

```
16:10:37.271214 IP 12.130.155.248.4863 > 192.168.149.201.https:
S 4007169894:4007169894(0) win 64512 <mss 1460,nop,nop,sackOK>
```

There are columns for each line representing a packet:

- timestamp
- network protocol
- source IP address.IP port
- destination IP address.IP port
- TCP flags
- selected information about the packet

The preceding example is a TCP/IP connection handshake packet showing the client-to-server initial SYN handshake.

There should be a second packet with the two IP.port columns reversed to indicate the server to client SYN handshake packet. If there is no second packet, then the SPAN port is not providing bidirectional traffic to the PCA's capture NIC.

Additional topics

For more information IBM Tealeaf CX Passive Capture Application topics, see "Appendix - Additional Passive Capture Topics" in the *IBM Tealeaf Passive Capture Application Manual*.

Locating and updating your license key

For new installs and upgrades, Tealeaf provides a license key, which must be entered during the installation or upgrade process. The license key determines the features that are available through the Tealeaf Portal.

Note: If you are retrieving your license key for submission to Tealeaf Customer Support and do not need to change it, select **Help** > **About IBM Tealeaf CX Portal** in the Portal menu. Copy the value next to License Key. For more information, see Tealeaf http://suppor.tealeaf.com.

For more information about how to apply the license key through TMS, see "Managing Your Tealeaf License Key" in the *IBM Tealeaf cxImpact Administration Manual*.

Applying the license key through the Windows registry

If the Portal and TMS are not available, you can apply the Tealeaf product license key through the Windows registry.

For more information about how to apply the license key through the Windows registry, see "Managing Your Tealeaf License Key" in the *IBM Tealeaf cxImpact Administration Manual*.

Default IIS permissions

When Tealeaf is installed or upgrade, the required permissions are automatically configured for IIS. If permissions are modified for enterprise reasons, Portal functionality can be compromised. Behaviors are varied.

If you believe that the IIS permissions in use for the Portal application or Portal API was modified from the following set, rerun TLIISUtility.

To run TLIISUtility:

Note: Rerunning this utility changes the permissions on the required folders to the above settings. Rerunning otherwise has no impact on the Portal.

1. Open a command line on the Portal server.
2. Navigate to the following directory:
   `<Tealeaf_install_directory>\Tools\Install`
3. Launch `TLIISUtility.exe`.
4. Configure permissions per the following requirements:
   a. **Required Permissions for IIS6:**
      - `http://support.microsoft.com/kb/812614/`
   b. **Required Permissions for IIS7:**
      - `http://support.microsoft.com/kb/981949/`
5. Restart IIS.

A 1970 date shows up in the "From date" in a chart

It is likely that a captured and processed hit contained `0` as one of the timestamp. A `0` timestamp corresponds to 1/1/1970. If the hit triggered an event, the event timestamp is also 1/1/1970.

NT-authenticated users cannot log in to the Portal

In some cases, Tealeaf users who are attempting to access the Portal using NT credentials cannot log in.

Note: This issue was observed on Safari browser 5.1 only. Safari is not supported by Tealeaf. For more information about supported browsers, see "Logging in to the Tealeaf Portal" in the *IBM Tealeaf cxImpact User Manual*.

The Safari browser and potentially other browsers may have issues with the Negotiate (Kerberos) authentication provider.

The workaround is to remove the Negotiate authentication provider from the NT authentication provider list that is served by IIS. After these steps are complete, IIS no longer uses the Negotiate authentication provider and instead uses the next one on the IIS list.

Note: These steps apply to IIS7 only. Steps for IIS6 may vary. Consult the documentation that is provided with IIS6.

1. To remove Negotiate from the providers list, login to the Portal Server as an administrator.
2. Run the following command from a command shell:

   ```
   %systemroot%\system32\inetsrv\appcmd set
   config /section:windowsAuthentication /-providers.[value='Negotiate']
   ```

3. When this command is first run, the output should look something like the following:

   ```
   Applied configuration changes to section"
   "system.webServer/security/authentication/windowsAuthentication" for
   "MACHINE/WEBROOT/APPHOST" at configuration commitpath
   "MACHINE/WEBROOT/APPHOST"
   ```

4. If the command is run again, the output changes, as the Negotiate entry was already removed:

   ```
   ERROR ( message:Cannot find requested collection element. )
   ```

5. Perform an IIS reset.
6. Log in to the Portal using Safari.
7. As an extra test, you should verify replay through BBR.
   - See "CX Browser Based Replay" in the *IBM Tealeaf cxImpact User Manual*.

Some users can log in to the Portal, while others get a 400 - bad request error

When Tealeaf is configured to use NT Authentication and IIS has Integrated Authentication enabled, some users receive a 400 Bad Request error when trying to access the Tealeaf Portal using a DNS name. The error in the HTTPSYS error log (`HTTPERRx.log`) is the following:

```
400 - RequestLength
```

If Integrated Authentication is enabled in IIS and it negotiates Kerberos, some users may have a very large Kerberos ticket that exceeds limits in the request header.

- In one case, users were unable to access the Portal due to a large number of group memberships.

To fix:

1. Add registry keys on the Reporting Server to increase the Maximum Field Length and Maximum Request Bytes to accommodate the larger Kerberos tickets:

   ```
   HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\HTTP\Parameters\
      MaxFieldLength = 32768 (default = 16384)
   HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\HTTP\Parameters\
      MaxRequestBytes = 32768 (default = 16384)
   ```

2. Restart IIS.

For more information

- `Http.sys` registry settings for IIS (Microsoft KB 820129) @ http://support.microsoft.com/?kbid=820129
- Kerberos Authentication Problem with Active Directory @ http://blogs.technet.com/surama/archive/2009/04/06/kerberos-authentication-problem-with-active-directory.aspx
- HTTP 400 Bad Request Discussion @ http://www.issociate.de/board/post/314237/HTTP_400_Bad_Request.html

Differences between Total and Archived session counts in Portal

In the Total Activity report in the Portal, the Sessions count in the upper left may not match the Archived Sessions count in the lower right. Depending on the circumstances, either count may be the higher one.

The Total Sessions count is incremented whenever a new website session starts. The Archived Sessions count is incremented when sessions are marked for saving. This difference in definition may cause differences in count.

Around midnight, there can be more Total Sessions than Archived Sessions because some active sessions were not yet saved to the Long-Term Canister.

Similarly, a day can have more archived sessions than total sessions, if a burst of sessions that are started before midnight and ended and were saved some time after midnight.

Display problems in Internet Explorer

By default in Windows Server 2003 and Windows Server 2008, Internet Explorer enables a set of Enhanced Security features. If Enhanced Security is enabled in IE, you might experience some or all of the following issues:

1. **Tealeaf Portal:**
   - Portal menus are not displayed.
   - JavaScript errors of `getElementByID is null or not an object`
2. **PCA:**
   - Clicking buttons does not work in the web console.
3. **RTV:**
   - Unknown impacts

If you are encountering display problems in Internet Explorer, you can either disable the Enhanced Security features or add the Tealeaf components as Trusted Sites.

Registering Tealeaf Components as Trusted Sites

To improve performance, you can register the URLs for Tealeaf components as trusted sites in Internet Explorer.

1. In the Internet Explorer menu, select **Tools** > **Options...**.
2. Click the **Security** tab.
3. Click **Trusted Sites**.
4. Click **Sites...**.

5. Enter each of the following URLs that pertain to your Tealeaf installation, where <host_name> is the host machine for the component:
   - Tealeaf Portal:

     http://<host_name>/portal/
   - PCA Web Console:

     https://<host_name>:8443

     Note: If your Portal or PCA are communicating over non-default ports, you must provide the port number as part of the URL.
6. Click **Add**.
7. When you finished adding all relevant URLs, click **Close**. Then, click **OK**.

Removing or disabling enhanced security from Windows

Removing for Windows Server 2003

These steps remove the Internet Explorer Enhanced Security Configuration.

1. From the Windows Start menu, select ***Settings > Control Panel > Add or Remove Programs**.
2. Click **Add/Remove Windows Components**.
3. Select **Removing Internet Explorer Enhanced Security Configuration**.

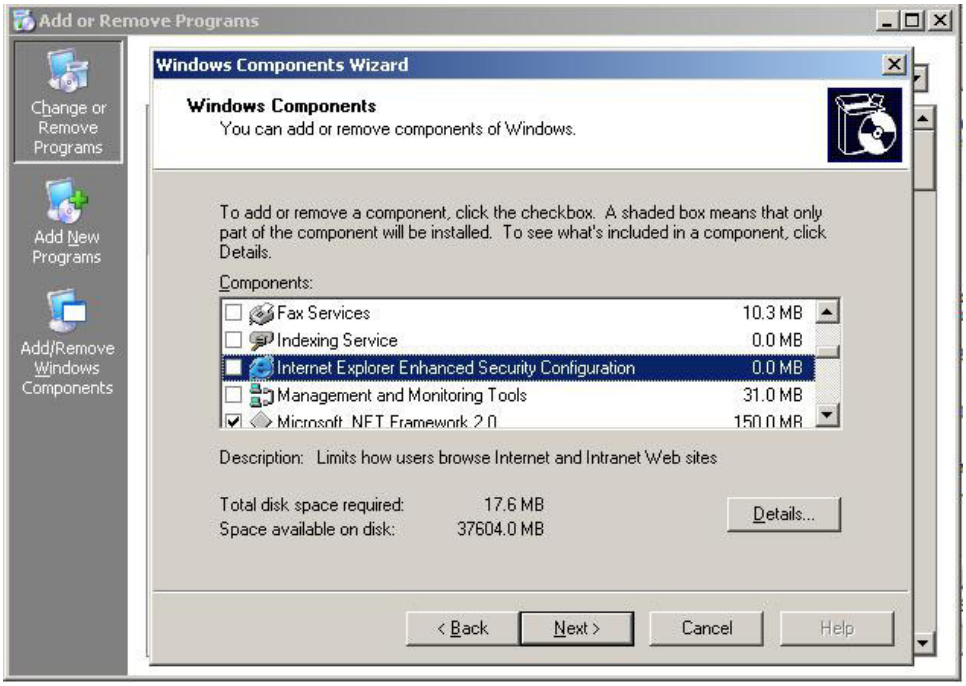

4. Click **Next**.
5. The component is removed.

Disabling for Windows Server 2008

These steps disable the Internet Explorer Enhanced Security Configuration within Windows Server 2008.

1. From the Windows Start menu, select **Settings** > **Control Panel** > **Administrative Tools** > **Server Manager**.
2. The Server Manager is displayed:

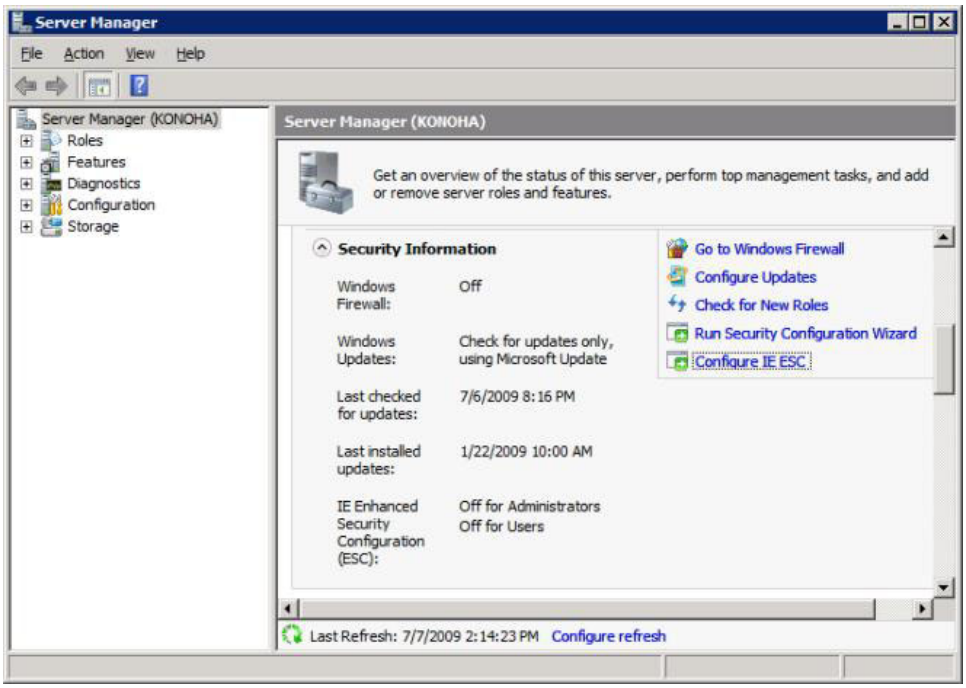

Windows Server Manager

3. Click the **Configure IE ESC** link.

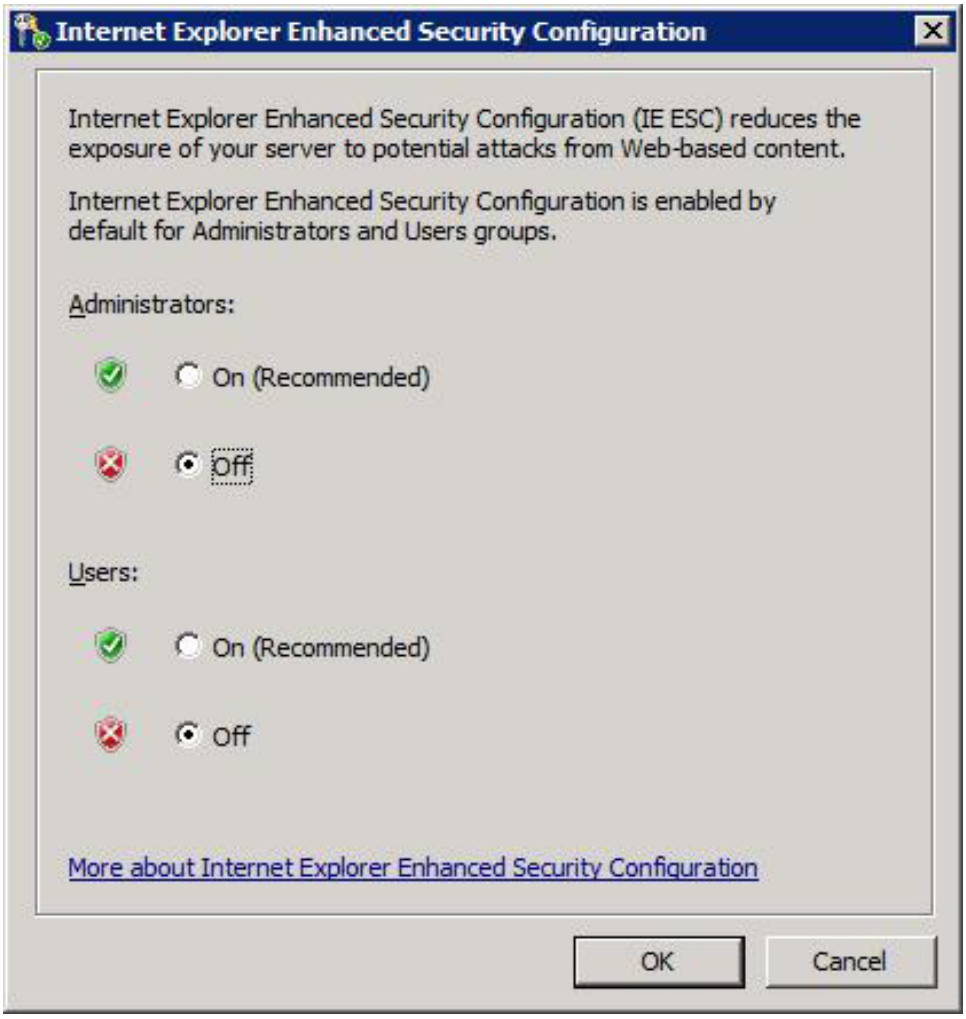

Internet Explorer Enhanced Security Configuration

4. Disable the ESC settings for administrators and users.
5. Click **OK**.
6. The features are disabled.

Stack overflow error when trying to log in to Portal on local machine using IE8

If you installed the Tealeaf Portal and Tealeaf databases on the local machine using Windows Server 2003 and a local instance of SQL Server 2005, you may receive the following error message when trying to browse to the Portal using Internet Explorer 8:

```
Stack overflow at line: 183
```

Verify that the following updates were completed:

- System is updated to .NET 3.5 SP1
- All required Windows updates were performed.
- All required and recommended IE8 updates were performed.

Enabling All Menu Access for Admin Group

If necessary, you can use the following SQL to enable access to all menu items in the Tealeaf Portal for the Admin user group.

Note: This SQL enables access to the menu items only. Some pages in the Portal have page-level permission restrictions, which are not enabled by this SQL.

Note: Performing these steps requires an IIS reset, which boots all users currently in the Portal. Perform these steps at an off-peak time.

1. Run the following SQL against the `TL_SYSTEM` database using the `TLADMIN` or `TLUSER` account:

```
USE TL_SYSTEM

DELETE FROM RS_PROFILE_ITEM WHERE PROFILE_ID = 1 AND OBJ_ID = 100
GO

INSERT INTO RS_PROFILE_ITEM
SELECT 1, 100, NODE_ID FROM RS_MENU
GO
```

2. Perform an IIS reset after the above is executed.
3. Users who are members of the Admin group can now access all available pages in the Portal menu.

Internet Explorer does not show Portal content below menu bar

(IBM Tealeaf cxImpact and IBM Tealeaf cxView only) Verify that the following two items were selected in the Advanced tab of Internet Explorer under Multimedia:

- Auto-resize of images
- Play animations in Web pages

The times in the Portal are off by some hours

The default Portal time zone is GMT-based, which may vary from what is expected for your Tealeaf implementation.

The steps to set the time zone are as follows:

1. Log in to the Portal as a Portal administrator.
2. Select **Tealeaf** > **Portal Management**.
3. Select the type of accounts to administer.
   - See "CX User Administration" in the *IBM Tealeaf cxImpact Administration Manual*.
   - See "cxView User Administration" in the *IBM Tealeaf cxImpact Administration Manual*.
   - See "cxReveal User Administration" in the *IBM Tealeaf cxReveal Administration Manual*.
4. Select the user who is experiencing time discrepancies.
5. From the Time Zone drop-down, select the user's correct time zone.
6. To complete the change, click **Save**.
7. To verify that time change was effectively applied, select **Active** > **Sessions** and check the timestamps for active sessions.

Portal Unable to Connect to Search Servers

If the Portal is unable to communicate with the Search Server(s), even when the destination server or servers is specified by IP netaddress or fully qualified domain name in the Portal, you should set the log level to 9 to debug the issue:

- **Pre-Release 7.2:** Select **Start** > **Programs** > **TeaLeaf Technology** > **TeaLeaf RealiTea Report Server** > **RealiTea Report Server Configuration** to open the TeaLeaf `ReportConfig.exe` utility to make the change.
  1. Then, restart the World Wide Web Publishing Service.
- **Release 7.2: or later** Make the change through TMS. In the Tealeaf Portal, select **Tealeaf** > **TMS**. For more information about TMS, see "Tealeaf Management System" in the *IBM Tealeaf cxImpact Administration Manual*.
  - See "Configuring the Report Server" in the *IBM Tealeaf CX Configuration Manual*.

When the log level is set to 9, the following error may appear in the Portal log:

```
Error creating the Web Proxy specified in the 'system.net/defaultProxy'
configuration section. ---> System.DllNotFoundException: Unable to load
DLL 'rasapi32.dll': A dynamic link library (DLL) initialization
routine failed. (Exception from HRESULT: 0x8007045A)
```

If the above error appears, .NET/ASP .NET may be inheriting a different set of proxy settings from those specified in the Portal. To fix this issue, the Portal must be configured to not inherit these settings.

Note: Verify that this change does not violate any enterprise requirements.

The solution is to add the following XML to the Portal's `Web.Config` file:

```
<system.net>
        <defaultProxy>
            <proxy usesystemdefault="false"/>
        </defaultProxy>
    </system.net>
```

The `Web.Config` file is stored in the following location:

```
<Tealeaf_install_directory>\Portal\WebApp
```

Note: After you fixed the issue, remember to reset the log level to its previous value.

Unable to display XML export to Excel

On a system with Microsoft Office installed, by default Windows attempts to open all XML documents using `MSOXMLED.EXE`. This application scans the document for the mso-application tag, which identifies the Office application to use to open the XML document.

If this association was changed to use, for example, `Notepad.exe`, then Windows cannot correctly open Office XML documents.

To verify the proper mapping:

1. Open Windows Explorer to any folder on your local computer.
2. In the Windows Explorer menu, select **Tools** > **Folder Options...**.
3. Click the **File Types** tab.

4. Select the **XML/XML Document** entry.
5. Click **Advanced**.
6. There should be two actions listed: `edit` and `open`.

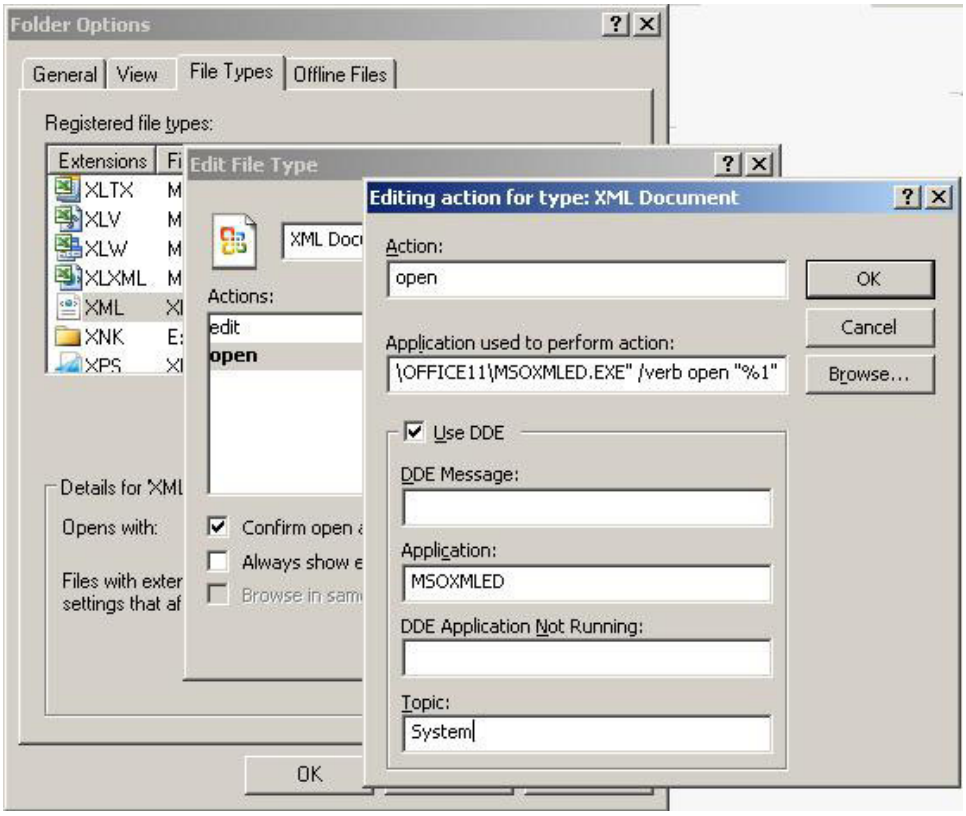

7. For each action, click **Edit** and review the value for `Application used to perform action`. They should similar to the following paths for 64-bit Windows:

```
"C:\Program Files (x86)\Common Files\Microsoft Shared\OFFICE12\MSOXMLED.EXE"
   /verb edit "%1"
"C:\Program Files (x86)\Common Files\Microsoft Shared\OFFICE12\MSOXMLED.EXE"
   /verb open "%1"
```

If the above paths do not point to the `MSOXMLED.EXE` application, XML documents cannot open properly in Excel. To remedy this problem, change the paths to point to the proper application.

Pre-Release 5.0 troubleshooting topics

The following topics apply to Release 4.x or earlier versions of Tealeaf.

Purging the erroneous date from the database

Note: In versions 4.5 and above, the drop-down menu that was populated from the database for selecting chart dates is no longer populated. The calendar popup is the only supported date input method.

Note: The remainder of this solution applies only to 4.0 and earlier systems.

To remove the erroneously dated record from the reports database, find the `UniqueID` of the event and then delete the records from the table MC DTS LOG:

1. To find the `UniqueID` of the event, do any of the following:

a. In the Portal's Chart wizard, hover over the link to create a chart of the event. The event's `UniqueID` is in the query string parameters.
b. Get the `EventID` value from the Event Editor.
c. Using the Viewer, select the event for which to search in the Basic search tab, and then copy the search to the Advanced search tab, where the search string shows the event's `UniqueID`.

2. To see the data, you must use SQL Server's Query Analyzer tool. Query Analyzer is part of a licensed SQL Server installation and is not included with MSDE. If your system uses MSDE, you can use ANY workstation with SQL Server installed to run SQL Query Analyzer and connect to the Tealeaf reports database:
   a. Start Query Analyzer.
   b. Issue the following query. Replace 1 with the correct `UniqueID` value:
      ```
      select dts from mc dts log where {{UniqueID}} = 1 order by dts
      ```
   c. If the "1/1/1970" record shows up, you can delete it with the following query. Replace 1 with the correct `UniqueID`:
      ```
      delete from mc dts log where {{UniqueID}} = 1 and dts < '1/1/1999'
      ```

Event Description Not Found when creating a chart

Note: This solution applies to Release 4.x or earlier versions of Tealeaf.

If you try to create a Portal chart of an event immediately after creating it in the Event Editor, you may receive this message:

```
Event Description Not Found
```

You may need to wait 5 - 30 minutes before you are able to create the chart. The length of time to wait depends on the Report Data Collector service's data collection time.

- If the system is busy, the Report Data Collector takes more than the default interval of 5 minutes to start a new cycle, because to the length of time it takes to collect the data from the Processing Server(s).
- A correctly sized system should not have a latency of more than 5 to 10 minutes to display new event descriptions.

If you examine the chart before the event description is present for a newly defined event, no data will be displayed in the chart until the next data collection run is completed and the event description is retrieved.

Portal home page generic error

Note: This solution applies to Release 4.x or earlier versions of Tealeaf.

The Portal home page may produce the following error message:

```
Portal Error
TeaLeaf Portal Version: 4.6.0.4655
An error has occurred in the TeaLeaf Portal Application.
Please contact the system administrator.
```

If you see this message, check the following for more information:

- The Windows Event Logs
- Set the log level up to 9:

1. **Pre-Release 7.2:** Select **Start** > **Programs** > **TeaLeaf Technology** > **TeaLeaf RealiTea Report Server** > **RealiTea Report Server Configuration** to open the TeaLeaf `ReportConfig.exe` utility to make the change. Then, restart the World Wide Web Publishing Service.
2. **Release 7.2: or later** Make the change through TMS. In the Tealeaf Portal, select **Tealeaf** > **Portal Management**. For more information about TMS, see "Tealeaf Management System" in the *IBM Tealeaf cxImpact Administration Manual*.
   - See "Configuring the Report Server" in the *IBM Tealeaf CX Configuration Manual*.

If you cannot diagnose the problem, do the following and check for information:

- Check whether the `TLAspUtil.dll` and `TLPrivacy.dll` files have Read & Execute permissions for Everyone.
- If using TeaLeaf's NT authentication feature, look for Active Directory groups that were specified in the Search Server Config utility for use in authentication.

Cannot get logon screen for Portal

Note: This solution applies to Release 4.x or earlier versions of Tealeaf.

When the following error message is displayed by the Portal:

```
TeaLeaf Portal Version: 4.5.0.4574
Error Number: -2147467259
Source:
File: /LM/W3svc/1/Root/Portal/global.asa
Description: .
Line: 34
```

Verify both `TLAspUtil.dll` and `TLPrivacy.dll` have Read and Read & Execute permissions for the IUSR <Portal server name> user or for the relevant Windows NT / Active Directory groups if using the NT authentication integration feature.

Error - 593-Non ADMIN user blocked from logon

Note: This solution applies to Release 4.x or earlier versions of Tealeaf.

This error occurs when `TLTMaint.exe` was locked the Canister for data validation or recovery.

To clear this error:

1. Stop all Tealeaf services using the Start menu shortcut. You may keep the Transport Service running.
2. Using a command prompt window, navigate to the Tealeaf installation directory.
3. At the command line, enter the following command:

   ```
   TLTMaint -NoServer
   ```
4. Wait for TLTMaint to finish execution. If it reports unrecoverable errors in the Windows Event Log, contact http://support.tealeaf.com/.
5. After TLTMaint completes, use the Start Tealeaf Services shortcut on the Start menu.
6. After the services were restarted, verify that Portal login works correctly.

Cannot log in to the Portal using any user ID

Note: This solution applies to Release 4.x or earlier versions of Tealeaf.

When you try to log in to the Portal, the following error message is displayed:

```
Database Error The web server experienced an error with the database.
Error Number: -2147467259
Error Description: The Portal cannot connect to the Report Server.
SQL Statement: SELECT PASSWORD FROM MC USER WHERE MC USER NAME = '<username>'
```

This message indicates that the reports database is not running. Verify the MSSQLServer service is running. If you are not sure, select the Start TeaLeaf services shortcut from the **Start** menu.

Portal giving 500 error on live sessions page

Note: This solution applies to Release 4.x or earlier versions of Tealeaf.

The `...\TeaLeaf\Portal\Web Application\temp chart` directory must be writable by the TeaLeaf Active Directory user groups or by the Portal machine's IUSR * user if you are not using NT authentication.

IIS error messages

This section describes some common error messages that are generated by IIS.

- For more information about IIS status codes, see `http://support.microsoft.com/kb/318380`.

Error message: 'The identity of application pool, 'DefaultAppPool' is invalid.

If you are using a named account, verify the following information:

- Set `User Cannot Change Password` to `true`. Setting the password to never expire does not seem to make a difference.
- User is a member of `IIS_WPG`:
  - `http://geekswithblogs.net/EltonStoneman/archive/2008/05/26/the-identity-of-application-pool-esbapppool-is-invalid.aspx`
  - `http://www.microsoft.com/technet/prodtechnol/WindowsServer2003/Library/IIS/f05a7c2b-36b0-4b6e-ac7c-662700081f25.mspx?mfr=true`

Error message: 'Failed to execute request because the App-Domain could not be created. Error: 0x80004005 Unspecified Error'

The AppDomain's identity does not have the correct permissions. Review the IIS permissions documentation for defaults.

Error Message: 'Failed to start monitoring directory changes' error message

There may be problems with the directory permissions for IIS/ASP.NET. For more information, see `http://support.microsoft.com/kb/317955`.

Error message: 'The current identity (NT AUTHORITY\NETWORK SERVICE) does not have write access to C:\WINDOWS\Microsoft.NET\Framework\v2.0.50727\Temporary ASP.NET Files'

Do the following:

1. Open a command shell (**Start** > **Run** > **cmd**) and navigate to one of the following directories with the {{50727)) varying depending on your OS build:
   - **For 64-bit O/S:**

     `C:\WINDOWS\Microsoft.NET\Framework64\v2.0.50727\`
2. Run the following command:

   `aspnet_regiis -ga "NT AUTHORITY\NETWORK SERVICE"`
3. Run the following command:

   `aspnet_regiis -ga IUSR_<machinename>`

Error Message: 'ASP.NET does not work with the default ASPNET account on a domain controller'

`aspnet_wp.exe` could not be started because the user name and/or password that is supplied in the processModel section of the config file are invalid.

- For more information, see `http://support.microsoft.com/kb/315158`.

Error message: 'CS0016: Could not write to output file c:\WINDOWS\Microsoft.NET\Framework64\v2.0.50727\Temporary ASP.NET Files\portal\d71063af\5e3121f8\App_global.asax.j8jlnnti.dll' - 'Access is denied. '

Verify that `TEMP system env. variable`:

- The folder exists.
- Network service and maybe the other ASP.NET accounts have full permissions to that folder

For more information, see `http://support.microsoft.com/kb/825791`.

Error Message: 401.2 'Logon failed due to server configuration'

No authentication methods were enabled on the Directory Security > Authentication and access control window

Common IIS 404 problems

See `http://support.microsoft.com/kb/248033/en-us`.

301 status code

Re-execute:

`aspnet_regiis -i`

No log entries in Windows/System32/LogFiles/W3SVC

IIS failed. Verify the required user account permissions. If you do not discover a permissions issue, reinstall IIS.

Service unavailable message

There might be an application pool permissions issue for the `Network Service` user. For more information, see `http://support.microsoft.com/?kbid=842493`.

Some NT users get a 400 - RequestRequest Length error message when they hit the Portal but others do not

If the user is a member of too many NT groups, IIS may have trouble acquiring the user's large request.

- If Integrated Authentication is enabled in IIS and it negotiates Kerberos, it is possible that some users have a large Kerberos ticket that hit request header limits.

To fix:

1. Add registry keys on the Reporting Server to increase the Maximum Field Length and Maximum Request Bytes to accommodate the larger Kerberos tickets:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\HTTP\Parameters\
   MaxFieldLength = 32768 (default = 16384)
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\HTTP\Parameters\
   MaxRequestBytes = 32768 (default = 16384)
```

2. Restart IIS.

- For more information about `Http.sys` registry settings for IIS, see `http://support.microsoft.com/?kbid=820129`.
- For more information about the Kerberos Authentication Problem with Active Directory, see `http://blogs.technet.com/surama/archive/2009/04/06/kerberos-authentication-problem-with-active-directory.aspx`.

Attempting to log in never leaves Default.aspx

It may be a problem with a cookie.

- Run `SimpleTest.aspx` to verify that cookies are working.
- If cookies are working, verify that the date on the client machine is correct.
  - If the date is too far off, it might think the Portal's cookie already expired, which results in the user remaining on the Login page.

Sessions die every 29 hours (1740 minutes or 1 Day, 5 Hours)

Disable the `Recycle worker processing (in minutes)` setting in the Portal's Application Pool's Recycling settings tab.

Error message: "The NT Login/DB Login and New User Dialogs all show on the Default.aspx page"

ASP.NET is not functioning properly.

Re-execute the following command:

```
aspnet_regiis -i
```

NT Credentials dialog pops up even in Internet Explorer

Verify the following:

- The user's machine and the Portal machine must both be in the same domain.
- In Internet Explorer, select **Tools** > **Internet Options** > **Advanced** > **Security**. Select the Enable Integrated Windows Authentication check box

- In Internet Explorer, select **Tools** > **Internet Options** > **Security zones**. Verify that either:
  - **User Authentication** > **Logon** is configured properly
  - The Portal is in the correct trusted zone.
- If `web.config debug=true` does not work, verify that the machine.config deployment retail=true.
- If right after installation, you receive a 500 compilation error, set `customErrors=Off` in the `Web.config` file in `/Portal/WebApp`. This setting disables the custom errors page and helps to narrow in on the problem. If the problem is in the `Web.config` file, you may be running the Portal against ASP.NET 1.1 and not 2.0:
  1. Verify .NET 2.0 is installed.
  2. Verify ASP.NET 2.0 was registered (Execute: `aspnet_regiis.exe /i`).
  3. Verify ASP.NET 2.0 is selected under the ASP.NET tab in the Portal virtual directory properties.
     - If no ASP.NET tab exists, you may need to manually register that virtual directory with `aspnet_regiis`.

Error starting BBRFilter.dll

Microsoft Visual C++ 2005 Redistributable Package (x86) is missing or corrupted. For more information, see `http://www.microsoft.com/downloads/details.aspx?familyid=32bc1bee-a3f9-4c13-9c99-220b62a191ee&displaylang=en`.

Locale error

To the locale specification for the Portal:

```
<system.web>
        <globalization fileEncoding="utf-8" requestEncoding="utf-8"
             responseEncoding="utf-8" culture="en-US" uiCulture="en-US" />
</system.web>
```

Error creating the Web Proxy specified in the 'system.net/defaultProxy' configuration section.

```
System.Configuration.ConfigurationErrorsException: Error creating the Web
Proxy specified in the 'system.net/defaultProxy' configuration section. --->
System.DllNotFoundException: Unable to load DLL 'rasapi32.dll':
A dynamic link
library (DLL) initialization routine failed. (Exception from HRESULT:
0x8007045A) atSystem.Net.UnsafeNclNativeMethods.RasHelper.
RasEnumConnections(RASCONN[]
lprasconn, UInt32& lpcb, UInt32& lpcConnections) atSystem.Net.
UnsafeNclNativeMethods.RasHelper.GetCurrentConnectoid() atSystem.
Net.AutoWebProxyScriptEngine.AutoDetector.Initialize() atSystem.
Net.AutoWebProxyScriptEngine.AutoDetector.get_CurrentAutoDetector() at
System.Net.AutoWebProxyScriptEngine..ctor(WebProxy proxy,
Boolean useRegistry)
at System.Net.WebProxy.UnsafeUpdateFromRegistry() at
System.Net.WebProxy..ctor(Boolean enableAutoproxy) at
System.Net.Configuration.DefaultProxySectionInternal..ctor
(DefaultProxySection
section) at System.Net.Configuration.DefaultProxySectionInternal.GetSection()
--- End of inner exception stack trace --- atSystem.
Net.Configuration.DefaultProxySectionInternal.GetSection() atSystem.
Net.WebRequest.get_InternalDefaultWebProxy() at
System.Net.HttpWebRequest..ctor(Uri uri, ServicePoint servicePoint) at
System.Net.HttpRequestCreator.Create(Uri Uri) at
System.Net.WebRequest.Create(Uri requestUri,
```

```
Boolean useUriBase) atSystem.Net.WebRequest.Create
(String requestUriString) at
TeaLeaf.SearchServer.RealSearchServer.RealCommunicate(String command,
ArrayList args, Int32 timeout)
```

Web.Config needs the following:

```
<system.net>
        <defaultProxy>
            <proxy usesystemdefault="false"/>
        </defaultProxy>
    </system.net>
```

Using IIS7 and seeing 401.2

If you are seeing 401.2 errors, verify that Anonymous and Forms authentications are enabled. Verify account for anonymous is IUSR or some other account that has the proper permissions.

Other Useful Microsoft Knowledge Base Articles

Default IIS/Account settings for IIS 6.0

http://support.microsoft.com/kb/812614/

How to set required NTFS permissions and user rights for an IIS 5.0 Web server

http://support.microsoft.com/kb/271071

"W3wp.exe could not be started" error message in the application event log

http://support.microsoft.com/default.aspx?scid=kb;en-us;833444

"Aspnet_wp.exe could not be started" error message

http://support.microsoft.com/default.aspx?kbid=811320

"403.1 (Execute) Access is Denied"

ASP.NET script mappings are broken

"This implementation is not part of the Windows Platform FIPS validated cryptographic algorithms."

http://blogs.msdn.com/shawnfa/archive/2005/05/16/417975.aspx

"Failed to map the path '/App_GlobalResources'."

http://forums.microsoft.com/MSDN/ShowPost.aspx?PostID=117013&SiteID=1

"403.18 - The specified request cannot be executed from current Application Pool"

http://blogs.msdn.com/rakkimk/archive/2006/09/01/735684.aspx
To avoid this problem, you have to create a registry key with the name "IgnoreAppPoolForCustomErrors", type "DWORD" and value 1 under the "HKLM\SYSTEM\CurrentControlSet\Services\W3SVC\Parameters\" folder.

"Error message when ASP.NET 2.0 is configured to run with a user account: "Unable to generate a temporary class"

http://support.microsoft.com/kb/908158
More detailed:
[InvalidOperationException: Unable to generate a temporary class (result=1).
error CS2001: Source file 'D:\WINDOWS\TEMP\d0lurtzx.0.cs' could not be found
error CS2008: No inputs specified

Default Web Site was removed

http://www.ssw.com.au/ssw/KB/KB.aspx?KBID=Q889546

Unable to perform searches of completed sessions

If you are performing a search of completed sessions from the Portal, you may return an error indicating that there are no indexes to search. You can perform the following steps to troubleshoot the issue.

1. Try searching through RTV. If you are able to search using RTV, then the Search Server is operational.
2. You may want to examine the state of individual indexes in the Search Server. To do so, you must configure RTV to use the same Search Server. For more information about connecting to a Search Server, see "RTV Search Setup" in the *IBM Tealeaf RealiTea Viewer User Manual*.
   a. If you are able to successfully connect to the same Search Server, the list of available indexes is displayed beneath the server node in the Search Setup window.
   b. To review a summary of the indexes for the Search Server, select the server node and click **Summary...**.
   c. In the Summary page that is generated in your browser, you can review the status of each index that is managed by the search server. Review the From and To timestamps and the values for Index, Size, and Session Indexes. If these fields have meaningful values, you should be able to search for session data between the From and To dates.
3. If Search Server appears to be operating properly through RTV, the issue may be with the Data Service. The Portal acquires its set of available indexes from the Tealeaf Data Service. Complete the following steps to verify that the Data Service is operational.
   a. In the Portal, select **Tealeaf** > **Portal Management**.
   b. Click the **Manage Servers** link.
   c. Select the **Data Service server**.
      - If the Data Service server does not exist in the list of servers, click the **Display Inactive** check box. If the server now appears, you must activate it.
      - If the Data Service server still fails to appear, you must create an entry for it.
      - See "Managing Tealeaf Servers" in the *IBM Tealeaf cxImpact Administration Manual*.
   d. When the Data Service server is selected, in the toolbar above the list of servers, click the **Ping** tool in the toolbar. Version and request information should be displayed in the **Data Service Server Status** panel below the list of servers.

4. If the Data Service server does not respond, try restarting the Data Service.

   Note: While the Data Service is restarting, the Portal is unable to connect to other Search Servers and the Tealeaf databases. Typically, restart takes only a few seconds.

   a. In the Portal, select **Tealeaf** > **TMS**.
   b. In the Tealeaf Management System, select the **WorldView** tab.
   c. From the View drop-down, select `Servers`.
   d. Click the **Data Service** node.
   e. In the Server Actions, click **Restart**.
5. If the restart does not resolve the issue, the problem may be related to any of the following:
   a. Inadequate permissions
   b. Special configuration for the Portal in IIS or ASP.NET required to connect to Search Server must also be applied to the Data Service.
   c. Use of raw IP addresses in the server addresses of the Portal Management page. Use server names instead.
   d. Proxy server between the client browser and the Portal. Disable the proxy if possible or configure it to manually connect to the Search Server.

Status Code 400 errors

If completed searches are failing with Status Code 400 errors, you might have a poorly specified path value in Search Server configuration.

The temp path must be a fully qualified path. If the path is not fully qualified, the path is treated as a relative path, and the resulting behavior depends on the current drive directory:

- Search server might refuse to start, and an event log error might indicate a problem with the temp path.
- Search server might return Status Code 400 errors on commands that use the temp drive, such as searching or retrieving the list of indexes.

The fix is to use full path names for values in Search Server configuration.

Note: In Search Server configuration, use fully qualified path values for all paths.

- See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

Diagnosing search performance issues

If you are experiencing long search times, search timeouts, or increases in search times, you should try the following tests to see if you can accelerate searches:

1. Obtain a disk I/O benchmarking utility.
   - Microsoft SQLIO is a useful one. For more information, search `www.microsoft.com` for "SQLIO Disk Subsystem Benchmark Tool".
2. Stop all Tealeaf services.
3. Run SQLIO against the Tealeaf disk(s).
4. You should observe transfer rates of higher than 7 megabytes per second. If you are experiencing slower transfer rates, there are issues with your storage device(s). Consult the appropriate administrator for details and further diagnostics.

5. Another area to check is the size of the indexes that are compared to the LSSN files in the Long-Term Canister. In the Tealeaf Portal, select **Tealeaf** > **System Status** > **Storage**. Select an individual server from the IBM Tealeaf CX drop-down.
   a. If the index sizes are larger than the size of the LTC files, then Tealeaf may be indexing unnecessary data.
   b. In several captured hits, you should examine the contents of the request and the response to ensure that all data that is being indexed is needed. Remove from indexing any unneeded data.
      - You can create privacy rules to strip out unwanted content so it is not indexed. For more information, see the Rules tab in "Managing Data Privacy in Tealeaf CX" in the *IBM Tealeaf CX Installation Manual*.

What does "All text" search for?

An "all text" search examines the request and the response for every term that was indexed. The all text search of the response includes every word that is considered "meaningful".

- To reduce disk space requirements for indexes, noise words, single letters, HTML tags, and JavaScript code are excluded from indexing.

1. To see approximately what is indexed and thus searchable in a response, open the IBM Tealeaf CX RealiTea Viewer and examine the response view of a typical page.
   a. Right-click the response view
   b. Select the Indexed view or use the drop-down menu on the right side of the **RSP** button on the toolbar.
2. This view indicates the words for which you can search in the response to find a specific page.

In a request, specific items are included by default. See "Configuring CX Indexing" in the *IBM Tealeaf CX Configuration Manual*.

Can't include the word NOT in a search

Tealeaf uses two different search engines:

- **dtSearch Indexer** is the more powerful one and is invoked when searching completed sessions through RTV or the Portal.
- **string comparer** is a simple string comparison search facility, which is invoked when you search for active sessions.

In the dtSearch Indexing engine, the word `NOT` is a reserved operator for Boolean expressions used to construct complex search terms. That word is not indexed, which means you cannot be searched for it.

- Likewise, hyphens and other punctuation characters are not indexed, as they are considered word separators.

Example: You want to search for the URL `slp-system-not-available.html`. Entering this string results in few if any matches. From Portal Search, search instead for `"slp system available"`, including the double quotation marks.

- The word `html` was left off at the end of the search to widen the net on the first search. If you find the three-word search is returning incorrect hits, then you can add the fourth word `html` to the phrase to narrow the results set.

For more information about search syntax:

- "RealiTea Viewer - Search Syntax" in the *IBM Tealeaf RealiTea Viewer User Manual*
- "Character Indexing" in the *IBM Tealeaf cxImpact Administration Manual*

Cannot search on fields with ~ or other punctuation characters in the field names

The Session Indexer service breaks words apart on special characters as defined in `alphabet.dat`. However, `name=value` fields would be broken up if they contained any names with special characters, making it impossible to search for the name. To avoid this situation, the Indexer replaces any special punctuation characters in the name (left of the = sign) with underscores (_). For example, to search for `~cc_num`, enter the field name as `_cc_num`.

- See "Character Indexing" in the *IBM Tealeaf cxImpact Administration Manual*.

Why do session fragments affect my search results?

Sessions can become fragmented under multiple conditions. Fragmentation affects searching and some event triggers.

For example, suppose the web pages for a business process consist of the `add an item`, `start checkout`, and `complete checkout` events. Assume that the website allows for a long application timeout and that the Tealeaf STC session timeout is configured to be less than the application timeout.

- This discrepancy is often necessary, as the Short-Term Canister usage directly affects RAM consumption in the Tealeaf server.

A visitor to the site could add an item, start checkout, go have dinner, and return to finish the checkout an hour later. This behavior results in one session fragment containing the first two events, and a second, later session fragment containing the final `complete checkout` event.

Tealeaf search looks for session fragments. A search for `events 1 and 2 and not event 3` would find session fragment 1 and not fragment 2. When replaying session fragment 1, the user can select `Find all fragments of this session`, which replays a single logical session consisting of the contents of both session fragments. The merged session has all three events, which are not what the original search requested.

For most sites, it is difficult to completely avoid fragmentation. To reduce fragmentation, you can increase the duration of the STC timeout, which may impact STC RAM consumption. Over-consumption of the STC RAM leads to spooling, which may cause fragmentation. Spooling should be avoided wherever possible. You can experiment with gradually increasing the STC timeout until fragmentation is reduced to a tolerable level without increasing spooling.

Another approach is to find a representative sample of the fragmented sessions, analyze them for the time duration that contains all events of interest, and set the STC timeout such that 95% or other desired percentage of the user sessions are not fragmented for the events of interest.

- For more information about configuring the session timeout, see "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

- For more information about managing fragmentation through RTV, see "RealiTea Viewer - Annotation, Find and Merge" in the *IBM Tealeaf RealiTea Viewer User Manual*.
- For more information about configuring how session fragments are merged for BBR, see "Managing Tealeaf Servers" in the *IBM Tealeaf cxImpact Administration Manual*.
- For more information about managing fragmentation through BBR, see "BBR Options" in the *IBM Tealeaf cxImpact User Manual*.

Differences in session counts as reported by Portal and Viewer (or Data Extractor / cxConnect)

Sessions are stored in the Long-Term Canister according to their last hit timestamp, as the archive must contain session-level compound events that cannot be evaluated until session timeout. These events and the session's Session-Time attribute are applied the timestamp of the last hit of the session.

The Viewer's archive search and Data Extractor / IBM Tealeaf cxConnect for Data Analysis searches both use the indexes and the date-time filter matches against the ending timestamp of the session. There may be small differences between the number of sessions in a day as reported by the Portal's Total Activity report and the number of sessions in a day as returned by the Viewer or the data extraction tools.

Configuration

For more information about configuration:

- "Configuring Search Templates" in the *IBM Tealeaf cxImpact Administration Manual*
- "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*
- "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*

How to tell if the disks are I/O limited

Suboptimal disk layout and allocation can lead to poor performance of the Tealeaf components. At http://support.tealeaf.com/, Solution 65 "Configuring the RAID Hard Disk Drive (HDD) Array" goes into detail about the best way to configure the physical disks into arrays and explains the reasons behind the recommendations. If your server is not laid out according to these practices, the Tealeaf system performance may be disk I/O-bound.

This solution explains how to determine if your disks are creating a performance bottleneck:

- Run the Windows PerfMon utility.
- If the `Average Disk Queue Length` counter is not already displayed, right-click the right pane and select **Add counters**.
- In the Performance Object drop-down, select **Physical Disk**.
- In the rightmost set of radio buttons, select **All instances**.
- In the left pair of radio buttons, select **Select counters from list**.
- Multi-select the following three counters from the list box:
  - `Avg Disk Queue Length`
  - `Avg Disk Read Queue Length`
  - `Avg Disk Write Queue Length`

- Click **Add**.
- Look at the results in the perfmon right pane.
- Focus on the disks that house the CANISTER, dbs directory, and the Indexes directory.
  - Average queue lengths should be less than 1.
  - Queue lengths of 4 or greater mean the disk is I/O bound.

Average queue lengths of 4 or more are a strong indication that the system's disk drives need repartitioning along the guidelines of Solution 65.

If your system is configured according to Solution 65 standards, does not use RAID, and has high average queue lengths, the solution is to move to a faster disk subsystem or add additional Tealeaf servers. The current system does not have the disk I/O bandwidth to handle the data processing load that is presented to it.

Repeated restarts after dtSearch hang

When running searches, you may receive the following entries in the application event log:

```
(18:49 Search Server) - TeaLeaf Search Server Ver: 6.2.0.6250 -
   RestartOnDtSearchHang
(18:01 Search Server) - TeaLeaf Search Server Ver: 6.2.0.6250 -
   RestartOnDtSearchHang
(16:21 Search Server) - TeaLeaf Search Server Ver: 6.2.0.6250 -
   RestartOnDtSearchHang
```

Search Server is unable to complete the requested search due to disk performance issues or insufficient time to complete the search.

Improving disk performance

This problem may be the result of performance issues with the disk hardware. Perform the following checks:

- Run the Windows PerfMon utility as described in "How to tell if the disks are I/O limited" on page 45.
  1. Check the average disk queue length of the disk where the Tealeaf Indexes directory resides. If this value is typically larger than 1, the disk cannot keep up with I/O requests.
  2. Close `PerfMon.exe`.
- You can use Microsoft's SQLIO benchmark utility to assess the disk speed. This utility is available here: `http://www.microsoft.com/downloads/details.aspx?familyid=9a8b005b-84e4-4f24-8d65-cb53442d9e19&displaylang=en`
- If the disk's throughput is less than 50 megabytes/sec, it is not fast enough for the processing volume that is required by the Tealeaf server.
- Consult with your IT department for ways to improve the disk I/O performance.

Edit Search Server Watchdog Timeout setting

If the disk hardware performance cannot be improved, change the Search Server Watchdog Timeout setting:

- **Release 7.2 or later:**

1. In the Portal, select **Tealeaf** > **TMS**. The Tealeaf Management System is displayed.
2. From the View drop-down, select `Servers`.
3. Click the Search Server node.
4. Click **Search Server configuration**.
5. Click **View/Edit (raw)**.
6. In the window, click the Default group.
7. Change the value for `WatchDogIndexSearchSeconds` to double or quadruple the current value.
8. Click **Save**.
9. Assign and push the configuration. See "TMS WorldView Tab" in the *IBM Tealeaf cxImpact Administration Manual*.

- **Release 7.1 or earlier:** Edit the following Windows Registry value on the Tealeaf server where the error message appears in the Windows Application Event Log.

  Note: The recommended value for this setting is `300` seconds.

  – For 64-bit Windows:

  ```
  HKEY_LOCAL_MACHINE\Software\(Wow6432Node)\TealeafTechnology\DataStore\
     SearchServer\WatchDogIndexSearchSeconds
  ```

- Double or quadruple the current value for the above key.
  – If necessary, you can set the value to `0` to disable this feature entirely.
- To apply the new value, restart Search Server.

Canister spacer file does not exist

The Windows Application Event Log may contain a warning message from the TeaLeaf Pipeline:

```
SADecoupleEx (TeaLeafCSS_1966): Canister spacer file does not exist
(C:\Program Files (x86)\TeaLeaf\Canister\Canister.spacer).
```

The Canister space file is used to reserve disk space for storage of sessions in the Canister. The above message indicates that this file does not exist. Without it, the Canister cannot save the processed sessions when disk space is low.

The Extended Decoupler session agent normally does not interact with this spacer file. However, when a low disk space condition is detected, the session agent deletes it and frees up disk space. When the message appears, DecoupleEx was unable to find the file.

- For more information about the session agent, see "Extended Decoupler Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

The spacer file is created when the TLTMaint utility is run. To correct this issue, you can run this utility to create the spacer file.

- See "Maintaining the CX System" in the *IBM Tealeaf CX Installation Manual*.

DecoupleEx warning occurs when starting services

The Windows Application Event Log may contain a warning message from the TeaLeaf Pipeline:

```
SADecoupleEx: Canister shared memory is invalid. Flow of hits to
the canister has been stopped.
```

If the Application Event Log contains many other TeaLeaf messages that are clustered around this same time, you may see that all Tealeaf services are shutting down and restarting. In most cases, the "TLBackup -CycleServices" Windows Scheduled Task is run during the early morning hours, typically at 2:05am. Part of the start of CycleServices is running the TLTMaint program, which is executed whenever the TeaLeaf Canister Server service is started. TLTMaint examines the entire Long-Term Canister for data consistency and can run from 2 seconds to 200 seconds or more. In most cases, it completes in less than a minute. While TLTMaint is running, no other TeaLeaf services can connect to the LTC.

The warning message you see from SADecoupleEx occurs when the Tealeaf Pipeline is attempting to connect to the Canister and is being informed that the Canister is not available because TLTMaint is still running.

A little later in the TeaLeaf startup messages in the Event Log, you should see that TLTMaint completed with no errors and another message from the TeaLeaf Pipeline, indicating that the flow of hits to the Canister resumed.

After Tealeaf is in place for a few weeks, the TLTMaint program takes longer to complete because the LTC is bigger. Therefore, the warning message might not appear until a few weeks after the initial installation yet becomes a standard entry in the Event Log thereafter.

If you see this warning in the Event Log and it is not part of a startup sequence, log a support case at http://support.tealeaf.com/.

TeaLeaf Technology Error Data length (xxxxxx) exceeds maximum size (yyyyyyy) body discarded

You may receive the following error:

```
TeaLeaf Technology Error: Data length (xxxxxx) exceeds maximum size (yyyyyy).
Body discarded.
```

Make the following changes to capture the body:

1. **IBM Tealeaf CX Passive Capture Application (PCA server) configuration:** Only apply the following if you have a IBM Tealeaf CX Passive Capture Application server: Adjust `Max response size` on the Pipeline tab of the PCA Web Console.
   - See "PCA Web Console - Pipeline Tab" in the *IBM Tealeaf Passive Capture Application Manual*.
2. **RealiTeaIBM Tealeaf CX Server configuration:** Adjust **MaxDataSize** in the `[Globals]` section of `TealeafCaptureSocket.cfg`. Restart the Transport Service afterward.
   - See "Configuring the Transport Service" in the *IBM Tealeaf CX Configuration Manual*.

Note: Completing these changes negatively impact performance and disk utilization by capturing these large hits.

FetchNSAT failed in WriteInteraction

FetchNSAT failed in WriteInteraction with code -160

If this error occurs repeatedly and frequently, you can increase the session timeout by a minute.

- Session timeouts can also be changed by triggered events. See "TEM Events Tab" in the *IBM Tealeaf Event Manager Manual*.

The most likely cause of this error is when a session times out just as the next hit in the session is being added. The result is that the session is fragmented, and the hit is placed in a newly created session.

FetchNSAT failed in WriteInteraction with code -127

Typically, the system is completing a cycle services or a CanTrim operation currently.

- If the Tealeaf Canister Server service fails while a hit is in the pipeline (a rare occurrence), the hit could be dropped with this message.
- If this message occurs only when the Canister service is being restarted, it can be ignored.

Configuration

See "Configuring the Transport Service" in the *IBM Tealeaf CX Configuration Manual*. For more information about cycle services, see "Maintaining the CX System" in the *IBM Tealeaf CX Installation Manual*.

Chapter 2. Troubleshooting Tealeaf databases

If you are experiencing an access or performance-related issue with your installed Tealeaf SQL database, review the following troubleshooting tips, which may help to resolve your issue.

Note: For some functions of the Tealeaf Database Manager, you must have System Administrator privileges for the SQL Server hosting the Tealeaf databases.

The Tealeaf SQL databases are installed using the Tealeaf Database Manager, which can be used to upgrade or reinstall the databases. Some of these steps may require reinstallation or upgrading. See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

Note: The Tealeaf Database Manager is used to install the SQL Reporting databases and IBM Tealeaf cxResults databases. It cannot be used to install or update the Canister databases, which are installed as part of the Processing Server through `Setup.exe` in the initial installation. For more information about those databases, see "CX Installation and Setup" in the *IBM Tealeaf CX Installation Manual*.

- For more information about configuring the Canister databases, see "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

SQL Server configuration

Unable to connect to SQL Server

If Tealeaf is unable to connect to the SQL Server database, try the following steps:

1. **SQL Server custom ports**
   - By default, SQL Server listens on port 1433, which is also the default communication port for Tealeaf, unless selected differently during installation. You can verify this value through the SQL Server Configuration Manager.
2. **Unable to connect through dynamic ports**
   - If your SQL Server implementation uses dynamic ports, you must sync your Tealeaf solution with them. See “Configuring use of SQL Server dynamic ports” on page 52.
3. **SQL network protocols**
   - By default, SQL Server does not enable TCP and Named Pipes, which are used by Tealeaf. Verify that these protocols are enabled through the SQL Configuration Manager.

Failure to install due to model database size mismatch

During installation of the Tealeaf databases, the Tealeaf Database Manager may report a failure because of a mismatch between the configured database sizes in Tealeaf and the model database sizes that are configured in SQL Server Management Studio.

In rare cases, a customer may change the size of the model database size in their SQL Server installation from the default value of 50MB. SQL Server Management Studio does not allow databases to be installed below the model database size. For

example, if the model database size is set to 250MB, when the Tealeaf Database Manager attempts to install the Tealeaf databases, some installations fail because they are smaller than this minimum limit.

The solution is to do one of the following:

1. Reconfigure the model database size in SQL Server Management Studio to 50MB. For more information, consult the documentation that was provided with your SQL Server product.
2. Reconfigure the database sizes in Tealeaf Database Manager to be slightly larger than the model database size in SQL Server Management Studio. For the above example, you could set the database size to 260MB.
   - Database sizes must be configured through individual filegroups for each database. See "Tealeaf Database Manager Reference" in the *IBM Tealeaf Databases Guide*.

Configuring use of SQL Server dynamic ports

Tealeaf does not natively support use of dynamic ports when accessing SQL Server. If possible, Tealeaf recommends disabling use of dynamic ports.

When SQL Server starts, a dynamic port is selected. This port is used during operations. During restart, the configured port remains, unless a conflict emerges.

If the ports must remain dynamic, complete the following steps.

1. Start all Tealeaf services.
2. Start the database.
3. Through SQL Server Management Studio, locate the port that SQL Server is using.
4. This port must be populated for each database in the Connection dialog of the Tealeaf Database Manager. See "Tealeaf Database Manager Reference" in the *IBM Tealeaf Databases Guide*.
5. In the Tealeaf Database Manager, you can set the ports for the databases.
   a. From the TDM menu, select **Mode** > **Info/Config** > **Report Server Configuration**.
   b. Populate the ports accordingly. When changes are applied, the registry keys below are updated for you.
      - See "Tealeaf Database Manager Reference" in the *IBM Tealeaf Databases Guide*.
6. Through Tealeaf Database Manager, install or upgrade the databases as necessary.
   - See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

On the Report Server, database port numbers are in the following registry keys. You can update these through the Tealeaf Database Manager Connecting screen.

Product (Database)
Registry Key

IBM Tealeaf cxImpact (Reports)
`Tealeaf Technology > DataStore > Report Server > Port`

IBM Tealeaf cxResults (Visitors)
`Tealeaf Technology > DataStore > Report Server > Visitor Port`

IBM Tealeaf cxReveal (Search)

```
Tealeaf Technology > DataStore > Database > Search > Database Port
```

Note: If SQL Server is forced to choose a new dynamic port, then the new port must be used whenever you access the databases through the Connecting screen in the Tealeaf Database Manager and repopulated in the appropriate registry key through the Tealeaf Database Manager using the above steps.

Migrating to a new SQL Server

If you are migrating to a new instance of SQL Server, there are more steps that are required.

See "Migrating Tealeaf Databases" in the *IBM Tealeaf Databases Guide*.

Restore failed when MDF file name is claimed

When restoring databases to a new instance of SQL Server, you may encounter an error similar to the following:

```
Restore failed for Server <MyServer>

Additional information:System.Data.SqlClient.SqlError: File:
'C:\Program Files\Microsoft SQL Server\
MSSQL10_50.MSSQLSERVER\MSSQL\DATA\RL_REPORTS.mdf' is claimed by 'DATA'(3) and
'PRIMARY'(1). The WITH MOVE clause can be used to relocate one or more files.
(Microsoft.SqlServer.Smo)
```

The above issue is caused by the SQL Server restore procedure that converts all of the Tealeaf database names to `RL_REPORT.mdf`. This situation is typically caused when the new location has a different directory structure or when the databases are moved to a new version of SQL Server.

The solution is to manually append the file name of each data file to include the Filegroup identifier for the file to the end of the file name.

For more information about the appropriate steps to restore the databases, see "Migrating Tealeaf Databases" in the *IBM Tealeaf Databases Guide*.

Poor SQL Server performance on 64-Bit operating systems

Note: If you are using Tealeaf version 8.8 or later, SQL Server 2005 is no longer supported. For more information, see "Supported SQL Server Versions" in the *IBM Tealeaf CX Databases Guide*.

Tealeaf databases are designed to take advantage of 64-bit operating system capabilities. However, because of a Microsoft issue, performance issues may arrive when running SQL Server 2005 on a 64-bit version of Windows. These issues can include the following:

- The performance of SQL Server 2005 decreases suddenly.
- SQL Server 2005 stops responding for a short time.
- A timeout occurs for applications that connect to SQL Server 2005.
- Problems occur when you run even simple commands or use applications on the system.

The issue may be that 64-bit Windows operating system is paging out the working set of the SQL Server process. When the SQL Server process reaches 50 percent of

the memory that is allocated for the process on SQL Server 2005 SP2 or later, error messages may begin appearing in the SQL Server error log.

To fix the issue:

1. Review the Microsoft troubleshooting steps. For more information, review "How to troubleshoot this problem" on `http://support.microsoft.com/kb/918483`.
2. If the problem persists, you can prevent Windows from paging out the buffer pool memory of the SQL Server process. You can lock the memory by assigning the `Lock pages in memory` user right to the user account that is the startup account for the SQL Server service.
   - This step prevents the operating system from paging out memory and instead allows SQL Server to do the paging. For more information, see `http://technet.microsoft.com/en-us/library/ms190730.aspx`.
3. Restart the SQL Server.

For more information about this topic including the troubleshooting steps, visit `http://support.microsoft.com/kb/918483`.

SuperSocket information: (SpnRegister): error 1355 in Windows application event log

From http://support.microsoft.com/kb/303411:

> This message is not an error message. This text is only a warning that SQL Server was not able to register a Service Principal Name (SPN), which indicates that the security mechanism used is Microsoft Windows NT Challenge\Response (NTLM) authentication instead of Kerberos.
> CAUSE
> The message usually appears because the SQL Server service account is running as a domain user who does not have requisite permissions to register SPNs.

SQL Server blocks access to procedure sys.sp_OAGetProperty

In some environments, the following error may be displayed in the SQL Server error log:

```
SQL Server blocked access to procedure 'sys.sp_OAGetProperty' of component
'Ole Automation Procedures' because this component is turned off as part of
the security configuration for this server. A system administrator can
enable the use of 'Ole Automation Procedures' by using sp_configure.
For more information about enabling 'Ole Automation Procedures',
see "Surface Area Configuration" in SQL Server Books Online.
```

`sys.sp_OAGetProperty` is an automation procedure that allows access to system details. Tealeaf uses this stored procedure to acquire details on disk space.

Note: If access to the `sys.sp_OAGetProperty` procedure is not enabled:

1. The Database Filegroup Size and Database Table Size reports in the Portal do not contain free space and unused space information.
2. Available disk space on SQL Server is not recorded and does not appear in the Tealeaf Event Log.

The error message is harmless. However, to avoid confusion, you may want to deny access to pr_ServerDiskSpace, which uses this stored procedure. When access is denied, error messages are reported from pr_ServerDiskSpace, instead of sys.sp_OAGetProperty. Error messages from the latter procedure may suggest possible hacks or malicious software, so unnecessary messages should be suppressed.

To deny access, run the following through SQL Server Management Studio:

```
use TL_STATISTICS
go
Deny execute on pr_ServerDiskSpace to TLAdmin;
Deny execute on pr_ServerDiskSpace to TLUser;

use TL_visreport
go
Deny execute on pr_ServerDiskSpace to TLAdmin;
Deny execute on pr_ServerDiskSpace to TLUser
```

Tealeaf database configuration

Unable to connect

If you are unable to access the Tealeaf database, you should verify that you are using the appropriate full-qualified host name.

If you are running against a named database instance, then the host name must follow a specific format:
<host_name>\<instance_name>
where
<host_name> is the IP address, the machine name, or localhost
<instance_name> is the name of the DB instance.

See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

During upgrade, Tealeaf Database Manager fails to upgrade the Result Set Extractor

The required views do not exist. The RSE Database must be re-created. See “Result Set Extractor database fails to analyze a result set” on page 57.

Failing to create database in clustered server environment

During installation into a clustered server environment, if you are seeing error messages during creation, you may have improperly specified your SQL path to a disk that is not part of the cluster group or dependency list.

To resolve the cluster dependency issue, see http://support.microsoft.com/kb/295732.

In the Tealeaf Database Manager, you should also verify the path where the database files are stored. If you do not know the path to a remote server, contact the appropriate person.

Note: The path must be the absolute path relative to the SQL Server. Network paths are not permitted.

See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

Tealeaf database performance

Memory allocation

If any of the following conditions are occurring:

1. Tealeaf is performing slowly
2. Reports are delayed in arriving
3. The box running SQL Server is having paging or disk I/O problems

The memory that is allocated for SQL Server may be set too low. You can verify and change the amount of memory that is allocated to the Tealeaf databases in the SQL Server Management Studio.

Fragmentation

Normal database operations may cause the hard disk drive media and the stored Tealeaf databases to become fragmented. Over time, disk fragmentation can significantly affect performance.

- As part of regular server maintenance, your DBA should perform disk fragmentation checks and defragment disks periodically.

The Tealeaf database indexes can also become fragmented. Periodically, these should be defragmented to ensure optimal performance.

Note: Generally, database index fragmentation should be kept below 20%.

- **Release 7.1 or earlier:** You can run fragmentation checks and defragmentation tasks through Microsoft SQL Server Administration Console. For more information, consult the documentation that came with the product.
- **Release 7.2:** Through the Tealeaf Database Manager, you can check the fragmentation status of your database indexes and run defragmentation tasks as needed. See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

Tealeaf Database Manager hangs when upgrading the System database

The Tealeaf Database Manager may hang when the System database is being upgrade if other components of the system are running.

While this situation may be caused by multiple factors, in many cases the cause is the presence of an active Stats Logger session agent in the Windows pipeline of the Report Server. This session agent is used to collect statistical information in the pipeline and submit it to the Statistics database for use in Portal-based reports.

- See "Statistics Logger Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

The active session agent may be keeping the Statistics database open and therefore unavailable to the Tealeaf Database Manager, which causes it to hang.

Possible solutions

To resolve this issue, you may pursue any of the following solutions:

1. Stop the Transport Service before you start the Database Manager.

   Note: As soon as you stop the Transport Service, hits are dropped and data is lost.

   a. After you use the Database Manager, you must remember to restart the Transport Service through TMS.

- See "Configuring the Transport Service" in the *IBM Tealeaf CX Configuration Manual*.
- See "TMS WorldView Tab" in the *IBM Tealeaf cxImpact Administration Manual*.

2. Remove or disable the Stats Logger session agent from your processing pipeline.

   Note: Tealeaf recommends disabling the Stats Logger session agent during Database Manager operations in an All-in-One system.

   a. When the session agent is removed from the pipeline, you may start the Database Manager.
   b. Remember to add it back through TMS after you exited the Database Manager.

   - See "TMS Pipeline Editor" in the *IBM Tealeaf cxImpact Administration Manual*.

3. If the Stats Logger session agent is part of a dedicated child pipeline in the Canister, you can disable the whole pipeline through TMS.
   - See "TMS Pipeline Editor" in the *IBM Tealeaf cxImpact Administration Manual*.

Issues with individual databases

Visitor database for cxResults

Note: The Visitor database applies to IBM Tealeaf cxResults. IBM Tealeaf cxResults is no longer available as a newly licensed product as of Release 8.7. Customers that licensed IBM Tealeaf cxResults in Release 8.6 and earlier may continue to use and receive support for the product in Release 8.7 and later. For more information, contact Tealeaf Customer Support .

Data connection timeout in Visitor database

Although there is a data connection timeout setting for the Visitor database, the default value is set high. It should not need to be changed.

Result Set Extractor Database for cxResults

Result Set Extractor database fails to analyze a result set

During an upgrade or reinstall of the RSE database, you may encounter the following error message, which is followed by a table create statement:

```
Failed to create the database views:
Invalid object name 'RSE_HITSTATISTICS'.
```

In this case, the RSE database may be missing some views. In this case, the RSE database empty. To resolve this issue, the database must be re-created:

1. Run the Tealeaf Database Manager. See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.
2. Select **Advanced mode**.
3. In the Database Install tab, click **Create**.
4. In the Database Type pane, click **RSE**.
5. Click **Run**. The RSE database is re-created.
   - Review and observe any generated error messages.
6. After the reinstallation is complete, you may upgrade the database. See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

Search database for cxReveal

cxReveal database install forces collation setting

If you attempted to install the IBM Tealeaf cxReveal search database, the database collation setting may be forced to the following:

```
SQL_Latin1_General_CP1_CI_AS
```

This setting is configured regardless of the collation setting in SQL Server, which presents a problem for international customers, specifically those in Europe.

The solution to drop and reinstall the TL_SEARCH database through the Tealeaf Database Manager.

Note: Before you begin, verify that the System database (TL_SYSTEM) and the Reports database (TL_REPORTS) were installed or upgraded to the version for which you are installing the TL_SEARCH database.

1. On the Portal Server, start the Tealeaf Database Manager.
2. From the TDM menu, select **Mode** > **Database Setup** > **Single Database Mode**.
3. Drop the Search (TL_SEARCH) database:
   a. From the Single Database Options screen, select Search from the Database drop-down.
   b. From the Action drop-down, select Uninstall.
   c. Click **OK**.
   d. The database is uninstalled.
4. Create the Search database:
   a. Use the Search database creation script, which is provided in the following location:
      <Tealeaf_install_directory>\SQL\DBCreateScripts\CREATE_TL_SEARCH.sql
   b. Start the Tealeaf Database Manager.
   c. From the TDM menu, select **Mode** > **Database Setup** > **Single Database Mode**.
   d. From the TDM menu, select **DB Creation** > **Create Databases** so that the option is not selected.
   e. From the Single Database Options screen, select Search from the Database drop-down.
   f. From the Action drop-down, select Install.
   g. Click **OK**.
5. The Search database is installed with the proper collation setting.
6. If you are performing further installation operations, you might want to re-enable the Create Databases option. From the TDM menu, select **DB Creation** > **Create Databases** so that the option is selected.

Configuration issues

Data collection

Poor data collection performance

If the data collection process is consuming large volumes of memory or failing to finish, you may need to adjust the Data Collector Batch Size setting, which defines

the maximum number of records to extract or load for the Data Collector in a single batch. The default value is set to `2000`.

Note: Do not set this value over `5000`.

See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

Slow data collection times

Typically, data collection should run in under 5 minutes. If the data collection process is more than 10 minutes, then you should increase the size of the Data Trim Interval setting.

The default value is `1 - Hourly`. Setting it to a larger value causes this process to run less frequently. However, you should see temporary spikes in the size of your database.
See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

Database connection timeouts

If your Tealeaf databases are experiencing repeated timeouts, you may need to adjust the Database Connection Timeout setting. The default connection timeout is set to `30` seconds. After the system was running for a while and traffic was increased, your hardware may not be able to manage the traffic level at the current setting.

Try setting this value to double its current value. If it happens again, double it again.

Note: If you must set the connection timeout above 300 seconds (5 minutes), you may have more issues in your environment.

See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

Data collection performance from multiple canisters

By default, Tealeaf can collect from two canisters at one time. If your environment has more canisters, you can increase the size of the `Data Collection - Max Concurrent` setting. The appropriate setting depends on your hardware environment.

- See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

If you notice any of the following:

1. Some slow-down in data collection
2. Errors in the data collection process
3. The canisters are busy

Then you should try decreasing this setting.

Database is getting too large

If the size of your database is pushing the limits of your storage, you may consider adjusting some settings.

1. If you do not use the Page Performance reports, you do not need to collect the Path Statistics. You can disable the Path Statistics collection to save significant database space.
   - See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

2. In SQL Server Management Studio, check the Recovery Model setting for each database. If the Recovery Model is set to `Full`, the databases and their log files can grow larger. See "Database Sizing" in the *IBM Tealeaf Databases Guide*.

Error - an existing connection was forcibly closed by the remote host

From time to time, the Data Collector may register the following log message:

```
An existing connection was forcibly closed by the remote host
```

This error indicates that SQL Server or the server hosting it closed the connection unexpectedly. Typically, SQL Server was restarted, or an operation that required closing all existing connections was run, such as a backup or restore operation.

When these errors occur sporadically, they do not affect data collection or the validity of the data.

Note: If these errors are occurring frequently, verify that data collection is actually completing. You can review the log information for the most recent Data Collector that is run through the Portal. See "Portal Logs" in the *IBM Tealeaf cxImpact Administration Manual*.

Visitor Database Extractor

The Visitor Database Extractor is used to extract session data from the database and to insert it into the staging tables for the Visitor database. This section provides some suggestions for how to troubleshoot issues with the VDB Extractor.

- See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.

Unable to search for visitors

If you are unable to complete searches of the Visitor database or are no longer getting updated results, the Visitor database extractor job may be disabled. For more information about enabling this job, see "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.

Intermittent Search Server connection errors report error code 12029

Periodically, Search Server may return errors with the following error message.

```
An error occurred while executing the search.
<ServerName> - Error: (12029) A connection with the server could
   be established
```

If the above situation occurs, the solution is to apply the steps that are listed in the section below.

TCP Connection errors in VDB logon Windows Server 2003 SP2

If the VDB Extractor process fails to process sessions, and you see errors of the following general format in the `TL_VDB_Extractor_<TIMESTAMP>.log`:

```
Extract from canister failed after <number of tries> retries.
   CanisterId: <Canister Id>. Exiting.
```

In the above case, the Visitor Database Extractor running on Windows Server 2003 SP2 is unable to proceed with extraction when reaching the maximum number of canister read errors.

This issue applies only to Windows Server 2003 SP2. The source may be one of several issues. Below is described some steps to help diagnose the root cause of the problem.

Complete the following steps.

1. From a Windows command-line shell, run the following command on the remote server hosting Search Server.

   ```
   netstat -an | findstr 19000
   ```

   - The above command checks the network status of the socket connections through port 19000, which is the default port that is used by Search Server. Change this value if Search Server is using a different port. See "Managing Tealeaf Servers" in the *IBM Tealeaf cxImpact Administration Manual*.
2. If the results returned from the netstat command include multiple sockets in the FIN_WAIT or LAST_ACK state, the Search Server is not yet acknowledged the connection termination request that is previously sent to the remote TCP client. These sockets are occupied and may be causing TCP/IP port exhaustion. More connections to Search Server may not be possible, leaving Search Server in a frozen, and unavailable state.
3. The solution may be caused by issues in how the server is configured. For more information about the Microsoft solution to this issue, see the Hot Fix at http://support.microsoft.com/kb/979230.
4. After you applied the Hot Fix, reboot the server hosting Search Server.

Troubleshooting database upgrades

Invalid filegroup specified error during database upgrade

During the upgrade of one or more Tealeaf databases, you may receive an error similar to the following:

```
Invalid filegroup <name> specified.
```

This error indicates that the database file (<name>) is missing.

To fix this issue, you must determine if the database containing the file is supposed to contain a single filegroup or not.

If the database contains a single filegroup

Note: The only Tealeaf database containing a single filegroup is the TL_REPORTS database if it was upgraded from a Release 4.x version of Tealeaf. All other Tealeaf databases contain multiple filegroups.

Try reinstalling the database through the Tealeaf Database Manager. See "Installing Tealeaf Databases" in the *IBM Tealeaf Databases Guide*.

If the database contains multiple filegroups, including the missing one:
If the database is not a single filegroup database, a missing filegroup indicates database corruption. This serious issue may require rebuilding the database or restoring from backup.

- For more information, contact Tealeaf http://support.tealeaf.com.

Checklist

If you are still experiencing difficulties with your Tealeaf databases, acquire the following information from your system:

1. Acquire the full specs of the box running the Reporting/SQL Server, including CPU, RAM, and information about each hard disk drive. Drive information should include number of drives, contents of each drive, and available disk space.
2. Acquire the version of Tealeaf that is in use. Improvements may be made in subsequent Tealeaf releases.
3. What is the traffic load? How many pages per day? How many events per page? How many pages per session?
4. Is the entire Portal sluggish or certain items in the Portal?
   - If the latter, the problem is likely caused by issues with the code or configuration of the individual item(s).
5. Set `Log Level = 9`. Acquire a day or two of `TealeafReportingSvc.log` messages.
   - This information is useful for troubleshooting.
6. On the SQL Server, open **Performance Monitor**. Watch the Avg. Disk Queue and Page/Sec counters for 5 to 10 minutes. Are either of them consistently pegged through the roof?
   - Does **Performance Monitor** show much paging on the box? This may indicate that the Data Service is starved for memory.
7. How large is the TL_REPORTS database?

When you acquired the above information, open a support ticket with http://support.tealeaf.com.

Tuning Canister performance

The Tealeaf Short-Term Canister (STC) is an in-memory datastore for processing session data from active sessions on your web application. When a session ends or is timed out, its data are moved from the Short-Term Canister to the Long-Term Canister, which is stored on disk.

- The Tealeaf Canisters rely on the c-tree database solution. For more information, see `http://www.faircom.com`.

Depending on system hardware, web traffic, and configuration options, the STC can become backlogged, use up all permitted memory, and begin spooling incoming hits to disk via the DecoupleEx pipeline agent for later processing. This section provides some guidelines for handling high memory usage conditions.

Monitoring Canister status

Through the Tealeaf Portal, you can monitor Canister status in the DecoupleEx System Status report.

1. Log in to the Tealeaf Portal Application as an administrator.
2. Under System Reports, there is a Canister status and a DecoupleEx status report:
   - **Canister Status:** Provides information about the current web traffic as well as STC memory/disk usage.
   - **DecoupleEx Status:** Provides information about current DecoupleEx status and spooling.

- See "System Status" in the *IBM Tealeaf cxImpact Administration Manual*.

In the System Status report for DecoupleEx, look for the `Canister Status` value for each Processing Server.

`Canister:Spooling` status

This status indicates that the Canister exceeded one or more of the performance thresholds, and DecoupleEx began spooling incoming hits. In this state, the Canister is processing the data already in memory, and new data arriving to the Canister is written to the disk spool to be read into the Canister later.

- If its `Persistence` setting was enabled for the DecoupleEx session agent, restarting services or rebooting retains the spooled data on disk.
- If `Persistence` is disabled, the disk spool is deleted when the Transport Service running DecoupleEx is started.
- See "Extended Decoupler Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

If the DecoupleEx System Status report indicates high memory usage for `Canister Reason`, the likely explanation is that there was a surge in captured web traffic. While you can adjust Tealeaf system settings to increase the data throughput in the Canister, making these changes provides sufficient increases only if the spooled backlog represents a time period considerably longer than the specified session timeout. For example, if the Canister is configured to have a `Session Idle Seconds` setting of 300 seconds, the spooled backlog must be over one hour old to be worth reconfiguring the Canister settings.

- See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.
- For more information about configuring session timeout settings, see “Setting the session timeout” on page 67.

`Canister:Real-Time` status

The STC is available to process incoming hits in real time.

`Canister:N/A` status

Canister status is not available currently. Canister might be offline or otherwise not ready to process data.

Canister spooling

During spooling, the Canister Manager service and its child processes continue to process the data that is inserted into the Short-Term Canister. When DecoupleEx stopped the flow of hits into the Canister, the c-tree memory utilization gradually decreases as sessions begin to time out and are moved out of the Short-Term Canister.

When the c-tree memory utilization level drops below the predefined threshold, DecoupleEx begins inserting data into the Short-Term Canister again. This condition is displayed as `Canister:Real-Time` on the DecoupleEx System Status report.

- Memory utilization levels can be configured by parameter in the Extended Decoupler Session Agent. See “Configuring c-tree memory utilization” on page 64.

The canister oscillates between `Real-Time` and `Spooling` states as long as the volume of spooled data exceeds the Short-Term Canister's memory capacity. Additionally, the Portal's Activity Reports indicate that website activity is occurring

later than the actual time of capture, because these statistics are recorded after insertion into the Canister by the STC processes. It is common for some hours on these activity reports to indicate lower than expected counts due to spooling. Hours after these artificially low-activity periods on the reports may show above normal counts because spooled hits were being re-inserted into the Canister from the spool at a faster rate, generally at the maximum safe rate that is configured in the Extended Decoupler session agent.

When spooling does occur, you should observe the Canister status for a sufficient period to identify behavior during periods of low web traffic. During these periods, the Canister should be able to catch up. New spool files will continue to be created as old spool files are processed, as spooling does not stop until the consumption rate of spool data exceeds the creation rate for long enough that the last spool file is consumed before a new one can be created.

If spooling is becoming a chronic condition, then one of the following solutions may alleviate the problem depending on the root cause:

- Reducing the workload
- Throughput optimizations
- Creating a faster and/or larger Tealeaf hardware cluster

Configuring c-tree memory utilization

In the configuration for the Extended Decoupler session agent, you can define the Canister memory threshold percentages at which the session agent begins spooling incoming hits to disk and at which the session agent resumes sending hits to the Canister after spooling began.

- These settings can be configured through the Pipeline Editor in TMS. The parameters for the DecoupleEx session agent to modify are listed in the Display Name column below. See "Tealeaf Management System" in the *IBM Tealeaf cxImpact Administration Manual*.
- You can configure these parameters in the `[DecoupleEx]` section of the `TealeafCaptureSocket.cfg` file. The parameters to change are listed in the Internal Name column.

Table 1. Canister spooling

Display Name	Internal Name	Description
`Canister Max % Memory Used`	`CanCheckMaxCtreeMemUsedPct`	Maximum allowed percentage of memory that is allocated by the Tealeaf Canister Server. When set to 0, the cache value is ignored, and the CanCheckMinCtreeMemUsedPct setting is also ignored. • The default value is `80` percent.

Table 1. Canister spooling (continued)

Display Name	Internal Name	Description
Canister Min % Memory Used	CanCheckMinCtreeMemUsedPct	Specifies the minimum percentage of FairCom Cache in use. • This value is checked only if the CanCheckMaxCtreeMemUsedPct setting is enabled. • If the maximum value was exceeded, the Extended Decoupler queues hits until the minimum value is reached. • The default value is 50 percent.

See "Extended Decoupler Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Canister spooling and the Archive Manager

Note: As of Release 7.1, Archive Manager was deprecated.

- **Release 7.0 or earlier:** If you are using the Archive Manager, spooling can cause it to delay the archiving of a day if any data from that day is still in the spool. The processing of session data for the day is delayed until the next one. Typically, this issue is non-critical. However, you should assess that whether web traffic levels drop low enough so that all spooled data can be reinserted into the Canister before the next scheduled archiving process.

If the data spool is not cleared within 24 hours, the archiving process may be postponed for multiple days.

Workarounds:

- You may increase the value of CanisterKeepDays by 1 in archiver.cfg. This setting indicates the number of days that a Canister data file is retained on the Processing Server before the Archive Manager service attempts to retrieve it. Increasing this value should eliminate this error. However, the storage on the Processing Server must have sufficient space to store an additional day's data.
- The LockOutPeriod value in archiver.cfg could be extended so that archiving begins later in the day immediately after the data was captured. However, the archiving process may still be running when Web traffic begins increasing for the current day and thus burdening the Processing Server.

Applying Canister settings

Note: After modifying Canister settings and restarting the Transport Service, hits may be spooled until the Canister is up and running. In some situations, the Extended Decoupler session agent does not receive notification that the Canister was restarted, and hits continue to be spooled even though the Canister is ready to receive them. In this case, a restart of all Tealeaf services is required. See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

Tealeaf sizing calculations

The following metrics help calculate the required sizing for the Tealeaf Canisters to process your web traffic without significant spooling.

- These values are also useful for making calculations for disk storage and database sizing.

Peak values can be calculated based on the DecoupleEx logs. These files are stored in the following files:

```
<Tealeaf_logs_directory>\CSS_*
```

Where:

`<Tealeaf_logs_directory>` is specified in the DecoupleEx configuration. See "Extended Decoupler Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

In the table below are the performance/sizing metrics. To assess your system requirements, the absolute minimum information that is required for sizing calculations are the first two items; the other information provides a more accurate assessment.

- Activity Reports are available through the Tealeaf Portal. See "Tealeaf Report Builder" in the *IBM Tealeaf Reporting Guide*.

Metric Data source

Average Number of Sessions (or visits) Per Day
: See the Session Count report in the Portal Activity Reports. Average this value over as many days as makes sense.

Average Number of Page Views Per Day
: See the Page View Count report in the Portal Activity Reports. Average this value over as many days as makes sense.

Average Number of Hits Per Day
: See the Hit Count report in the Portal Activity Reports. Average this value over as many days as makes sense.

Average Session Duration
: See the Session Duration report in the Portal Activity Reports. Average this value over as many days as makes sense.

Average Page Size
: See the Page Size report in the Portal Activity Reports. Average this value over as many days as makes sense.

  Note: If possible, filter on the Tealeaf reference dimension values to remove image files and other static content from the report.

Average number of hits/pages per session
: See the Session Avg Hits report in the Portal Activity Reports. Average this value over as many days as makes sense.

Number of Sessions during Peak Hour
: See the Session Count report in the Portal Activity Reports. Use the peak hour from the peak day.
  - As an alternative, you can use the DecoupleEx log files, whose location is listed above.

Number of Page Views during Peak Hour
: See the Page View Count report in the Portal Activity Reports. Use the peak hour from the peak day.
  - As an alternative, you can use the DecoupleEx log files, whose location is listed above.

Number of Hits during Peak Hour
See the Hit Count report in the Portal Activity Reports. Use the peak hour from the peak day.

- As an alternative, you can use the DecoupleEx log files, whose location is listed above.

Number of Days Persisted
This setting is defined in the Canister configuration. See `Number of Days to Retain Data` in "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

Session Timeout
This setting is defined in the Canister configuration. See `Session Idle Seconds` in "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

Growth Factor
You can use either expected or measured growth of the data volume as the metric.

Note: After you acquired the above information, contact Tealeaf Professional Services to assess your system health and hardware requirements.

Setting the session timeout

The Short-Term Canister determines when to end a session that is based on the `Session Idle Seconds` setting for the Canister. When no new hit was inserted into the Canister for a session for a period greater than the value of `Session Idle Seconds`, the STC closes the session and flags it for downstream processing. This timeout scheme parallels standard web application practices.

- For more information about configuring this setting, see "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

The following secondary settings may contribute to deciding when to end a session:

- Events may be defined to close sessions that are based on the occurrence of the event. For example, if there is a logout link or button in your web application, you might define a close session event that is triggered when that link or button is pressed.
- To protect the STC, the Canister pipeline agent in the Transport Service imposes limits on session size and duration. By default, these limits are set to 2048 total hits, 5 MB of total memory, and 1 hour of total session duration. After changes any of the above settings, the Transport Service must be restarted.
  - For more information about the total session settings, see "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

Why are sessions fragmented?

Any of the following three causes may create fragmented sessions:

- A visitor's session is fragmented when the Tealeaf inactivity timeout period is exceeded. For example, if a visitor leaves a web page open during lunchtime, browsing again after lunch in the same browser window may cause a fragmented session. Because of the long period of inactivity, the first few pages are moved out of the Short-Term Canister (STC) into the Long-Term Canister (LTC). The next pages after lunch appear in the STC and are given a unique

session fragment ID. However, since the new session is the same browser window, the TLTSID session cookie value does not change.

- A visitor session that exceeds any of the three 'safety' limits (time duration, number of hits, or bytes) is moved out of STC into LTC to prevent unnecessary consumption of STC space. After the session is moved to LTC, the next hit with the same TLTSID generates a new session fragment in the STC.
- A session opened anytime the Canister services are shut down for maintenance is moved to the LTC. Subsequent hits with the same TLTSID create a new session fragment in the STC.

Tealeaf services cannot connect to canister as TLUSER or ADMIN

If the TLUSER and/or ADMIN userids are not present in the Canister datastore, the other Tealeaf services are prevented from connecting to the Canister.

The following steps rebuild the base Canister files, saving the existing Canister session and search index data:

1. Stop all Tealeaf services.
2. Rename `CANISTER.dbs` directory to `CANISTER.dbs.old`.
3. Rename `Indexes` directory to `Indexes.old`.
4. Run `CanRebuild.exe`, choosing to rebuild only the full Canister and nothing else.
5. Rename `Indexes.old` directory back to `Indexes`.
6. Move (do not copy) the following files from `CANISTER.dbs.old` to the new `CANISTER.dbs` directory:

   ```
   LSSN * PEVT * NDLY.dat LDLY.dat PATH.dat SRVR.dat EVNT.dat SEVT.dat
   ```

7. Start all Tealeaf services.

Canister Services continue to restart

If you disabled Canister Services through the Windows Services Control Panel and the service continues to restart, you should verify that the network monitoring and infrastructure diagnostic tools such as Big Brother are not performing these restarts.

Assessing canister storage allocation

If you are monitoring multiple sites or areas of your website using a single canister, you can use the following process to assess the canister storage requirements for each site:

Note: These steps assume that reference dimensions was properly configured to identify each individual site or area. See "Tealeaf Event Manager" in the *IBM Tealeaf Event Manager Manual*.

1. On the canister storage device, find out how much total storage space is used per day.
2. In the Tealeaf Portal, select **Tealeaf** > **System Status** > **Canister**.
3. The ratio of the session counts for each host that is compared to the total session count should give a rough estimation of how much space is used.

If the sessions from a specific host are much bigger than average, then this estimate may not be accurate. It should provide a good estimate.

Assessing Canister session agent performance

The Canister session agent supports multiple threads, which enable you to increase the hit insert rate by routing hits through multiple Tealeaf pipelines into the short-term canister. Before adding pipelines, you should assess the current Canister insertion performance.

Typically, bottlenecks in the Canister session agent occur while compressing the hit to better use STC memory. By enabling the display of the session agent's performance statistics, you can monitor the hit insertion and compression rates.

In the Canister session agent configuration section in `TealeafCapturesocket.cfg`, set the following property value:

```
DisplayPerf=true
```

This option enables the output of the maximum processing rate for the session agent to the **Queued** field of the Tealeaf Pipeline Status report.

- For more information about Tealeaf Pipeline Status, see "Tealeaf Management System" in the *IBM Tealeaf cxImpact Administration Manual*.

When enabled, the **Queued** field of the Canister session agent in the report displays a number like `30360558` as shown below.

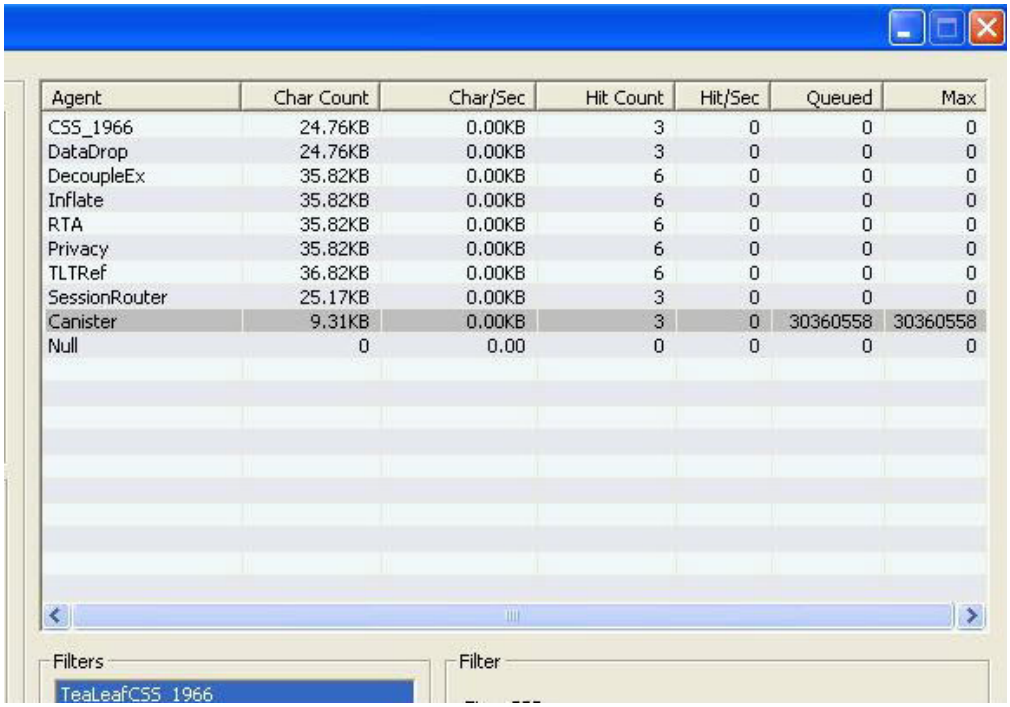

You can use this figure to determine if you have a bottleneck in the Canister session agent and the number of pipelines that are required to service the wanted hit rate. In the number `30360558`, the last four digits indicate the maximum compression rate into the STC, and the first four or five digits indicate the maximum insertion rate into the session agent. In the above example:

- **Maximum insertion rate** = `3036` hits per second
- **Maximum compression rate** = `558` hits per second

Backing up Canister files without Archive Server

If your Tealeaf system does not use Archive Server, you can use the following method to back up your canister LSSN files.

Note: When backing up the session data files, the canister must not be able to access the data files, or those files require a rebuild.

To back up the files with minimal processing server downtime:

1. Verify that no data is still to be added to the LSSN file of interest. For example, if you are backing up data files for yesterday, verify that no spool data is destined for yesterday.
2. Shut down the Canister Services.
3. Create a temporary directory.
4. Move the lssn data files (`lssn*.dat` and `lssn*.idx`) of interest into the temporary directory.
5. Restart the Canister Services.
6. Back up the files in the temporary directory.
7. Back up the corresponding dtSearch indexes.
8. Shut down the Canister Services.
9. Move the LSSN files in the temporary directory back into the active `CANISTER.dbs` directory.
10. Restart the Canister Services.

When restoring the LSSN files to the live canister, verify that CanTrim is either disabled or the setting for number of days to retain is large enough so that the restored data files are not trimmed on the next startup.

Moving a canister

Moving Tealeaf canisters requires moving the session data files and their indexes and updating the appropriate configurations of the changes.

Note: This procedure should be scheduled during off-peak hours. During the period in which the canister is disabled and disconnected, all captured hits are dropped.

Note: If your Tealeaf solution was upgraded to Release 7.2 before upgrading to your current version, you may receive an Error 160 in the Application event log if the Release 7.2 version was not allowed to run for an entire day. See “Error 160(0) in NextLssnRec()” on page 71.

Moving the Canister Search Index data

1. Through TMS, open the Canister configuration:
   a. Log in to the Tealeaf Portal as an administrator.
   b. From the menu, select **Tealeaf** > **TMS**.
   c. Expand the Canister node.
   d. Click **Canister configuration**.
   e. In the Config Actions panel, click **View/Edit**.
   f. The Canister configuration is displayed. See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

2. On the Canister Services tab, change the value for `Location of Files to be Indexed` to the new location of this directory. Do not edit any other fields.
3. To save changes, click **OK**.
   a. In TMS, expand the Session Indexer node.
   b. Click **Index Service configuration**.
   c. In the Config Actions panel, click **View/Edit**.
   d. The Session Indexer configuration is displayed. See "Configuring CX Indexing" in the *IBM Tealeaf CX Configuration Manual*.
4. On the Indexing Options tab, clear `Build Temporary XML In Memory`.
5. Edit the value for `Path for Temp XML File` to be the new directory path to the temporary XML file location.
6. Select `Build Temporary XML In Memory`.
7. To save changes, click **OK**.

Moving the Canister Session data

1. Open the Canister Server node in TMS.
   a. In TMS, expand the Canister Server node.
   b. Click **Canister Server configuration**.
   c. In the Config Actions panel, click **View/Edit**.
   d. The Canister Server configuration is displayed. See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.
2. Set the value of `LOCAL_DIRECTORY` to be the full path to the parent directory of the `CANISTER.dbs` directory. Be sure to include the final backslash (\ ) in the value.

Completing the canister move

After making either or both of the above configuration changes:

1. Back up the existing canister directories.
2. Through the Windows Start menu, select **Stop all TeaLeaf services**.
3. Through the Windows Services Control Panel, stop the TeaLeaf Transport Service.
4. Move the Indexes directory and/or the parent directory of the `CANISTER.dbs` directory to the new locations specified in the above configuration settings.
5. After the files and directories were moved, you should remove the original directories that you backed up.
6. Use the `Start all TeaLeaf services` Start menu shortcut to restart all Tealeaf services.

ctree errors

The following types of errors may be generated by ctree database operations.

Error 160(0) in NextLssnRec()

Unindexed sessions

The following error may appear in the Application event log:

```
Event Type: Error
Event Source: TeaLeaf Session Indexer
Event Category: Indexer
Event ID: 9864
```

```
Date: 10/9/2009
Time: 2:41:05 PM
User: N/A
Computer: TLDB01
Description:
Failed to retrieve record from Canister. Error 160(0) in NextLssnRec():
retrieving first batch of unindexed LSSN records.. (Type: FAIRCOM; Code: 160).
```

The above error indicates the session indexer is querying for a unindexed session that is current being updated.

If the Event Source is the TeaLeaf Session Indexer as indicated above, the sessions are skipped for the indexing run and are picked up on the next one.

Note: If you upgraded a system to Release 7.2 and an LSSN canister file was already created for the current day, you may be receiving this error consistently. In Release 7.2, Tealeaf introduced transaction support in the LSSN Canister files. The existing current day's LSSN file does not include support for transactions, and the above error is a result of this lack of support.

The sessions that are waiting for indexing are still indexed; there is no risk of data loss. However, the eventlog continues to receive these errors messages until the current day's data is fully indexed. If these errors continue to show up in the eventlog after a couple of days, contact Tealeaf http://support.tealeaf.com.

System crash breaks reference issue between data files and indexes

When retrieving the first batch of unindexed sessions, the Canister is unable to pull any sessions.

- No errors are reported in the TLTMaint logs.
- Automated indexing processes are not able to run at all.

The following error may appear in the Application event log:

```
(11:29 Session Indexer) - Failed to retrieve record from Canister.
Error 160(0) in NextLssnRec(): retrieving first batch of unindexed
LSSN records.. (Type: FAIRCOM; Code: 160).
```

Fix Indexes:
In this case, a system crash may have added some bad file information to the indexes. To fix this issue, please do the following:

1. Through TMS, shut down Canister services. See "TMS WorldView Tab" in the *IBM Tealeaf cxImpact Administration Manual*.
2. On the Canister data volume, create the following directory at the root of the volume:

   `lssn\Canister.dbs`
3. From the Canister datastore, move the `LSSN_<Date>.` files to the above directory, where `<Date>` is the date of the server crash.

   Note: You must move the files out of the Canister datastore.
4. Restart Canister Services.
5. Check that the indexer error message is not longer present.

Recover missing data:
If these steps fix the indexing problem, then perform the following steps to recover the session data file moved into the `lssn\Canister.dbs` directory:

1. Delete the `*.idx` files from the `\lssn\Canister.dbs` directory.
2. Run the following command in a command shell:

   ```
   tltmaint -v -noserver -archiver -localdir \lssn
   ```

   Note: Depending on the file size of the session data, it may take some time to complete the above command.
3. When the command prompt returns, the last line indicates the status. A "no errors" message indicates that all is well.
4. Copy files with `LSSN_` prefix from the `\lssn\Canister.dbs` directory back into the canister data directory.
5. Restart Canister Services through TMS. See "TMS WorldView Tab" in the *IBM Tealeaf cxImpact Administration Manual*.
6. The session data is now available.
7. To test, search for sessions from the date of the server crash. See "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.

Error 69(0)

The following error may appear repeatedly in the Application event log:

```
Failed to retrieve record from Canister. Error 69(0) in ProcessSesn():
Could not update session CANISTER.dbs\LSSN_20120815_MyServer.dat 469778184.
isam_err = 69. (Type: FAIRCOM; Code: 69).
```

This error may appear even after recycling services and rebooting the Portal, and the Canisters are appearing to operate normally.

The above error typically indicates that the indexes in the specified `.dat` file are corrupted.

To repair this problem, complete the following steps to rebuild the indexes.

Note: Rebuilding Canister indexes can take a long time. You should run these steps during off-peak hours.

1. Log in to the server hosting the Canister as an administrator.
2. Navigate to the directory where the above file is located. Typically, this directory is the following:

   ```
   <Tealeaf_install_directory>\Canister\Canister.dbs
   ```
3. Move all files (there should be 3) matching the following name pattern out of the `Canister.dbs` directory:

   ```
   LSSN_20120815_MyServer*.idx
   ```
4. Restart Canister Services through TMS. See "TMS WorldView Tab" in the *IBM Tealeaf cxImpact Administration Manual*.
5. When the Canister restarts, it attempts to rebuild the indexes by reindexing all sessions that are stored on the Canister. Depending on the volume of data, this process can take multiple hours.

Configuration

For more information about configuration, see "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

Changing the logging level to troubleshoot exceptions

When an exception is detected in the Indexer log, raise the logging level for the indexer to 4 or 5, which enables the capture of a wider set of debugging data.

- When reporting an error to Tealeaf Customer Support , submit a log file that is set to level 4 or higher.

Steps:
Complete the following steps to change the logging level.

1. Log in to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **TMS**.
3. The Tealeaf Management System is displayed. See "Tealeaf Management System" in the *IBM Tealeaf cxImpact Administration Manual*.
4. From the **View** drop-down, select Servers.
5. Click the Session Indexer node.
6. Click **Index Service configuration**. In the **Config Actions** panel, click **View/Edit**.
7. The Index configuration is displayed.
8. Click the Scheduling/Diagnostic tab.
9. In the Diagnostics group, increase the level to 4 or 5.
10. Click **Save**.
11. Push the configuration to each server. A service restart should not be performed currently. See "Configuring CX Indexing" in the *IBM Tealeaf CX Configuration Manual*.

    Note: If you raise the indexing level for troubleshooting purposes, be sure to reset the logging level to a more appropriate value. The recommended value is 2.

"Sessions waiting to be indexed" broke its threshold

When using direct-pull indexing, "Sessions Waiting to be Indexed" is a fairly innocuous threshold, as no disk space or other machine resources are consumed for unindexed sessions.

In other cases, the system may be configured to write files to disk for the Session Indexer service to process. Before the advent of direct-pull indexing in versions 4.0 SP 2, a high "Sessions waiting to be indexed" was a sign of impending problems, as sessions to be indexed were written to disk as physical files.

For Tealeaf versions from 4.0 SP 2 to Release 7.1, you may want to disable checking this threshold in PortalStatus.cfg or increasing the threshold value so it does not trigger as often. To double-check the setting, run the IndexConfig.exe utility and click the Operation Times tab to see whether the "Enable direct communication with Canister" check box is selected.

- For Release 7.2 or later, Index configuration is managed through TMS. See "Configuring CX Indexing" in the *IBM Tealeaf CX Configuration Manual*.

You can either increase the WaitToBeIndexed Error threshold value in ...\TeaLeaf\Tools\PortalStatus.cfg or set it to 0 to completely disable checking of this threshold.

Some sessions seem to be missing from the Indexes

If the Session Indexer service is in the stopped state, sessions are not indexed when they move to the Long-Term Canister. If you suspect some sessions were not indexed or that the Indexes may not be up to date, you can use the Search Server's "Check and Fix Indexes" feature.

Note: Before applying this fix, verify that the file `<Tealeaf_install_directory>\IndexCheck.exe` was upgraded or patched.

1. Open a Web browser to the following address:
   `http://<hostname or IP address of Tealeaf server>:19000/`
2. Click the Canister/Indexer Check link. The table on the resulting page shows how many sessions exist in each day's data and how many of those sessions were
   indexed.
3. Click **Check and Fix**.
4. The system indicates the number of sessions that are needed to be indexed and begins the process. The process continues in the background even if you close the browser window.
   a. To check the status of the process after the browser window is closed, open a new browser window and enter the following URL:
      `http://<hostname/IP addr of Tealeaf server>:19000/CanisterIndexCheckerStatus`

Failed to get number of sessions to be indexed from Canister

This error indicates that the TeaLeaf Session Indexer service cannot connect to the Canister Server datastore service. This situation may occur if the Canister Server service is in the stopped state or in maintenance mode.

When the Canister Server service is started, it begins maintenance mode by running runs the TLTMaint program. TLTMaint does an exhaustive check of the Canister data files. Although TLTMaint usually completes in less than a minute, it can sometimes take quite a long time to complete if there was an abnormal shutdown that invalidated any of the Canister's `.dat` or `.idx` files.

When the other Tealeaf services start, they try to open communications with the Canister Server. The services are well-behaved. If they cannot achieve communication right away after starting, they retry every 10 seconds or so and wait at least five minutes before reporting any connection errors to the Event Log.

If you look at the actual Event Log, you should see messages after the Session Indexer service error message that TLTMaint completed without errors, and the Session Indexer error messages should stop being logged.

Index error recovery

When an index job creates a new index or performs a merge, a new index is created and added to the `TLL` file immediately to ensure that the index name is reserved, but the index is marked as not valid. After the index or merge operation is complete, the index is marked as valid. If the index or merge crashes, then the index remains marked as not valid and remains locked. A lock is broken if it is in place for more than two hours. If during the `TLL` check, Index Check finds an index with no lock that is marked as not valid, the index is deleted.

If an error is returned during indexing, the index is marked as requiring an index check. This index is not used for indexing until the flag is removed. If the index check fails, the index is marked as corrupt and removed from the TLL. If the check succeeds, then the check required flag is cleared and the index is returned to service.

Corrupted index directories are renamed with a .CORRUPT file extension for later inspection.

Re-creating indexes

If necessary, you can re-create indexes by deleting the index directories and using the index check option in the Search Server. For more information, contact Tealeaf http://support.tealeaf.com.

Configuration

For more information about configuration, see "Configuring CX Indexing" in the *IBM Tealeaf CX Configuration Manual*.

RSE Service Fails

When analysis is performed on result sets, the RSE fails, generating the following error message:

```
Error getting session info for session 290610439: System.Xml.XmlException:
Invalid character in the given encoding. Line 3035, position 351. at
System.Xml.XmlTextReaderImpl.Throw(Exception e) at
System.Xml.XmlTextReaderImpl.InvalidCharRecovery(Int32& bytesCount, Int32&
charsCount) at System.Xml.XmlTextReaderImpl.GetChars(Int32 maxCharsCount) at
System.Xml.XmlTextReaderImpl.ReadData() at
System.Xml.XmlTextReaderImpl.ParseCDataOrComment(XmlNodeType type, Int32&
outStartPos, Int32& outEndPos) at
System.Xml.XmlTextReaderImpl.ParseCDataOrComment(XmlNodeType type) at
System.Xml.XmlTextReaderImpl.ParseElementContent() at
System.Xml.XmlReader.ReadString() at
TeaLeaf.SearchServer.XmlParsing.SessionInfoFromMemStream(MemoryStream
memStream, TLSession sessn, Boolean parseSummary) at
TeaLeaf.SearchServer.TLSession.GetSessionInfo(Boolean rsp, Boolean
interestingHits, Boolean parseSummary, Int32 timeout, Boolean validateXML,
Boolean SessionXML, Boolean log)
```

This problem results from failing to have the proper UTF-8 encodings in the request buffer. To support UTF-8 encoding, the following name/value pairs must appear in the request buffer:

```
REQ_BUFFER_ENCODING=UTF-8
REQ_BUFFER_ORIG_ENCODING=ISO-8859-1
RESP_BODY_ENCODING=UTF-8
```

To fix this issue, one of the following upgrades must be completed:

- Upgrade to PCA Build 3330 or later. I18N must be enabled.
- IBM Tealeaf cxImpact Build 7080 or later. Specifically, the pipeline session agent SAInflate must be version 7080 or better.

Query Server fails to shut down and restart properly

Upon restart of the Query Server, you may see errors like the following in the event log:

```
Errors:
7/28/2008 10:50:00 AM QueryServer ERROR An error occurred in the listener for
10.10.32.252:21000 : Only one usage of each socket address (protocol/network
address/port) is normally permitted
7/28/2008 10:50:00 AM QueryServer ERROR An error occurred in the listener for
127.0.0.1:21000 : Only one usage of each socket address (protocol/network
address/port) is normally permitted
7/28/2008 10:50:00 AM QueryServer ERROR at
System.Net.Sockets.Socket.DoBind(EndPoint endPointSnapshot, SocketAddress
socketAddress)
at System.Net.Sockets.Socket.Bind(EndPoint localEP)
at System.Net.Sockets.TcpListener.Start(Int32 backlog)
at TeaLeaf.Reporting.Server.TLHttpServer.Listen(Object obj)
```

If you are seeing the above error messages, the issue may be caused by a failure in the RSE Service to shut down properly. In some cases, the service remains operational for a few seconds after the shutdown command is issued.

When a restart of the RSE Service is attempted, the Service fails to bind to the assigned port. When this error occurs, the service must be killed manually through the Windows Task Manager (or equivalent tool). Then, the RSE Service can be restarted as part of normal Tealeaf startup procedures.

This issue may be addressed in subsequent builds of your Tealeaf version. For more information, contact Tealeaf http://support.tealeaf.com.

Troubleshooting - Data Collector

Acceptable error messages

In the Portal Data Collector log, one or more of the following error messages may appear:

```
A transport-level error has occurred when sending the request to the server.


An existing connection was forcibly closed by the remote host.


A network-related or instance-specific error occurred while establishing a
connection to SQL Server. The server was not found or was not accessible.
```

In all almost all scenarios where the above messages appear, the underlying issue is that SQL Server is rebooting or entered a maintenance mode. These errors can be ignored, unless they reappear frequently or at irregular and unexpected intervals.

Performance issues

The Tealeaf Data Collector runs at least one run every 5 minutes against all active and available Canisters when it is operating optimally. Due to the high number of transactions it performs, the Data Collector can run into performance issues that may or may not be sourced with Tealeaf components.

This section provides some information about how to begin troubleshooting the Data Collector's performance.

Periodical stalling

Periodically, it may appear that data collection is stalling. Data collection runs may fail to complete, with the following type of log message:

```
3:05:26 TeaLeaf Data Collector ERROR Data Collection appears to be stalled.
Data Collection has been running since: 2012-06-20 02:05:26
```

When this error message appears, verify through the logs the duration of the data collection run. If the data collection run takes longer than one hour, the process restarts automatically and a message is logged, similar to the following:

```
2012-06-20T10:40:26 TeaLeaf Data Collector STATUS The TeaLeaf Data Collector
is stopping.
The TeaLeaf Data Collector has been running for: 0 days, 8 hours, 6 seconds.
```

To lower the time to complete a data collection run, you can experiment by raising the following values:

Setting Recommended Adjustment

Data Collection - Max Concurrent
: Maximum number of Canisters that the Data Collector can query at the same time. Try bumping this value by 2, if you have more Canisters than the current value.

Data Collection - Batch Size
: Maximum number of records to collect in a single batch. Try increasing this value by 2000.

You may need to experiment with these settings through a number of collection runs to effectively tune the Data Collector's performance.

- For more information about these settings, see "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

Timeouts

Some data collection operations may time out depending on data load and system performance. When timeouts occur, messages similar to the following one appear in the Data Collector log:

```
Error occurred while executing the non-query command: Timeout expired. The
timeout period elapsed prior to completion of the operation or the server is
not responding.
```

Note: To resolve timeout issues, underlying system performance should be analyzed.

In the short term, you may try the following setting to allow the process to complete.

1. Log in to the Portal as an administrator.
2. From the **Portal** menu, select **Tealeaf** > **Portal Management**.
3. Click the IBM Tealeaf CX Settings link.
4. Click the **Data Collector** category.
5. Locate the setting `Database Connection - Timeout (seconds)`.
6. Click **Edit**.
7. Raise this value sufficiently to complete the process.

What happens to statistical data if the Data Collector Service or reports database goes down?

Every time events fire in a Processing Server, the Canister datastore on that Processing Server creates one or more records. In a functioning system, the Report Data Collector (RDC) service collects these data from every Processing Server and

stores the information in the reports database. The RDC usually runs on the same server that hosts the Portal. By default, it is scheduled to run a data collection 5 minutes after the completion of its previous run.

If the RDC cannot communicate with the reports database or cannot communicate with the Canister on the Processing Server(s) or if the RDC itself cannot run, the records accumulate in the Canister on the Processing Server(s). The next time the RDC does run successfully, all records in the Canister on the Processing Server(s) are collected and stored in the reports database.

The Canister on the Processing Server holds records locally for 48 hours by default. If the RDC is unable to read and reap these records for more than that amount of time, the Canister on the Processing Server(s) begins to delete the oldest records. Old records are deleted in 1 hour chunks.

- The length of time that records are retained is configured by parameter in the Services Perform tab of the Canister configuration. See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

Mismatches between counts in event charts and counts in drill-down searches

In an event chart or dashboard, the event counts for a given period may be reported. When you click the link to drill down to the underlying sessions, the number of returned sessions differs from the number reported in the chart.

- The above assumes that all events in the charts are configured to count only one time in the session. Otherwise, the counts between the charts and the number of returned sessions may differ.

The discrepancy between these values may not be answered through the Portal, as the issue may be related to how the Data Collector is collecting event counts from the Canister. Complete the following steps to see if the event counts match.

Note: Completing these steps requires the use of regedit and running an SQL query. If you do not feel comfortable with these tasks, contact Tealeaf http://support.tealeaf.com.

1. Log in to the Canister Server.
2. Through the **Windows Start** menu, open regedit.
3. Acquire value from the following registry key. Retain this value:
   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\TeaLeaf Technology\
      DataStore\Canister\Time-Statistics Table Trim Seconds
   ```
4. Set this value to `86400`, which is 24 hours. This setting effectively disables deletion of aggregated event data on the canister.
5. Insert the following registry key and value into regedit:
   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\TeaLeaf Technology\
      DataStore\ReportServer\Canister Trim Disabled
   ```
6. Set this value to `true`.

   Note: This registry setting is different from `Can Trim Enabled`, which is part of the default configuration for each Canister. Do not change that setting.
7. Save your settings.
8. Let the Data Collector run for several hours or until you find a discrepancy in the chart/search counts.

9. When a discrepancy is detected, complete the following steps to query the EVNT table on the Canister to find out if the event chart counts were collected properly.
   a. Open a command prompt using the following shortcut on the Canister Server:
      ```
      Start > All Programs > TeaLeaf Technology > Portal Server > Server >
         Client Tools
      ```
   b. At the command prompt, enter the following command:
      ```
      csql -p ADMIN
      ```
   c. Run the following SQL query:
      ```
      SELECT * FROM EVNT
      ```
10. Verify the counts that are returned from the query against those reported in the Portal.
11. Revert the value for the following registry setting in the Tealeaf Technology hive:
    ```
    \DataStore\Canister\Time-Statistics Table Trim Seconds
    ```
12. Disable or remove the following registry entry in the Tealeaf Technology hive:
    ```
    \DataStore\ReportServer\Canister Trim Disabled
    ```
13. Save your regedit settings.

Issues of Data Collector connectivity with the Canister

If the log files contain any of the following messages:

```
Could not connect to lpwsa3000:CANISTER Canister + {lpwsa3000, 5597}.
... CANISTER is not active
... Data Collection error on server ...
```

Then, the connection between the Data Collector and the Canister is not operational.

To resolve the issue, do the following:

1. Restart the Tealeaf Data Collector Service. If that does not resolve the issue, continue with the following steps.
2. Review recent changes to the system. If you upgraded Tealeaf recently, for example, some configuration changes may impact the Data Collector.
3. Check the name of the Canister server in the Portal Management page.
4. Verify the Canister user name and password through the Tealeaf Portal.
   - If you are using custom passwords for the `TLADMIN` and `TLUSER` accounts on the canister, verify that the passwords were properly entered in the Tealeaf Portal.
   - See "Managing Tealeaf Servers" in the *IBM Tealeaf cxImpact Administration Manual*.
5. If the above does not resolve the issue, verify that the Portal server can ping the canister server.
   - See "Managing Tealeaf Servers" in the *IBM Tealeaf cxImpact Administration Manual*.
6. If the above works, then verify that telnet from the Portal server to the canister server on port `5597` works.

Note: If you make changes to the Data Collector settings through the Portal, you should restart the service so that the settings can be applied immediately.

SQL exception "Connection reset by peer: socket write error" in event log

This error may occur when the Tealeaf Report Data Collector service is trying to communicate with SQL Server running on a remote server. It indicates that the remote SQL Server did not respond to the request by the Report Data Collector service to insert data. The error condition could be either network or SQL Server related.

Restarting the Tealeaf Report Data Collector service should clear up the problem, if it was a temporary issue with the remote SQL Server. If the error persists after restarting the Tealeaf service, contact the DBA responsible for the remote SQL Server or the network staff responsible for the connection between the Tealeaf server and the remote SQL Server. The root cause is probably found in one of these two areas.

Database Filegroup Report shows no data

After installation, the Database Filegroup Report may not display data. However, investigation of the underlying database tables may show that the data is present.

- This issue applies to new installations of Release 8.4 or later.

This problem is caused by a bug during installation in which the values for the Host and Visitor Host are left as `localhost` and are not mapped to the machine name. When the report attempts to query for the report data using `localhost`, the Portal fails to connect to the database, rendering an empty report.

- This issue can also occur if customers modified the above registry settings, although this situation is rare.

To fix this problem:

1. Log in to the Report Server as an administrator.
2. Navigate to the following registry hive (64-bit):
   `SOFTWARE\Wow6432Node\TeaLeaf Technology\DataStore\ReportServer`
3. Change the following settings:

Table 2. Database Filegroup Report shows no data

Setting	Current Value	New Value
`Host`	`localhost`	Full name of machine
`VisitorHost` (if IBM Tealeaf cxResults is installed)	`localhost`	Full name of machine

4. Save the registry.
5. Restart the Report Server.

Troubleshooting - data service

In the event of issues with the Tealeaf Data Service, provide the reporting log information for Customer Support. See “Reporting logs” on page 85.

- See "Configuring the Tealeaf Data Service" in the *IBM Tealeaf CX Configuration Manual*.

Failed to write ReportQueryLogEntity error message

Periodically, you may see an error message similar to the following in the Data Service log:

```
Failed to write ReportQueryLogEntity
```

The above error is recorded a transaction issue with an internal logging table. It does not impact users or performance in any way and can be ignored.

Troubleshooting - Reporting Service

Note: As of Release 8.0, the Reporting Service is a fully integrated component of the Tealeaf Data Service. See "Troubleshooting - data service" on page 81.

Troubleshooting - Report Builder

"There are no matches for your search" on drill-down from Report Builder

In the Report Builder, if drill-down is enabled, you can click links in the Report Builder detail table to review the list of underlying sessions that provided the source data for the item in the table.

In some cases, when you click a non-zero hyperlinked number, you may not receive any sessions in return, instead receiving a message similar to the following:

```
There were no matches to your search, as it may be one of XX session(s)
   pending indexing. Please wait a moment and try again.
```

where

- `XX` is a number

It is possible that this issue is caused by a backlog of sessions waiting to be indexed. Typically, however, the appearance of this message when a Report Builder drill-down is clicked is caused by an embedded limitation in the number of sessions that can be returned.

To prevent Search Server from crashing when retrieving too many sessions, a limit of 16,384 maximum sessions can be returned from a query, such as a drill-down query. If the number of requested sessions exceeds this limit, Search Server silently fails the request. As a result, the Portal receives no sessions back from Search Server and posts the above message.

- In the Search Server logs, a `Nothing Found` message is posted.
- A more descriptive message is included in the Search Server log in a later build:

  ```
  Max results = XX is too high for a same-page search. The limit is 16384.
  ```

  Where `XX` is a number.

When it was introduced in Release 8.0, the Tealeaf Report Builder integrated drill-down searches using And on Same Page queries, which can generate large numbers of results from Report Builder. In many cases, drill-down queries from relatively simple reports can exceed this 16,384 session limit.

To patch the issue, you can reduce the maximum number sessions that are returned in a session segment.

Note: The parameter in the following set of steps is also used by the Result Set Extractor for creating session segments, which may be analyzed through IBM Tealeaf cxResults. You should review any available information as to why the parameter is set to its current value.

1. Log in to the Portal as an administrator.
2. From the **Portal** menu, select **Tealeaf** > **Portal Management**.
3. From the left navigation pane, click the IBM Tealeaf CX Settings category.
4. Click the **Search** link.
5. For the `Session Segment Max Sessions Limit`, set the value to `1000` or less.
   - The default value is `1000`.
   - For customers who upgraded from Release 7.x or earlier, this value may be set to a considerably higher value for use in creating large session segments for IBM Tealeaf cxResults. This setting, however, is applicable to the Tealeaf Report Builder use case above and should be resized accordingly.
6. Click **Save**.
7. Try the drill-down link again.
   - If the link still returns 0 sessions, you should resize it to 1. If that fails, then reset the value to `1000` and contact http://support.tealeaf.com.

Troubleshooting - Top Movers

Top Movers not getting updated

When you configure a Top Mover, it may take a period for data to be collected for the Top Mover so that they can be displayed in the Portal. Hourly Top Movers are collected once per hour, and daily Top Movers are collected once per day.

Note: A newly created daily Top Mover does not begin displaying data in the Top Movers report until the following day's Data Collector run for Top Movers. By default, this run occurs at 5:30 in the morning.

For more information about creating Top Movers, see "TEM Top Movers Tab" in the *IBM Tealeaf Event Manager Manual*.

- For more information about reporting for Top Movers, see "Analyzing Top Movers" in the *IBM Tealeaf Reporting Guide*.

If you have waited 24 hours and are still not seeing data in your configured Top Movers, it may be caused by the following issue. In Release 8.2, it is a requirement that the time zone of the Report Server matchs the time zone for the Tealeaf system. If the time zones do not match, then no Top Movers data is collected. Nothing is displayed in the Top Movers report.

To address this issue, you must change the time zone of the Report Server to match the Tealeaf system time zone. The following sections provide the required steps.

Acquiring Tealeaf system time zone

1. Log in to the Portal.

   Note: The Tealeaf system time zone should be configured during initial installation. It should not be changed afterward. For this workaround, do not change the Tealeaf system time zone, which can affect many other components of the Tealeaf system.

2. From the **Portal** menu, select **Tealeaf** > **TMS**.
3. The Tealeaf Management System is displayed. See "Tealeaf Management System" in the *IBM Tealeaf cxImpact Administration Manual*.
4. Click the **WorldView** tab.
5. From the View drop-down, select `Servers`.
6. Click the master server.
7. Click the **Tealeaf node**.
8. Click **Tealeaf global configuration settings**.
9. Click **View/Edit**.
10. Review the value for `Roll Time Zone`. This value is the Tealeaf system time zone.

Review Top Mover calculation times in the Portal

Note: Depending on the number of Top Movers and the volume of uncollected data, the first recalculation may take a long time to complete. Before you begin, you should verify that the configured times for hourly and daily Top Mover recalculation, as expressed in Tealeaf system time, allows ample time to complete the recalculation before peak hours resume.

1. Log in to the Tealeaf Portal.
2. From the **Portal** menu, select **Tealeaf** > **Portal Management**.
3. From the left navigation panel, click the IBM Tealeaf CX Settings panel.
4. Click **Data Collector**.
5. Review the value for the `Top Movers - Time for Daily Calculation`. This value is the time in the Tealeaf system time zone when daily top movers are recalculated. Change it, if needed, to occur during an off-peak hour.
6. Hourly Top Movers are recalculated 10 minutes after the top of the hour. The first recalculation of hourly Top Movers may take a long time. If possible, you should schedule it during off-peak hours.
   a. To disable calculation of all Top Movers, set the value for the `Top Movers` setting to `Disabled`.
   b. After hours, this setting can be re-enabled.
      - Hourly Top Movers are recalculated at 10 minutes after the top of the next hour.
      - Daily Top Movers are recalculated at the `Top Movers - Time for Daily Calculation` time in the Tealeaf system time zone.

Changing the time zone of the Report Server

If you change the time zone of the machine hosting the Report Server to the Tealeaf roll time zone, then Top Movers are calculated normally afterward.

Note: Before you begin, verify that you have administrator access to the machine hosting the Tealeaf Report Server.

1. Login to the machine hosting the Tealeaf Report Server as an administrator.
2. In the Windows toolbar, click the time value in the lower-right corner.
3. In Windows Server 2008, click **Change date and time settings...**.
4. Click **Change time zone...**.
5. From the Time zone drop-down, select the time zone that corresponds to the Tealeaf system time zone.

Note: If this change is temporary, note the selected time zone before you make the change.

6. Click **OK** twice.
7. Top Movers are recalculated at the time in the Tealeaf system time zone that is based on the Portal settings.

Reporting logs

If with the Tealeaf Data Service or Reporting Service, provide the following information for Customer Support:

- All Data Service log files for the relevant dates. They are in the following location:

  `<TeaLeaf_install_directory>\Logs\TLDataService*.log`
- All Tealeaf Database Manager log files for the relevant dates. They are in the following location:

  `<TeaLeaf_install_directory>\Logs\TeaLeafDatabaseManager*.log`
- Any errors or warnings that appear in the Windows Application event log.

Configuration

For more information about configuring the Data Collector:

- See "Data Aggregation and Retention" in the *IBM Tealeaf cxImpact Administration Manual*.
- See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

For more information about configuring the Data Service, see "Configuring the Tealeaf Data Service" in the *IBM Tealeaf CX Configuration Manual*.

For more information about configuring the Report Server, see "Configuring the Report Server" in the *IBM Tealeaf CX Configuration Manual*.

"Auth failed, no Authorization header on cmd" errors in Search Server log

When Portal or NT authentication is enabled for Search server, errors similar to the following may appear in the Search Server log:

```
Auth failed, no Authorization header on cmd:
   /EventDefinitions</Authentication>
Auth failed, no Authorization header on cmd:
   /IndexSearchParallelStart</Authentication>
Auth failed, no Authorization header on cmd: /SessionInfo</Authentication>
Auth failed, no Authorization header on cmd: /ResultSet</Authentication>
Auth failed, no Authorization header on cmd: /TLLList</Authentication>
```

Some number of these messages is expected during normal use of the product.

Under NT authentication or Portal authentication, each time that a user logged in and attempts to query Search Server for data that is protected by this authentication, one instance of this message is generated.

When Search Server is configured for NT or Portal Authentication, a series of headers is used to validate requests made by Tealeaf services on behalf Tealeaf users.

An Authentication header cannot be added by the Tealeaf service calling to Search Server until the server responds with the challenge string. This string is then used in the header of the subsequent query to Search Server. Since Search Server cannot anticipate whether the initial query is followed by the subsequent query, the initial request is rejected, and these messages are recorded in the log, if they are, in fact, errors.

- The reject response from Search Server contains a set of challenge headers that are used to generate valid subsequent requests with the proper authentication information.

Note: If these errors messages appear at very short intervals or appear with an inability to perform some search functions, they may be indicative of a problem with how authentication is configured and should be investigated. See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

Under normal conditions, one of these entries is almost always followed by a normal log entry for the same command, which indicates that the client successfully supplied the correct authentication header on the second try.

Unable to add Active Directory groups through Search Server configuration

If you added a new domain or enabled new Active Directory groups in your domain, you may not be able to add them through Search Server configuration. When attempting to add the groups through TMS, you may receive the following error message:

```
Error getting domain groups!
```

The Search Server log may contain an error message similar to the following:

```
Authorization error: NetGroupGetUsers() error
   (domainServer: \\MyServer group: corp\TealeafPortal): Access is denied.
```

This issue may be caused by access permissions to the set of groups on the domain. In many environments, new domains are defined to permit any client, such as Search Server, to be able to access the group information anonymously. This "Anonymous Binds" security policy in Active Directory enables Tealeaf to use the group information for authentication.

In some environments, this anonymous access is not permitted. You may be able to access the group information by running the Search Server configuration utility on the authentication master server using a named user. Complete the following steps.

1. If you do not know it already, locate the name of the authentication master through TMS:
   a. Log in to the Portal as an administrator.
   b. From the **Portal** menu, select **Tealeaf** > **TMS**.
   c. The Tealeaf Management System is displayed. From the WorldView tab, select **Servers** from the View drop-down.
      - For more information about TMS, see "Tealeaf Management System" in the *IBM Tealeaf cxImpact Administration Manual*.
   d. Expand the Search Server configuration node.
   e. Click **Search Server configuration**. In the Config/Actions panel, click **View/Edit**.

f. The Search Server configuration is displayed. In the **Authentication** window, review the value for the Auth Master setting.
g. This value is the Authentication master in your Tealeaf environment.
h. Cancel your changes, and exit TMS.

2. On the Authentication Master server, you must configure the Search Server configuration utility to run as a named user that has access to this user group information.
   a. Log in to the Authentication Master server as an administrator.
   b. From the Start menu, open the Windows Services control panel.
   c. Right-click **TeaLeaf Search Server**. Select **Properties**.
   d. On the Login tab, select to use `This Account`.
   e. Enter the username and password for a user that is permitted access to domain group information.
   f. Save changes.
   g. Restart Search Server.

Chapter 3. Troubleshooting - Replay Server

This section provides troubleshooting solutions that are related to the Replay Server, which serves sessions for replay to Browser-Based Replay clients.

Troubleshooting - rendering issues

Steps:
If you believe that the Replay Server is struggling to render the content, complete the following steps to identify render times per page.

1. Log in to the Replay Server interface.
2. Click the **SessionList** link.
3. Select a session that was loaded.
   - If no session was loaded, you can use the home page of the interface to load one from your local desktop.
4. Select a session with a link like the following: `SessionNNNN_NNNNNNNNNN`.
5. Click the **NavList** link. All pages, their render times, and other information are displayed.
6. Select individual pages to review their rendered HTML, looking for irregularities.
7. You can also check the Page Load Details (PLD) link for each page, where render times and load times are published. Less than 10 seconds for each page is considered a reasonable time.
8. If you see load times that significantly exceed 10 seconds:
   - If the excessive load time applies to a single page, there may be problems with accessing content referenced on the individual page. Check out the references on the page, and try to load them in your browser.
   - If the excessive load times apply to multiple but not all pages, see if you can identify whether the slow times apply to individual servers, which may indicate connectivity or network issues.
   - If the excessive load times apply to all pages on the server, verify that the issue applies to multiple sessions. If so, there may be a problem with Replay Server. Contact Tealeaf http://support.tealeaf.com.

BBR Replay view displays response with encoded characters

In Replay view, BBR may display the response of the hit. However, some of the characters are encoded (for example, < is displayed instead of <. Other characteristics of replay from the same desktop:

- RTV works fine
- BBR Response view works fine

Below is an example output from BBR Replay view showing Response view with encoded characters:

```
&lt;!doctype html&gt;&gt;!--[if ie]&lt;&gt;![endif]--&gt;

<html lang="nl">
<head id="ctl00_Head1">
  <meta http-equiv="X-UA-Compatible" content="IE=8" />
  <meta charset="utf-8" />
```

```
<title>example.com - Here's my web site</title>
  <meta http-equiv="Content-Language" content="nl" />
  <meta name="copyright" content="MyCo - All rights reserved" />
  <meta name="author" content="MyCo" />
  <meta content="text/javascript" http-equiv="Content-Script-Type" />
  <meta id="ctl00_metaRobots" name="robots" content="NOODP" />
  <meta name="verify-v1" content="VQALIO5JJtratMWQdcnDiNmlqN/vZObY5iifah8=" />
```

The source of this output irregularity could be one of a number of issues, including the following:

1. Portal times out waiting for the page to return
2. Replay Server renderer is waiting for content to be delivered from origin servers
3. Replay Server does not have connectivity to static content
4. Other issues

See “Troubleshooting - rendering issues” on page 89.

Configuration

For more information about configuration, see "Configuring the Replay Server" in the *IBM Tealeaf CX Configuration Manual*.

BBR Sessions do not replay in BBR due to Proxy Errors in Replay Server

BBR may hang during replay of sessions or fail to display images and the correct page formatting. The Replay Server may be unable to access content on the original site, due to proxy configuration issues.

Diagnose

To diagnose this issue, do the following.

1. Replay a session in BBR.
2. Log in as an administrator to the server hosting the Replay Server. Typically, this server is the Portal Server, too.
3. Open a browser window. Navigate to the following:
   ```
   http://localhost:38000
   ```
   Note: If you are connecting from the localhost and if authentication is enabled, you may use ssadmin/ssadmin for the username and password.
4. The Replay Server interface is displayed.
5. Click **Cache**.
6. Review the Cache listing. If you see multiple entries with a byte length of 0 and no content type, then Replay Server is having difficulty communicating with the origin server through the internal Tealeaf proxy.

If the above information is present, then the Replay Server was configured to use WinInet to access the server of origin. This method of connection cannot use a proxy when running with the local system account, which is the default setting for Replay Server.

- WinInet is configured based on the settings that are used by the Internet Explorer instance on the hosting server. These settings cannot be configured for a service running as the local system account.

Fix

To fix this issue, you must configure the Replay Server to use a named account in the domain.

1. **Configure Replay Server:** Log in to the server hosting the Replay Server as an administrator.
2. In the Administrator's Control Panels, open the Services panel.
3. Right-click **Tealeaf Replay Server** and select **Properties**.
4. Click the **Log On** tab.
5. Click the **This Account** radio button.
6. Enter the domain and user name of the named account in the following format:

   ```
   SomeDomain\SomeUser
   ```
7. Enter the password twice.
8. Click **Apply**.
9. Restart the Replay Server service.
10. **Configure Internet Explorer:** Open the instance of Internet Explorer.
11. Select **Tools** > **Options...**.
12. Configure a Proxy server for use in Internet Explorer.

    Note: For more information, the product documentation that is accessible through the installed version of Internet Explorer.
13. Clear the **IE cache** within the browser.
14. Click **OK** to close the **Options** window.

    Note: After you configured Internet Explorer, you must restart Replay Server through the Services control panel so that the cache is cleared.
15. **Verify configuration:** To verify configuration, you may test the cache settings by using the following URL. It is processing a static content URL through the Replay Server on the local machine:

    ```
    http://localhost:38000/GetCacheFile?href=http://www.tealeaf.com/images/
       home/slide-deck-1.jpg
    ```
16. If the above displays an image, then the configuration is working.
17. If not:
    a. Try substituting a different URL after `href=`, which points to a known static content object on your web application.
    b. If that fails, revisit the configuration settings with your IT staff.

The browser window containing the Portal is not opened when you create an event or hit attribute in BBR

Through BBR, you can highlight data and create events or hit attributes through the context menu. These events or hit attributes are drafted for you in the Event Manager and pre-populated with the data to locate the items in a session.

In some cases, using the context menu fails to force open the browser window containing the event or hit attribute that is predefined in the Event Manager. This issue could be caused by one or more of the following items:

- By default, Firefox 4 or later prevents window raising and lowering. To enable, do the following:
  1. Open Firefox.

2. In the menu, select **Tools** > **Options...**.
3. Click the Content tab.
4. Select the Enable JavaScript option. Next, click the **Advanced** button.
5. Select the Raise or lower windows option.
6. Click **OK** twice.

   Note: Raising and lowering windows does not work if the browser window has multiple opened tabs and the Portal tab is not the front one.

- Check the JavaScript settings in your web browser.

Error rendering the page

During replay, the response content may be replaced by the following error message:

```
Error rendering this page! This may be caused by AJAX responses with a
text/html content type, or malformed HTML. Consider adding a replay
rule to remove this page.
```

If the above message appears, then the Replay Server was unable to properly render the page. There could be a number of different reasons for why the page was unable to be displayed. The simplest solution is to remove the page from replay.

To remove a page from replay:

1. In BBR, right-click the currently selected page in the Navigation List.
2. Select **Remove This Page from Replay**.
3. When the session next replays for any Tealeaf user, it is skipped during replay.

For more information about creating replay rules, see "BBR Replay Rules" in the *IBM Tealeaf cxImpact User Manual*.

Chapter 4. Troubleshooting - RealiTea Viewer

Note: As of Release 8.0, the Event Editor was moved to the Tealeaf Portal and is no longer available in RTV. See "Tealeaf Event Manager" in the *IBM Tealeaf Event Manager Manual*.

Note: For browser-related issues with RTV, verify that you installed a supported version of Internet Explorer. See "RealiTea Viewer Overview" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Unable to install RealiTea Viewer

Error - "TeaLeaf RealiTea Viewer requires Microsoft IE V5.0 or greater"

While trying to install the IBM Tealeaf CX RealiTea Viewer, you may receive the following error:

```
TeaLeaf RealiTea Viewer requires Microsoft IE V5.0 or greater.
```

Adding the Registry entry below may fix the problem:

```
[HKEY LOCAL MACHINE\SOFTWARE\Microsoft\Internet Explorer\Version Vector]
"VML"="1.0"
"IE"="6.0000"
```

IE 7 installed

If Internet Explorer 7 was added to the machine after its operating system was originally installed, there are multiple instances of IE on the machine. You must add another entry to the Registry, as follows:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Internet Explorer\Version Vector]
"VML"="1.0"
"xIE"="6.0000"
"IE"="7.0000"
```

After running the Viewer installer, remove the `"IE"="7.0000"` entry from the Registry to avoid causing problems when upgrading to future versions of Internet Explorer.

"Would you like to remove" continues to appear even after RTV uninstalled

If you have uninstalled RTV and are attempting to reinstall, you may receive the "Would you like to remove" dialog, even though RTV components were already removed.

The issue is caused by the presence of an InstallShield installation information folder; Windows Explorer may fail to display the folder, making it easy to miss when trying to manually clean up after the uninstall. To fix:

Note: You must have administrator access to your local machine in order to make this change.

1. On your local machine, navigate to the following directory:

   ```
   C:\Program Files
   ```

2. In the **Windows Explorer** menu, select **Tools** > **Options...**.

3. Click the **View** tab.
4. Underneath the Hidden Files and Folders node, select `Show hidden files and folders`.
5. Click **OK**.
6. In the `Program Files` folder, the `InstallShield Installation Information` folder should now be visible.
7. In each subdirectory of the `InstallShield Installation Information` folder, open and read the contents of the `setup.ini` file. When you find the subdirectory containing the RTV installation information, delete that entire subdirectory.

Improving RTV performance

Here are two areas where you can accelerate performance of RTV. Typical performance problems in RTV concern the loading of external files, such as JavaScript, style sheets, and images.

Set RTV advanced options

1. In the **RTV** menu, select **Tools** > **Options**.
2. Click the **Advanced** tab.
3. Set the following property values:

   Advanced Option
   Value

   Javascript Autopage advanced
   OFF

   Store Images during Replay
   ON

   Use Redirect for external image
   OFF

After the initial page of a session loaded, the above settings force static content to be stored locally. These local files are used in subsequent page loads, which should accelerate the loading of the rest of the pages of the session. See "RealiTea Viewer - Advanced Options Tabs" in the *IBM Tealeaf RealiTea Viewer User Manual*.

- You can configure these settings for user profiles. See "RealiTea Viewer - Profile Options" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Creating external file database

If the previous solution is inadequate, you can create an external file database. A `.tli` file can be used to store static content, which is referenced during replay. Load times for the first page of a session improve with this method.

Note: If changes are made to your site, this solution may cause RTV to reference obsolete version of stored files, which could cause various replay problems. When site changes are made, you should remove the `.tli` file from the configuration, and repeat these steps to create a new one.

1. Apply the settings in the previous solution. See “Set RTV advanced options.”
2. In the Options screen, select the **Static Files** tab.
3. Select the `'Get Images' goes to Static File Database` option.

4. Click **New**. You can either accept the suggested name or pick a preferred file name and location.
5. Click **OK**.
6. External files are now stored as you replay a session. As more pages are replayed, fewer requests are required from the source server.

Images not displaying in RealiTea Viewer

The SERVER_NAME in a REQ must be resolvable to a name that can reach the original Web server from the desktop where the Viewer is running. For example, if SERVER_NAME is `company.com`but the desktop must use

```
http://www.company.com/
```

to reach the website, then the Viewer is not able to get the image.

To fix this issue, you can use the Viewer Profile server remapping option:

1. In RTV, select **View** > **Options** > **Profile** > **Edit Profile**
2. Create a stanza that is named `company.com`. This value should match the string found in SERVER NAME.
3. Under this stanza, create a line:

   ```
   SERVER_NAME=www.company.com
   ```
4. Save the profile changes and exit the options dialog. The replay should redraw, and the images are present.

Every captured Web server must have a stanza in the profile. If there are 5 servers being captured, each of five stanzas must have `SERVER NAME=www.company.com` added to the profile. For example, if the SERVER_NAME in the REQs are `web1`, `web2`, `web3`, etc. then there must be a section for each; `[web1]`, `[web2]`, `[web3]`, etc.

Scripts not replaying in RealiTea Viewer

Error - CallJS: JavaScript function not found: TeaLeaf_Client_tlGetNodeFromXPath

During replay of a session, the following error may be displayed:

```
Error - CallJS: Javascript function not found:
TeaLeaf_Client_tlGetNodeFromXPath
```

This error may be caused by one of the following causes:

1. **Scripts that are not permitted to run in IBM Tealeaf CX RealiTea Viewer:** IBM Tealeaf CX RealiTea Viewer must be configured to replay scripts embedded in a session. Verify the following:
   a. From the **RTV** menu, select **Tools** > **Options...**.
   b. Click the **Replay** tab.
   c. Select the `Allow Scripts to Run` check box.
   d. Click **OK**.
      - See "RealiTea Viewer - Replay Options" in the *IBM Tealeaf RealiTea Viewer User Manual*.
2. **Improperly formatted HTML pages:** In the session, the above error message may appear on only some of the responses. Check the following:
   a. Open the session.

b. Click the **Replay tool** in the toolbar.
c. In the **Viewable Pages** list, select one of the pages where the error occurs.
d. Right-click the replay pane and select **View Document Source...**.
e. The response is displayed. If this page is an HTML page, the page should have the following basic structure. Verify that the following tags are present and listed in the displayed order, ignoring the content between them:

```
<HTML>
<HEAD>
(header content)
</HEAD>
<BODY>
(body content)
</BODY>
</HTML>
```

f. If the above structure is not present in the displayed page, the web application may not be properly constructed the selected page. Please save the source page displayed in `Notepad.exe` to your local computer and provide it to your web application development team.

"No sessions could be loaded from result set" while retrieving Sessions in RTV

When completing a search in RTV for completed sessions, you may receive an error similar to the following:

```
No sessions could be loaded from result set RS_20091124-992233-000 on
server cximpact:19000
```

If you receive the above error, enter the following URL in your browser:

```
http://cximpact:19000/version
```

If the above URL fails to return XML containing the version information of the Tealeaf Search Server (`TLSrchSrv`), then try the following URL:

```
http://cximpact.mycompany.com:19000/version
```

If the above returns XML, then you must provide the fully qualified name of your Search Server as an alias in Search Server configuration. Complete the following steps.

1. Log in to the Portal as an administrator.
2. In the **Portal** menu, select **Tealeaf** > **TMS**.
3. The Tealeaf Management System is displayed.
   - See "Tealeaf Management System" in the *IBM Tealeaf cxImpact Administration Manual*.
4. Select the server that contains the Search Server being accessed by your local instance of RTV.
5. Click the **Search Server node**.
6. Click **Search Server configuration**.
7. In the **Actions** panel, click **View/Edit**.
8. In Search Server configuration, enter the fully qualified name for the Search Server, which may be in the following format:

   ```
   cximpact.mycompany.com
   ```

   - No `http://` is required.
   - No port number (`19000` in examples) is required.

9. Click **OK**.
10. In TMS, configure a task to push this Search Server configuration to other servers.
    - See "TMS Jobs Tab" in the *IBM Tealeaf cxImpact Administration Manual*.
11. When the job completes, re-execute your search to see if results are returned.

Error during Replay of Archived Sessions

For archived sessions, you may see the following error message:

```
The session <session number> in Canister LSSN_* cannot be loaded:
LSSN open failure on <Server name>.
```

1. To fix, you can request the corresponding TLX files from Search Server from the index/Portal machine or the Canister machine, or you can use the Canister session ID to access.
2. If neither of the above works, request the TLC file from the Search Server.
3. If that works, quit IBM Tealeaf CX and verify that the IBM Tealeaf CX Server settings for the two-machine configuration have correct values through the Portal.
4. Try to replay again.

What causes a status code of 0?

HTTP StatusCode=0 is associated with incomplete capture of a hit or page and often with a labeling of the hit as "request canceled" (`ReqCancelled=Client`, `ReqCancelled=Server`, or `ReqCancelled=True`").

Request cancellation as interpreted by the passive capture software can occur in a number of ways:

1. The visitor clicked the **Back** button or clicked a visible link before a page finished rendering.
2. The visitor could explicitly click the **Stop** button in the browser or press the `Esc` key on the keyboard or similar.
3. The actual length of the request or response data may differ from the value that is specified in the HTTP Content-Length header. The client or server closing the TCP/IP connection, miscalculation and misreporting of the content length by the Web server, or complete absence of the Content-Length header can cause this effect.

Viewer incorrectly showing Back button pressed on several pages

The Viewer assumes that the browser's **Back** button was used whenever the **Referer** of the current page is not the URL of the page immediately previous in the session.

1. To turn off **Back** button insertion, select **View** > **Options**.
2. Select the **Replay** tab
3. Clear the **Insert Back Pages** check box.

HTTP headers being displayed in Replay view

1. In RTV, select **View** > **Options**.
2. Click the **Replay** tab.
3. Set the HTTP Header Skip setting to `Auto`.

Some pages seem to be missing when I replay sessions

For sessions that appear to be missing pages, the following are the most likely causes:

- The Viewer may be configured to not replay the page. To resolve this issue:
  - Acquire the full URL of the page, including file extension.
  - Check the Viewer's option for Interpreted Pages.
  - Verify the file type that you want to replay has a check mark by its extension.
  - Review the Viewer's option for Profile to ensure that the URLS of the missing pages are not listed in an IGNOREURL line.
- If you are using a IBM Tealeaf CX Passive Capture Application server (PCA server):
  - There may be a Web server that the PCA server was not configured to capture on the Interface page of the PCA servers Web UI.
  - The missing pages are of Binary file type. To determine:
    - Create a session for yourself. On the pages that are missing from replay, what is the file extension of the pages? Does it match an extension found in the ExcludeExtensions list on the Miscellaneous page of the PCA servers Web UI?
    - If so, remove the extension from the ExcludeExtensions list and click **Save Changes** (preferably when Web traffic is low to minimize the disruption to data capture).
  - There may be an issue with the input data to the PCA servers capture NIC(s).
    - If you can define a session-level event that recognizes the absence of the page(s), then you can do a focused TCPDump to determine whether there is such an issue. Start a TCPDump run (output to a file) at a time of day when this issue is most likely to occur. If you can filter the traffic by Web server IP address or at least IP port number, that can help keep the size of the output file from growing too quickly. See “Before you begin” on page 5 for instructions on using TCPDump to record network traffic to a file.
    - After making the above changes, observe the Portal for the occurrence of the missing-page event. After the event was seen, the TCPDump run can be stopped. You can then retrieve the session that fired the event, get the REMOTE ADDR (client IP address) from it, and filter the TCPDump output file to create a much smaller file for analysis of whether the page is missing because the input data to the PCA server is missing some data (see “Before you begin” on page 5 for instructions).

This procedure is more complicated if the Web traffic of interest is HTTPS. In that case, you must verify that the beginning of the captured session that occurred after the beginning of the TCPDump run to ensure that the TCPDump output contains the initial SSL handshake. Capture of this handshake data is crucial to decrypting the captured data for analysis. If using IIS capture filter:

- There is a Web server whose capture filter is not installed/not working/can't connect/has insufficient bandwidth to Tealeaf server/experiences intermittent network faults to the Tealeaf server. In this case, the solution is to review the Windows Application Event Log of each Web server to ensure the filter is running correctly on each.
- The missing pages are of the Binary file type. To determine if this is the case, manually drive a session. On the pages that are missing from replay, what is the file extension of the pages? Does it match an extension found in the `TealeafIIS.cfg` file's ExcludeExtensions list? If so, remove the extension from

the ExcludeExtensions list in the `.cfg` file on each Web server and restart IIS to reload the `.cfg` file (preferably when Web traffic is very low).

HTTP status code 304 and cached objects

As part of a GET request, the browser can include a header that is called `If-Modified-Since`.

- The Request view of a Tealeaf-captured hit displays it as `HTTP IF MODIFIED SINCE`.
  Along with this header is a date indicating when the object was cached. It is up to the Web server to look for this header and compare the date in it to the last-modified date of the requested resource. Most Web servers support this option.

If the Web server supports this option, it can return an HTTP status code 304, to tell the browser to use its local cached copy. This interaction between Web browser and server still results in a REQ/RSP pair with a zero-length response body. This sequence of events is different from hitting the browser's Back button and getting a page from local memory cache with no REQ being issued to the Web server.

The browser only includes the `If-Modified-Since` header in its request if it finds a copy of the object in its local cache, so this behavior should only apply to files that can be cached.

The cache is not necessarily the in-memory cache, although it is entirely up to each browser how it is handled. In the case of Internet Explorer, the cached files are stored as Temporary Internet Files on the local disk and remain until you explicitly delete them in IE's Internet Options. So, quitting out a browser and reopening does not avoid 304's.

If you have selected the `Load Remote 304 Pages` check box in the Replay options tab, it runs a simple GET request that does not include the `If-Modified-Since` header, so that the Web server returns the object rather than a 304 response. However, the Viewer may not retrieve the same object that the original visitor saw. Any redirection that occurs as a result, such as to the home page, is not immediately obvious when replaying a session, which is why that option is usually disabled.

- See "RealiTea Viewer - Replay Options" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Why are the .css files missing from my sessions?

Cascading Style Sheets (`.css`) files are considered static content. Depending on how your Tealeaf environment is configured, some or all of the following file types may be dropped from capture:

```
.au, .avi, .bin, .bmp, .cab, .class, .css, .dcr, .doc,
.exe, .gif, .gz, .htc, .htrc, .jar, .jpeg, .jpg, .js,
.mov, .mp3, .mp4, .mpe, .mpg, .pdf, .png, .ppt, .ra, .ram,
.rar, .rm, .rtf, .snd, .swf, .tif, .tiff, .wav, .xls,
.zip, .ico
```

The above list of file extensions identifies many common types of binary files, which may be managed through the following mechanisms:

- **PCA:** Through the Pipeline tab in the PCA Web Console, you can configure the following methods of managing binary content:

  - Captured files whose extensions are present in the Excluded Extensions list are automatically dropped by the PCA. See "PCA Web Console - Pipeline Tab" in the *IBM Tealeaf Passive Capture Application Manual*.
  - Beginning in PCA Build 3502, you can configure the PCA to drop responses that are identified as images. See "PCA Web Console - Pipeline Tab" in the *IBM Tealeaf Passive Capture Application Manual*.
- **Windows Pipeline:**
  - When enabled in the Data Drop session agent, DelImages discards image data in the Windows pipeline, so that repeating static content is not stored in the session. See "Data Drop Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
  - Optionally, you can configure a static archive to capture and store static content, so that you retain a permanent snapshot of captured sessions.

    Note: When a static archive is deployed, DelImages must be disabled. See "Managing Static Archives" in the *IBM Tealeaf cxImpact Administration Manual*.

Depending on how your Tealeaf environment is configured, static content, including style sheets (`.css` files) may be dropped during processing. When a file extension was configured to be dropped, Tealeaf discards the response.

- If the request has an HTTP status code of 200, the request is dropped. As a result, requests for static content do not appear in the list of pages of a session unless there was an error fulfilling the request.

Each page may make multiple requests for these files, so storing every successful hit of these types can waste space in the database and the indexes. These static files rarely change. Unless there is a specific error, it is more sensible to assume that they were successfully delivered.

Error during replay of Static HTML pages

`Missing anchor attributes in the <A> link HTML tag`

Anchor attributes that link to items on the same HTML page are not downloaded with the pages in the session. These items are a client-side tag feature, so you may see several highlighted links to the same page.

Replay issues with applications built for leading versions of Internet Explorer

Some web applications may be created to leverage leading-edge features that are introduced in the latest versions of Internet Explorer. Applications such as RTV that use an embedded version of the IE web browser control require a specific registry setting to enable use of these features.

Note: Before making changes, you should verify that you have the latest version of RTV compatible with your IBM Tealeaf cxImpact solution installed on your desktop. For more information on downloading Tealeaf software, see Tealeaf CX Downloads.

By default, the web browser control in RTV reverts to IE 7 compatibility mode, which may result in quirks during replay. Symptoms may include:

- Text size differences
- Form alignment

- CSS oddities

For compatibility purposes, the web browser control runs in IE7 Standards Mode by default if no Feature Control Key value is defined for an application.

To ensure proper replay, you must configure the IE web browser control that is used by RTV to default to the appropriate IE browser version. To run the RTV web browser control in IE8 Standards Mode, add the following new value into the registry:

```
[(HKEY_CURRENT_USER or HKEY_LOCAL_MACHINE)\Software\Microsoft\
   Internet Explorer\Main\FeatureControl\FEATURE_BROWSER_EMULATION]
    "RealiTeaPro.exe" = dword 8000 (Hex: 0x1F40)
```

See http://msdn.microsoft.com/en-us/library/ee330730%28v=vs.85%29.aspx#browser_emulation.

Proxy Issues with IE6

If you are connecting through a proxy server to the web application using IE6, RTV may encounter replay issues. During replay, these issues might appear as Status Code 404 error pages that are displayed in the replay window. Response data may be accurate, and image download appears to work but does not actually capture the content.

RTV uses an embedded version of the Internet Explorer control for the version of IE that is installed on your local computer. RTV creates an internal HTTP server for use with replay. If the IE control is not configured to bypass the proxy server and to use the local server instead, RTV may not be able to replay the content.

IE7 or later is automatically configured to perform this bypass. In IE6, you must manually configure Internet Explorer to bypass the proxy in favor of the local HTTP server addresses.

1. On your local system, open Internet Explorer 6.
2. In the **IE** menu, select **Tools** > **Options...**.
3. Click the **Connections** tab.
4. Click **LAN Settings...**.
5. In the Proxy Server pane:
   a. Click the `User a proxy server for your LAN` check box.
   b. Click the `Bypass proxy server for local addresses` check box.
6. Click **OK**.
7. Click **OK**.
8. Close IE6.
9. Restart RTV.

Viewer says IndexWizard.dll is incompatible

An RTV dialog displays the following:

```
IndexWizard.dll is incompatible, Viewer cannot find .tlx file.
```

1. Open a command prompt window in the RTV Viewer install directory.
2. Run the following command:

   ```
   regsvr32 IndexWizard.dll
   ```

Embedded .pdf documents opening in separate window

In the Internet preferences of Adobe Reader 8, there is a setting to open the `.pdf` within the browser window. This setting should be enabled.

Event icons unavailable

If the Event icons are not available, do one of the following configurations depending on your Tealeaf deployment.

All-in-One Tealeaf or stand-alone Portal Server

1. Run RegEdit on the Portal server.
2. Navigate to the following:
   `HKEY LOCAL MACHINE\ SOFTWARE\TeaLeaf Technology\DataStore\SearchServer`
3. The EventImagesPath should be the full path of your `...\TeaLeaf\Portal\WebApp` directory.

Multi-Server deployment

1. Log in to the Portal as an administrator.
2. From the **Portal** menu, select **Tealeaf** > **TMS**.
3. In the **WorldView** tab, select a server that hosts a Canister (Processing Server).
   a. Open the Search Server node.
   b. Click the **Search Server configuration** node.
   c. Click **View/Edit**.
   d. Enter the appropriate value for the Portal Server.
   e. Save your changes.
      - See "TMS WorldView Tab" in the *IBM Tealeaf cxImpact Administration Manual*.
4. If there are multiple Canisters in the environment, repeat the above TMS steps for each Canister.
5. Configure a job to push the changes to the Canister server or servers.
   - See "TMS Jobs Tab" in the *IBM Tealeaf cxImpact Administration Manual*.

Replay rules

IgnoreURL rule causes event list to be merged with page list

If you have created a new IgnoreURL rule, you may notice that the event list merged with the page list. This happens infrequently.

The issue can be corrected by completing the following steps:

1. If you have unsaved changes in your session, save it.
2. Close the session.
3. Reopen it.
4. The panes are separated, as normal.

Consistent crashes of the RealiTea Viewer (RTV)

RTV downloads all hits for a session when doing replay. When doing searches, it must download the specific hits that match the search.

If you open RTV in the morning and spend hours searching and replaying sessions, RTV consumes large amounts of disk space in its temporary directory. After hours of use, the RTV application window may suddenly disappear. If you are experiencing this problem:

- Try to make your searches more restrictive. The benefits are twofold:
  - Searches finish faster
  - The result sets consume less memory on the computer running RTV
- After searching, wait for the hit results to be displayed before starting replay. Click **Cancel** in the upper right of the RTV search results screen while RTV is downloading individual hits to jump straight to replay.
- Take a break every hour or so or after replaying a score of sessions to close RTV and reopen it. This restart allows RTV to free up space from its memory and disk caches.

POST data matching fails for DWR POSTs

If you upgraded to Release 8.4 and your site uses DWR to transfer messages to the web server, you must apply a patch and make a change in your RTV replay profile in order for POSTs to be properly matched.

- See "RealiTea Viewer Overview" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Using RTV across multiple versions of CX

In some customer installations, multiple versions of IBM Tealeaf cxImpact are available. How can RTV be configured to access each version?

- Known Version(s): 46xx and later

The problematic area of RTV is the Event Editor. You must make available to the latest version of RTV the Event Editor DLL for each version of IBM Tealeaf cxImpact in use. In the following example, there are two versions of IBM Tealeaf CX in use: 6050 and 7075.

1. Install a full version of the latest release of RTV.
2. From Release 6050, find the `EventEd.dll` file in the Tealeaf Install Directory.
3. Rename this file to `EventEd6050.dll`.
4. Copy this file into the RTV directory for Release 7075.
5. To edit events on the 6050 canister:
   a. Close RTV.
   b. From the **Windows Start** menu, select **Run**.
   c. Run the following command:
   `regsvr32 EventEd6050.dll`.
   d. Reopen RTV.
6. To edit events on the 70755 canister:
   a. Close RTV.
   b. From the **Windows Start** menu, select **Run**.
   c. Run the following command:
   `regsvr32 EventEd.dll`.
   d. Reopen RTV.

Troubleshooting - Replay cookbook

This section outlines a replay cookbook of preliminary steps that you can follow to diagnose common issues with replaying Tealeaf sessions through IBM Tealeaf CX RealiTea Viewer. Outlined as a sequence of steps, the sections below contain several examples to help to alleviate common problems associated with the specified replay situation.

Overview of Replay

During replay, RTV loads a version of the target web page into an instance of Internet Explorer over which it has control. When objects required by the target page are requested, RTV intercepts the request in the IE instance and decides if it has the requested object in the loaded session or if it must get it from the original server, the web application.

If RTV can locate the object in the session data, it creates a version of it and requests it through a proxy server that it uses to deliver the requested object as if it came from an external web server.

- When you look at the headers of the request, these proxied requests typically originate from `localhost:<some_port_number>`.

Images, JavaScript, and CSS are typically not captured and stored in the session data, as these items do not change often and require significant storage overhead. During replay, these items are commonly requested from the origin server. Other items such HTML, scripts that are created based on customer data, JSON, and binary data are usually captured in session data; if they were to be requested live at replay time, they may have changed or are unavailable.

Note: Any data that **must** be the same during capture and during replay must be captured into session data.

For more information about troubleshooting common problems including the Replay Trouble Tables, see "Next steps" on page 108.

Check 1. Verify everything you need is in the session

In the Page Load Details page in RTV, you can see each object that the web page is requesting from the session and the origin server.

- If any item is, Page Load Details may give insight into the elements that are not replaying correctly.
- To access Page Load Details, select **View** > **Page Load Details**.

Many seemingly complex replay issues can be rectified by verifying that all required data is in the session data or accessible from the server. Often, a UI element that is missing or fails to replay is corrected as soon as a POST or JSON request missing from the session became available.

Examples

- If JavaScript or CSS is not retrieved successfully during replay, the site may not replay correctly or may render poorly.
  - These types of errors are easily corrected.
- Some files are created based on user credentials, like cookies or login accounts.

  - If these files are not captured in the session data and are requested during replay, they usually do not complete successfully. In most cases, a generic version of the file is returned from the origin server, or a redirect occurs during the request, returning another unassociated page, such as the home page. This issue results in erratic replay.
- If POST data are required to dynamically affect UI elements of the web page, this data must be captured in the session data.
  - Common examples include drop-down lists that auto-populate based on previous choices or radio button groups toggled based on user input.
  - Usually during replay, a request is issued when a UI element is changed. This request is usually a POST request that instructs the server to send the data that is required for the page. This request and the returned data must be in the session for replay to accurately portray what happened during capture.
  - UI Capture may need to be implemented in order to capture these data.
- There are certain cases where JavaScript must be captured in the session data for replay to work correctly.
  - Good examples are ashx files for .NET websites, as these contain script code that is needed to control UI elements of the page.
  - If these data are being dynamically created on the server during the user's interaction with the web page, replay must reference versions identical to the version used during capture. In this instance, the custom-built JavaScript must be captured in the session data to replay correctly.
- When in doubt, check Page Load Details!

Replay for Sites Requiring Sign-In

If the customer has most of their static content, such as images, script, or objects, blocked from users who did not logged in, RTV is likely to have issues retrieving this content during replay.

Options

- You can acquire the content and place it on a proxy server, such as the Tealeaf Portal server, which can store the content so it may be served during replay. This scenario requires a host/port remap rule to map the URL's for static content to your proxy.
- You can enable Business IT Mode in the PCA so that all images and static content are captured.
  - You must remove blacklisted extensions and MIME types and add the file name extensions to the whitelist (like `gif`, `jpg`, etc.).
  - You must also verify that the Processing Servers have `DelImages` disabled in the DataDrop pipeline agent so that the content is not dropped. See "Data Drop Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

    Note: Beginning in PCA Build 3502, this functionality can be enabled in the PCA pipeline, which limits the volume of data that is processed and transmitted to the Processing Servers. See "PCA Web Console - Pipeline Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

    Note: This configuration captures all content in the session. Operating in this mode dramatically increases the storage requirement and should only be used during POC or other temporary situations where Replay must be demonstrated. This solution is for temporary use only.

Check 2. Verify needed data are not being blocked by Privacy

Privacy is used to obscure sensitive customer data. While privacy secures customer data, it can result in some data needed for replay being blocked.

A common rule that causes replay issues is `BlockURLFields`. This rule blocks all the parameters that are appended to the end of URL's following a `?` character. If during replay you notice many of the URL's in the Navigation list (upper left pane) have lists of `X` characters following them, you may have privacy settings issues.

- If the URL's in the Navigation list have any repeated character following them, privacy is usually the culprit.

The obscured URL parameters are needed by RTV to match the requests with the data in the session. If the requests cannot be matched, then they are requested from the origin server. The returned hits are not identical to the originally captured hits, which results in poor replay.

Examples

- When the main page request for the session has URL parameters, obscuring them can force RTV to request them from the origin server, which results in poor replay and artificially inflation of the hit count for the page.
- If `.axd` pages in .NET server sessions are blocked by privacy, crucial parts of the web page that is required for reacting to JavaScript events are not available.
  - Because these pages commonly have the same name with different content (e.g. `WebResource.axd?xxxxxxx`) for each request, the URL parameters are required by RTV to decide which version of `WebResource.axd` in the session is needed for the page.
  - There are several types of files where this can happen. Any file containing specific JavaScript for UI controls must be available during replay through the session data.
- If POST data are blocked by privacy or the POST request has URL parameters that are blocked by privacy, those POST requests are not matched correctly, and UI elements fail to replay.
- Replay is impeded for any URL if the parameters used to make that page unique were blocked by privacy. At one customer site, dynamically generated GIF images of application forms were requested using a `document` and `page` parameter. When blocked, these parameters caused RTV to fail to find the image. The image failed to display, even if it existed in the session.

Check 3. Client UI Capture for Ajax is in use and its events are captured

Tealeaf offers the IBM Tealeaf CX UI Capture for AJAX to be deployed to the customer website and served to the visitors' browsers when they navigate the website. This JavaScript hooks in to the UI elements on the client's web page and sends events back to the Tealeaf system when the visitor interacts with the UI elements. The events are converted to XML data and are POSTed back to the website using Ajax.

- See "UI Capture for Ajax Guide" in the *IBM Tealeaf UI Capture for Ajax Guide*.

After the UI events are captured, they are saved in the session to be used by RTV during replay. The events come in bundles of XML and can be seen in the REQ

view of the session in RTV. Using timestamps and other parameters, RTV sorts these events and associates them with the pages in the session.

Because the events were captured while the user interacted with the web application, they can replay how the user interacted with each UI element on the web page and in the order of interaction. For replay of this type to work correctly, however, events must be correctly matched with the parent page, and the UI elements to which they pertain must be available at replay time in the UI of the page.

Note: Replay is sequential. If an event does not replay correctly, you can expect the next set of replay events to fail as well.

For more information about including and setting up the library, see "UI Capture for Ajax Guide" in the *IBM Tealeaf UI Capture for Ajax Guide*.

Examples

- A web page contains a form and is referencing the Tealeaf IBM Tealeaf CX UI Capture for AJAX. The form is either shown or hidden based on the values of a set of radio buttons. If replay of the radio button event fails, the form is not displayed during replay.
  - In this case, all replay events to fill the form fields following that radio button click fail to replay as well, since the form does not exist in the UI. This common situation has many different variations.
- A web page has a set of drop-down lists. The choice from list 1 affects the content of list 2. For example, when you pick `Honda` from list 1, the page requests the models `Honda` manufactures to fill list 2. If the request from list 1 fails due to privacy blocking needed data, the request resulting in an HTTP 404 error, or the event for list 1 failing to fire, then list 2 fails to replay. For example, if there is a list 3 dependent on list 2, events should fail for it as well. In this case, follow Checks 1 and 2 to verify the UI event data are captured in the session and that the requests are being matched up correctly (for example, no privacy blocking).
- There are several events that say `xpath` in front of them, and they have a long list of members and may have no "value" data. These events do not do anything during replay, and all subsequent UI events don't do anything.
  - In this case, the `XPATH` event is missing a crucial piece of information that most likely cannot be matched to the content on the page, and RTV cannot send the event to the UI element.

    Note: This situation is complicated. Contact Tealeaf Engineering for assistance.

Check 4. Verify only one main page is requested for each page in the Navigation list

A common issue regarding replay occurs when the RTV user clicks a page in the Navigation list. The page displays and then disappears, replaced by an HTTP 404 or another, incorrect page.

In this scenario, RTV ran JavaScript or HTML when it highlights the elements of the UI during replay. RTV can be configured to highlight fields and UI elements during replay to denote the element clicked by the visitor during their visit to the web application.

In RTV the replay highlighting option allows scripts associated with the currently highlighted UI element to run and to carry out all actions that are specified in the script that is intended for the origin website. This option is not always needed if the origin site is simple and does not toggle elements through UI controls. However, if the site uses complex JavaScript behaviors to control UI elements, these script actions must be run during replay.

- In RTV, select **Tools** > **Options** > **Replay tab**. The options is labeled, `Invoke JavaScript while highlighting`.

Examples

- If the parent site has a **Next** button to navigate away from the current page, this button most likely causes a navigation when RTV highlights it. You must remove this anchor tag or JavaScript so that it cannot cause navigation.
  - See Chapter 5, “Replay Trouble Tables,” on page 109.
- There is an overlay page that makes the rest of the screen inactive and displays a form with which you can interact.

Check 5. Verify frame content is configured correctly

In websites that use frames, replay can require some work to display the pages properly. A page in the session is likely to be the parent frame and subsequent pages after it in the Navigation list (upper left list) are the child frames that must be placed within the parent.

To allow proper replay of frames, you must enable the following RTV features:

- **Tools** > **Options** > **Replay** > **Replay Frames** (required)
- **Tools** > **Options** > **Advanced** > **Replay Remote Frames** (optional)
- **Tools** > **Options** > **Advanced** > **Aggressive Frame Placement** (optional)

Then, you must tell RTV the hierarchy of the frameset:

1. Find the page that you expect to be the parent page. Click this page in the Navigation list and click the **RSP** button to examine the HTML code of the response, which should contain a `<FRAME>` tag. Inside this tag, the URL's of the subsequent pages in the Navigation list should appear.
   - After they are found, you can right-click the parent page in the Navigation list and select **Replay Rules** > **Place this page in a frame** > **_top**. The parent page is configured to hold all subsequent child pages.
2. Locate each page that must be placed in the parent page from the Nav list. Right-click each page and select **Replay Rules** > **place this page in a frame** and then select the frame ID that pertains to this element.
   - Usually, the `src` parameter that is located in the `FRAME` element of the parent page contains the URL of the child element to be assigned. Then, use the `name` or `ID` parameter of the `FRAME` element in the parent page to assign the child.

Next steps

If you are still having challenges, see Chapter 5, “Replay Trouble Tables,” on page 109.

Chapter 5. Replay Trouble Tables

This page is your entry point to finding solutions to issues affecting replay of sessions through RTV.

To use this page

- Review the scenarios in the Links below. They begin with S-##.
- When you find the applicable scenario, step through the questions.

  Note: If possible, review these pages while connecting to the Internet. Solutions on this page may point to solutions listed elsewhere in the Replay Cookbook or Tealeaf Online Help.

S-01. The page in replay does not display or is not the correct page

Q-01. Have you checked page load details to see if the page is being requested correctly?

NO:

- Go to RTV > View > Page Load Details.
- Look at all the requests and see if the main page you are expecting to see is resulting in a 404 or a status code

other than 200. Verify that the page is being returned from the session and NOT the remote site. In the event it is coming from the remote site, it was not properly caught in the session data. Verify that the PCA is set up to capture pages of the type you are attempting to replay. In the event it is returning some code other than 200, it is almost certain it was not captured correctly. If the page results in another page being served, often times it is a redirect after the remote server could not serve the page. This once again points to the session not being captured correctly.

S-02. Element does not show up in the webpage during replay

Q-01. Do any of the pages in the Navigation list (upper left pane) have lists of a single character after them (such as `XXXXX` or `@@@@`)

Yes:

Privacy is blocking URL field data that may be needed to match the requests with the UI elements. Alter the privacy settings to allow URL fields to come through. Usually the culprit is the rule `BlockURLFields` in the Windows pipeline on the canister server.

Q-02. Are there any requests resulting in 404 status code OTHER than images and CSS?

Yes:

If the missing files have an extension like `AXD`, ASHX, `ASPX` or otherwise, they contain JavaScript needed to display elements on the page. They were not captured in the session and RTV is making a request to the origin server for them and failing. On the PCA, make sure the extension of the pages that show 404 are white-listed. Also verify they are not in the "Block" list.

Note: In the event the page has NO extension, the PCA decides to keep or throw away pages based on the MIME type contained in Content-Type in the header. Verify that the MIME type for this missing page is white-listed and not blocked as well.

Q-03. Are there multiple requests to the home page or multiple redirects in page load details?

Yes:

Most likely, pages that are needed for replay were not captured and are being requested from the origin server during replay. Because the server needs cookies or a valid session id to serve these pages, it is simply returning a redirect to the home page or other generic page. On the PCA, verify that the page originally requested has an extension that is white-listed and not blocked.

Note: In the event the page has NO extension, the PCA decides to keep or throw away pages that are based on the MIME type that is contained in Content-Type in the

header. Verify that the MIME type for this missing page is white-listed and not blocked as well.

Q-04.Are there loading graphics or a large section of the page is blank?

Yes:

See S-05.

S-03. A UI element is present on the page but does not populate with data. Subsequent replay does not work.

Q-01. In page load details, are there any 404 requests OTHER than images and CSS?

Yes:

See S-02, Q-02.

Q-02. In page load details, are there any requests being redirected or returning empty data?

Yes:

See S-02, Q-02.

Q-03. In page load details, are there any requests completing with a 200 that are of type POST AND are being requested form origin server?

Yes:

Save the session (TLS) as a TLA, and do a text search inside the session to see if the missing URL is in the session file.

Q-03.NO. Is the URL present in the TLA?

NO:

The POST request was not captured in the session data. This can happen because the MIME type or the extension is not white-listed or is blocked in the POST section of the PCA. Add this extension or unblock it if blocked. Go ahead and add or unblock it from the MIME type section as well.

Yes:

Most likely there is some issue with matching the captured POST request in the session during replay. If the URL has URL encoded values (like _, or %36) in it, go to RTV > Tools > Options > Advanced and clear Strict Post Data Matching. This does not do a simple string comparison when matching POST requests.

Note: As of Release 8.3, this setting is removed. It may be an issue of the post data matching plug-in. See "Managing POST Data Matching plug-ins" in the *IBM Tealeaf CX Configuration Manual*.

S-04. A page shows during replay, but then is immediately replaced by a 404 or some other page that clearly isn't the right one.

Q-01. In the page load details, do you see a request for more than one MAIN page (page with an "asp, jsp, html, htm, php etc...extension)?

Q-01.YES Does the initial page have a submit button or other page transition button on it (like "next")?

Yes:

Most likely the highlighting used when replaying is causing the button to navigate. If you don't require RTV to run JavaScript when it highlights (like there are no dynamically shown/hidden elements) then you can turn off `Invoke JavaScript when highlighting` in RTV > Tools > Options > Replay.

If you need the JavaScript to run while highlighting for other elements to replay, then you can write a response mod rule to replace the JavaScript that fires that button. Running a subsearch for the URL that is being requested is usually sufficient. Then replace the URL with the '#' character (anchor).

NO:

It may be that there is a javascript that is attempting to load a new page. Searching for phrases like "window.location.href " or "window.navigate" or "document.location" or "window.location" can help you pinpoint what code is doing the navigation.

In the event you don't find these, it is possible that the code to do highlighting is simply an anchor tag "<a href=". Because there will be tons of these, you need to search for the last "main" page url you see in the page load details. Once you see the url that is causing the problem (the one that is causing the 404 or is the wrong page), you can search for this url and hopefully the anchor tag that is associated with it.

In either of the above cases, a response mod rule requires to replace the code or anchor tag so it does not fire when highlighted.

S-05. A portion of the window content is missing (loading graphic(s) or obvious blank area on page).

If this setting is available, verify that the following was selected: **RTV** > **Tools** > **Options** > **Advanced** > **Two Phase Replay**.

Q-01. Are there multiple subsequent pages in the navigation list (upper left list) that have unformatted HTML or weird looking data in them?

Q-01. YES Does the content in the following pages look like it should be the content in the area with the loading graphic or blank area?

Yes:

The content is not being matched correctly but is in the session. This usually means the URL in the parent page (the one with the blank area or loading

graphic) is using URLs with a host included. This causes a cross site scripting issue where this content is not loaded. RTV loads its content through a proxy server located at localhost:190X, this is expected to be the host and when the replayed page requests something from a different host (the host of the original site) security in the browser blocks the request.

To fix this, you must search for URLs that have the host that is contained in them. An example would be URLs that look like "www.thecustomersite.com/interesting_content.asp. "
The URL matches one of the subsequent pages that contain the unformatted HTML mentioned above.

A response mod rule should be added to change this URL to "/interesting_content.asp". This relative URL should allow RTV to place the content in the missing area.

Q-02. In page load details (see S-01) are there any requests of type POST that are returning 404?

Yes:

The content to fill the missing area is being requested by the page to the origin server because it may not be in the session or is not being matched correctly. See S-02, Q-02.

Q-03. In page load details (see S-01), does everything complete correctly (no 404's) and, you can see the content that is supposed to be in the blank area but it still shows up as another page (in navigation list)?

Yes:

Try turning off RTV->Tools->Options->Advanced->Strict Post Data Matching. This option does an exact string compare for the content. This may cause it to fail the comparison and the content is placed incorrectly.

Note: As of Release 8.3, this setting is removed. It may be an issue of the post data matching plug-in. See "Managing POST Data Matching plug-ins" in the *IBM Tealeaf CX Configuration Manual*.

Q-04. Does a rectangular area (possibly blank or with a message) cover a portion of the page?

Yes:

This may be a frame placement problem. The problem may appear more frequently when navigating backward through the session.

In this case, to correct the problem, try clearing `Aggressive Frame Placement` in the Advanced Options tab in RTV.

Q-05. When you look at the RSP (hit the RSP button), does it look like there is a whole page there (meaning opening and closing html tags, body tags, and content) that looks like valid HTML?

Yes:

In some cases, the content of the page is hidden using CSS or JavaScript and when the "loading" portion is finished, the page is shown. However, in RTV, the page is already there (no Ajax request to complete) and so we want it to show. An example would be a session where there is a DIV that is shown on load of the page. You can write a response mod rule to replace the JS code to allow it to show. This requires some slogging through the customer JavaScript.

Customer example:

The body tag looked like this:

```
<body
onLoad="javascript:onloadExpressICMS();
   javascript:setOnLoadFlag();
   javascript:hideProgres();
   javascript:validateAllForOnLoad('A');
   javascript:myFormLoad();
   javascript:setFocus('agentRecord1');
   javascript:releaseOnLoadFlag();
   javascript:controlInsuredInformationBlock
      ('F','false');
   javascript:handleIcmsCookie('1250291211465');
   javascript:displayLinks();">
```

This tag contained a call to "hideProgress()" that I then looked for (because there was only a progress bar displaying during replay of the page). I realized it might help to call that function so I wrote this replay rule:

```
<HostProfile name="eagent.farmersinsurance.com" id="31">
          <ResponseModify id="34"
              url="/PLA/eAgent/eAutoE/view/info/
                   premiumsummary/
                   premiumSummaryCancelAll"
              pattern="onLoad=&quot;"
              replacementString="onLoad=&quot;
                  hideProgres()&quot;
                  onLoaded=&quot;" occurrences="all"/>
     </HostProfile>
```

This rule converted the tag to the following:

```
<body onLoad="javascript:hideProgres();">
```

The page now shows by calling the hideProgress() function first.

S-06. I keep getting popups that mention javascript errors.

Note: If the popup asks if you want to continue running scripts, NEVER hit "no". This disables the scripting engine in RTV and it does not run any scripts until you restart RTV.

Q-01. Do the popups ask if you want to debug?

Yes:

Try turning off the script debugger. RTV->Tools->Options->Advanced->Disable script debugger. This should block the popups that ask you to debug.

If this still does not block them, Open IE->Tools->internet options-> advanced and check "disable script debugging (Internet Explorer)" and "disable script debugging (other)"

If neither of these help, see the resolutions under S-06, Q-02.

Q-02. Do the popups ask if you want to continue running scripts?

Q-01. YES Does the popup mention "access denied" AND the current page is using HTTPS?

Yes:

This is often caused by javascript trying to set or change a variable that is on a secure page when RTV is using localhost over a non-secure connection to serve the files. You must find the offending line (usually in a separate file) and make sure the request cannot occur.

See the next topic as well even if HTTPS.

Q-02. YES Does the popup mention "access denied"?

YES:

This is often caused by javascript trying to set the domain through "document.domain = ". You should do a subsearch for this text to see if this is the issue. If it is not found, you should use RTV->Tools->Get Images and then do another subsearch. This gets dynamically requested/generated pages that are not in the session and search them. When you find it, write a response mod rule to replace this text with "//document.domain=", this comments out that line of javascript.

If this does not work, it is also possible there is some other javascript member that is domain-specific and is not allowing RTV to set it during replay. At this point, you must look for `document.`, `window.` and investigate. You will must toy around with response mod rules to modify this offending line.

Example:

At one customer, they had a review partner named Bazaarvoice loading in an IFrame and getting content from `reviews.epson.com`. Because RTV was loading everything from localhost:1901 (our internal proxy), there was a cross site scripting issue. To fix it, I wrote this replay rule:

```
<ResponseModify id="159" url="" pattern=
    "reviews.epson.co.uk"
    replacementString="localhost:1901" occurrences="all"/>
```

This rule allowed the IFrame URL to be changed to the localhost:1901 that we use, eliminating the cross site scripting.

Q-03. YES Does the popup mention "'x' is undefined"?

YES:

This is often caused by javascript requiring an object that was not collected in the session, or cannot be collected from the origin server at replay time.

This can be because a javascript file did not load successfully. Check page load details and see if any JS files are returning a 404.

- If the file returning 404 is being requested from the origin server, then it means it is either password that is protected and cannot be retrieved without a session, or it is behind a firewall or some other blocking network condition. Either way, this file should be collected during capture. See S-02, Q-02.
- If there are no 404's, then it means something is different between the way the page runs during replay and the way if runs in the wild. It is possible that some javascript is being successfully retrieved from the origin server during replay (200 status) but that its content is different than when it was requested during capture. If there are any redirects or pages in the page load details with empty content length, this is most likely the reason. Make sure these pages are collected during capture See S-02, Q-02.
- If the object it complains about seems to be a Tealeaf related variable, it is likely the TeaLeaf.js file was not included by the customer correctly. Verify that it was included in a static way that can be cached. Usually this means at the top of the main page and not in any subsequent included requests. If it is included in ANY dynamically generated page, this is a serious issue and must be resolved. Consult the UI Capture library installation manual for details as to how the script must be deployed.
- If none of the previous issues apply, it means that some javascript variable is causing the script to fail on replay. To track the issue down, do a subsearch for the variable the popup mentions, and investigate.

S-07. RTV replay is significantly slower than the live site

Q-01. Are you blocking any URL's by replacing part of them (that is, changing "my.site.com" to "blocked")?

Yes:

We have seen internal proxy servers have issues with requests going to unknown host names or URL's. Usually the issues manifest with very slow performance of all content. The way to fix this is to go to the "hosts file" of IIS and add the new verbiage to the file as the same entry as localhost. An example would be to add "blocked" as above into the hosts file as the same entry as localhost.

NO:

It is possible that the internal proxy or other software (anti-virus) does not recognize the RTV User Agent. This means that when RTV requests something, unknown User

Agent strings (i.e. not Internet Explorer or Firefox, and so on) can be blocked from requesting content or experience poor performance. The customer must make sure that their internal infrastructure knows and allows RTV to request content for replay to be reliable. This issue can also result in images, CSS, and other portions of the site to be blocked when RTV requests them, resulting in poor replay.

S-08. Some client UI events do not replay

Q-01. Do UI events replay up to a certain point?

Q-01. YES Do you see any UI events of type "Exception"?

Yes:

Validate that your TeaLeafTarget page is set up correctly. Hit the REQ button in RTV and inspect the REQ buffer of the Exception type event. If the event contains the "FailedUrl" parameter, and it contains `TeaLeafTarget.asp`included in the URL, it is misconfigured. This does not apply to Client Side Capture.

If the ResponseType=unknown, this is also an indication that the target page is misconfigured.

Q-02. YES As you click through, do events that look like they should do something fail to affect the page?

Yes:

If replay works up until a certain point and then subsequent events do not modify the page or highlight, it usually means some content is missing. Events are based on ID's associated with form elements on the page such as fields or buttons. If these UI elements do not exist or the IDs are not unique or do not exist, replay cannot use them.

Right click the UI element in the main window you believe should be replaying with the event and choose "view element source". This shows you the HTML for the UI element in the page. Obviously, if the element is missing on the main page, you must consult S-02.

Once you see the ID of the element, make sure that the event ID is the same as the UI element ID in the HTML. If it does not have an ID, or the ID is different in some way, this is why the element is not replaying. One missing UI element can cause all subsequent events to fail, if the event was supposed to trigger display of the UI elements for subsequent UI events.

Q-03. YES Does the page contain form fields but the events seem to be missing?

Yes:

In the event the main page contains a form with fields and there don't seem to be UI events to show the user interaction with the fields, UI events may be missing or not

captured. This can happen because other javascript on the page is "unhooking" our listeners that help us capture user interaction on the page.

If you have verified Q-01 above and the target page is correctly configured, contact engineering if you believe UI events are missing that should be caught.

Q-04. YES Is the failing event of type "change" in the Nav List?

Yes:

If the event is of type "change", you must add this type of event to the list of events RTV is replaying. This event is needed in situations where other associated events are missing such as "KeyUp" in a form field. To enable this event, go to RTV->Tools->Options->UI Events and add the word "change" into the field next to the "Add New Type" button. Once the "Add New Type" button becomes active, click it to add this type. Then click the checkbox next to "change" in the list.

Q-05. YES Does the broken event begin with "xpath" in its entry in the Nav List?

Yes:

Events of type `xpath` are needed when UI elements on the replayed page do not contain IDs. The `xpath` denotes a walk of the HTML DOM to find the element and use it during replay.

The xpath events can be tricky, in that they sometimes require changes to the TeaLeaf UI Capture library to handle special cases.

Contact engineering when xpath events fail during replay.

S-09. I have pages that are messing up replay.

Q-01. Do the pages display as a bunch of (seemingly) nonsensical characters?

Q-01. YES Do the pages have a unique query string or URL that does not match pages you want to display?

Yes:

You can right click them and choose Replay Rules->Remove this page from replay. Remember, this causes all pages with this URL to be removed.

Q-02. Do the pages all have the same URL or query string that also matches pages you want to display?

Yes:

You need to remove them using a unique REQ value. Click the **REQ** button on one of the offending pages and search for a REQ name-value pair that is in the pages you do not want, but is not in the pages you do want. A good example of this is "HTTP_X_REQUESTED_WITH=XMLHttpRequest" which usually denotes AJAX requests you do not want to show.

Highlight this in the REQ and right click choosing "Remove page with this request value from replay". This hides these pages from now on.

Q-02. Do the pages display portions of a webpage?

Q-01. YES Do the pages look like parts of a frameset?

Yes:

See. Q-03.

cxReveal error "Thread was being aborted."

Note: This issue applies to IBM Tealeaf cxReveal from Release 7.1 and earlier.

```
Error
Source: sesnList.aspx.cs:line 119
Search:
Description: Thread was being aborted.
```

This error message may result from failing to configure IBM Tealeaf cxReveal `spanelcfg.xml` file to enable authentication through the `<Authentication>` node.

Unable to replay sessions that were found in cxReveal database

When results are returned from a IBM Tealeaf cxReveal search, it may not be possible to replay the session. When you select the session to replay from the session list, the session may fail to load in Browser Based Replay.

While there are multiple possible reasons as to why the session cannot be retrieved, the issue may be caused by some of the timestamps in the session data.

When a hit is reassembled by the IBM Tealeaf CX Passive Capture Application, the PCA attempts to locate and generate timestamp information. If this information is malformed or missing, the PCA writes the value `01/01/1970` by default.

When this hit is passed to the Processing Server, the following things happen in the order listed below:

- The session attribute information in the hit is processed and written to the IBM Tealeaf cxReveal search database.
- The hit is written to a session archive file with the above timestamp.
- Once per hour on the hour, the session archive file is purged, as the data in it is considered to be too old for storage in the Processing Server.
  - See "Performance Measurement" in the *IBM Tealeaf Passive Capture Application Manual* section of the "Passive Capture Application Manual" in the *IBM Tealeaf Passive Capture Application Manual*.

So, when a IBM Tealeaf cxReveal user performs a search for this session data, the Search database indicates that the data exists in the Processing Server, yet it might be purged. When IBM Tealeaf cxReveal attempts to access these session files, they are not found because of the bad timestamp values.

To test:

1. Through TMS, locate the directory where the session files are stored.
   - For more information, see “Tuning Canister performance” on page 62.
2. Near the end of an hour, log in to the Processing Server as an administrator.

3. Navigate to the directory that you found through TMS.
4. Locate any files that contain the 1970 timestamp in their file name. The file name pattern should look like the following:

   ```
   LSSN_19700101_<hostname>
   ```

5. Using search, try to retrieve a session from this archive.
   - See "cxReveal - Searching Sessions by Session Attribute" in the *IBM Tealeaf cxReveal User Manual*.
6. If you are able to locate these sessions, replay the session through BBR and attempt to ascertain why the timestamp value in the request is corrupted.
   - Timestamp information is stored in the `[timestamp]` section of the request.
   - For more information about these values, see "RealiTea Viewer - Request View" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Chapter 6. Troubleshooting the Tealeaf Management System

Beginning in Release 7.0, Tealeaf enables centralized management of Tealeaf server configurations. Through the Tealeaf Portal, the Tealeaf Management System (TMS) allows administrators to make configuration changes in one central location and then to push those changes to all affected servers. For more information about TMS, see "Tealeaf Management System" in the *IBM Tealeaf cxImpact Administration Manual*.

If you are encountering difficulties with TMS, review these troubleshooting symptoms and solutions.

Startup

TMS slave server fails to appear in the server list on the TMS master

The slave server may be configured as a TMS master server. Complete the following steps:

1. Identify the name of the server that is supposed to be the TMS master.
2. Stop TMS on the mis-configured server.
3. From the command line on the slaver server, run the following command:

   ```
   TLMgmtSrv.exe -setmaster <master_server>
   ```

   where `<master_server>` is the name of the actual TMS master.
4. Start TMS again.
5. Check the event log for errors.

At startup, event logon TMS slave says it is running as a TMS master

You have multiple TMS masters. See “TMS slave server fails to appear in the server list on the TMS master.”

The self-signed TMS certificate has expired

The Tealeaf Management Service (`TLMGmtSrv.exe`) uses a self-signed TMS certificate for SSL communications. If the TMS certificate has expired, you can replace the certificate with another self-signed certificate, or with a certificate from a trusted certificate authority.

You can use the openssl utility on the PCA to generate a private key and a self-signed SSL certificate by using the key. You can then copy the certificate to the IBM Tealeaf server where the TMS service is running.

Complete the following steps to generate a self-signed SSL certificate.

1. Open the command prompt, and navigate to the following directory:

   ```
   /usr/local/ctccap/bin/
   ```
2. Use the **genrsa -out** command to generate the private key. For example, the following command generates the `TLMgmtSrv.key` file, which is a 2048-bit RSA key file:

openssl genrsa -out TLMgmtSrv.key 2048

3. Remove the passphrase from the private key. For example, the following command removes the passphrase:

   copy TLMgmtSrv.key TLMgmtSrv.key.orig

   openssl rsa -in TLMgmtSrv.key.orig -out TLMgmtSrv.key

4. Use the private key that you generated to create the self-signed certificate. Use command prompt and generate the self-signed certificate. For example, use the following command to generate the `TLMgmtSrvNoKey.crt` certificate by using the `TLMgmtSrv.key` private key that you generated:

   openssl req -x509 -days 3650 -newkey rsa:2048 -key ./TLMgmtSrv.key –out TLMgmtSrvNoKey.crt

   The **-days 3650** option specifies the validity of the `example.crt` certificate for the next 365 days (one year).

   The **openssl req** command interactively prompts for various values. The following table displays the prompts and sample replies:

Table 3. Sample values for **openssl req** *command*

Prompt	Sample reply
Country Name (2 letter code)	US
State or Province Name (full name)	California
Locality Name (for example, city)	San Francisco
Organization Name (for example, company)	Tealeaf®, an IBM® Company
Organizational Unit Name (for example, section)	Release Engineering
Common Name (for example, YOUR name)	cx.Tealeaf.com
Email Address	root@cx.Tealeaf.com

5. Combine the two files into a single certificate that contains its own private key by using a **cat** command. For example, use the following command:

   cat TLMgmtSrvNoKey.crt TLMgmtSrv.key > TLMgmtSrv.crt

6. Transfer the new certificate file from the PCA to the Windows-based Tealeaf server whose TMS certificate is to be replaced.

   Make sure that the new certificate has the same name (`TLMgmtSrv.crt`) and is in the same directory as the expired certificate. The default location is `Tealeaf\TLMgmtSrv\SSL`.

7. Stop the TMS Service before installing the new certificate file and restart the service after putting the file in place.

If you have configured the non-SSL mode for the TMS service, an error occurs when you try to access Tealeaf Management System (TMS).

If you have configured the non-SSL mode for the TMS service, the following error is displayed when you access TMS: `There was an error connecting to the TMS service.`

Complete the following steps to resolve this problem:

1. Log in to the Tealeaf Management System (TMS).
2. On the **Worldview** tab, expand **Tealeaf Management Server**, and click **Tealeaf Management System Configuration**.

3. In the **Config Actions** area, click **View/Edit**.
4. Verify that the value of the **UseSSL** parameter to 0.

   By default, the value of the **UseSSL** parameter is 1.

   Make sure that you verify the value on the master and the slave TMS servers.
5. Save and apply your changes.
6. On the TMS master server, go to `$Tealeaf_Home\Portal\WebApp`.
7. Open the `web.config` file, and add the following entry in the `<appSettings>` section:

   ```
   <add key="TMS_USE_SSL" value="false"/>
   ```

   The following section is an example of the `<appSettings>` section in the `web.config` file:

   ```
   <!-- AppSettings -->
   <appSettings>
   ......
   <add key="TMS_USE_SSL" value="false"/>
   ......
   </appSettings>
   ```
8. Save your changes in the `web.config` file.
9. On the TMS master server, go to `$Tealeaf_Home\Portal\WebService`.
10. Open the `web.config` file, and add the following entry in the `<appSettings>` section:

    ```
    <add key="TMS_USE_SSL" value="false"/>
    ```

    The following section is an example of the `<appSettings>` section in the `web.config` file:

    ```
    <!-- AppSettings -->
    <appSettings>
    ......
    <add key="TMS_USE_SSL" value="false"/>
    ......
    </appSettings>
    ```
11. Save your changes in the `web.config` file.
12. Restart the TMS and World Wide Web publishing services on the master server, and the TMS service on the slave servers for your changes to apply.

Access

Request from unknown master

In the `TLMgmtSrv_access_YYYMMDD.log` from a slave server, you may see a message similar to the following:

```
10/10/10 13:55:35 | W | Auth: Request from unknown master at aaa.bbb.ccc.ddd!
```

Where:

- `aaa.bbb.ccc.ddd` is an IP address.

In this case, the TMS master is attempting to connect to TMS on a slave server using an IP address that is different from the address that is used by the slave server to connect to the master. This mismatch can occur if the name for the TMS master on the slave server resolves to a different IP address.

To fix this issue, use the IP address as the name that the TMS master uses to communicate.

1. Log in to the slave server.

2. Open a command-line shell.
3. Navigate to the Tealeaf install directory.
4. Run the following command:

   ```
   TLMgmtSrv.exe -setmaster aaa.bbb.ccc.ddd
   ```

   Where:

   - `aaa.bbb.ccc.ddd` is the IP address of the TMS master.
5. Repeat the above steps on each affected slave server.

Chapter 7. IBM Tealeaf documentation and help

IBM Tealeaf provides documentation and help for users, developers, and administrators.

Viewing product documentation

All IBM Tealeaf product documentation is available at the following website:

https://tealeaf.support.ibmcloud.com/

Use the information in the following table to view the product documentation for IBM® Tealeaf:

Table 4. Getting help

To view...	Do this...
Product documentation	On the IBM Tealeaf portal, go to **?** > **Product Documentation**.
Help for a page on the IBM Tealeaf Portal	On the IBM Tealeaf portal, go to **?** > **Help for This Page**.
Help for IBM Tealeaf CX PCA	On the IBM Tealeaf CX PCA web interface, select **Guide** to access the *IBM Tealeaf CX PCA Manual*.

Available documents for IBM Tealeaf products

Use the following table to view a list of available documents for all IBM Tealeaf products:

Table 5. Available documentation for IBM Tealeaf products

IBM Tealeaf products	Available documents
IBM Tealeaf CX	• *IBM Tealeaf Customer Experience Overview Guide* • *IBM Tealeaf CX Client Framework Data Integration Guide* • *IBM Tealeaf CX Configuration Manual* • *IBM Tealeaf CX Cookie Injector Manual* • *IBM Tealeaf CX Databases Guide* • *IBM Tealeaf CX Event Manager Manual* • *IBM Tealeaf CX Glossary* • *IBM Tealeaf CX Installation Manual* • *IBM Tealeaf CX PCA Manual* • *IBM Tealeaf CX PCA Release Notes*

Table 5. Available documentation for IBM Tealeaf products (continued)

IBM Tealeaf products	Available documents
IBM Tealeaf CX	• *IBM Tealeaf CX RealiTea Viewer Client Side Capture Manual* • *IBM Tealeaf CX RealiTea Viewer User Manual* • *IBM Tealeaf CX Release Notes* • *IBM Tealeaf CX Release Upgrade Manual* • *IBM Tealeaf CX Support Troubleshooting FAQ* • *IBM Tealeaf CX Troubleshooting Guide* • *IBM Tealeaf CX UI Capture j2 Guide* • *IBM Tealeaf CX UI Capture j2 Release Notes*
IBM Tealeaf cxImpact	• *IBM Tealeaf cxImpact Administration Manual* • *IBM Tealeaf cxImpact User Manual* • *IBM Tealeaf cxImpact Reporting Guide*
IBM Tealeaf cxConnect	• *IBM Tealeaf cxConnect for Data Analysis Administration Manual* • *IBM Tealeaf cxConnect for Voice of Customer Administration Manual* • *IBM Tealeaf cxConnect for Web Analytics Administration Manual*
IBM Tealeaf cxOverstat	*IBM Tealeaf cxOverstat User Manual*
IBM Tealeaf cxReveal	• *IBM Tealeaf cxReveal Administration Manual* • *IBM Tealeaf cxReveal API Guide* • *IBM Tealeaf cxReveal User Manual*
IBM Tealeaf cxVerify	*IBM Tealeaf cxVerify Administration Manual*
IBM Tealeaf cxView	*IBM Tealeaf cxView User Manual*
IBM Tealeaf CX Mobile	• *IBM Tealeaf CX Mobile Android Logging Framework Guide* • *IBM Tealeaf Android Logging Framework Release Notes* • *IBM Tealeaf CX Mobile Administration Manual* • *IBM Tealeaf CX Mobile User Manual* • *IBM Tealeaf CX Mobile iOS Logging Framework Guide* • *IBM Tealeaf iOS Logging Framework Release Notes*

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Bay Area Lab
1001 E Hillsdale Boulevard
Foster City, California 94404
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample

programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

Privacy Policy Considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. A cookie is a piece of data that a web site can send to your browser, which may then be stored on your computer as a tag that identifies your computer. In many cases, no personal information is collected by these cookies. If a Software Offering you are using enables you to collect personal information through cookies and similar technologies, we inform you about the specifics below.

Depending upon the configurations deployed, this Software Offering may use session and persistent cookies that collect each user's user name, and other personal information for purposes of session management, enhanced user usability, or other usage tracking or functional purposes. These cookies can be disabled, but disabling them will also eliminate the functionality they enable.

Various jurisdictions regulate the collection of personal information through cookies and similar technologies. If the configurations deployed for this Software Offering provide you as customer the ability to collect personal information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for providing notice and consent where appropriate.

IBM requires that Clients (1) provide a clear and conspicuous link to Customer's website terms of use (e.g. privacy policy) which includes a link to IBM's and Client's data collection and use practices, (2) notify that cookies and clear gifs/web beacons are being placed on the visitor's computer by IBM on the Client's behalf along with an explanation of the purpose of such technology, and (3) to the extent required by law, obtain consent from website visitors prior to the placement of cookies and clear gifs/web beacons placed by Client or IBM on Client's behalf on website visitor's devices

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Online Privacy Statement at: http://www.ibm.com/privacy/details/us/en section entitled "Cookies, Web Beacons and Other Technologies."

Printed in USA