

IBM Tealeaf CX
Version 9 Release 0.2
June 18, 2015

IBM Tealeaf CX Installation Manual



Note

Before using this information and the product it supports, read the information in "Notices" on page 133.

This edition applies to version 9, release 0, modification 2 of IBM Tealeaf CX and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 1999, 2015.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Chapter 1. Overview 1

Overview of the installation process	1
Installation roadmap.	2
IBM Tealeaf overview	3
IBM Tealeaf CX	4
IBM Tealeaf products	8
Data processing overview	11
IBM Tealeaf services	11
IBM Tealeaf user accounts	13
Data privacy and IBM Tealeaf CX	14
How data sessionization works.	22
Internationalization support	25
Support for IPv6.	30

Chapter 2. Planning the installation . . . 43

Installation topologies for IBM Tealeaf CX	43
Topology variations.	43
Support for virtualized server environments	46
System requirements for installing IBM Tealeaf	47
Operating system requirements.	47
Software requirements.	49
Hardware requirements	51
User account requirements	52
High-availability considerations	53
High availability: CX Passive Capture Application server	53
High availability: Processing servers	55
High availability: Portal web application	56
High availability: Reporting server	57
Example: High availability configuration	57
Web application development considerations	58

Chapter 3. Preparing to install the software 65

Identifying available resources	65
Recording machine identities	66
Installing and configuring prerequisite software	66
Installing .NET framework	67
Registering IIS	68
Setting up SQL Server to work with Tealeaf	68
Preparing the installation environment	71
Configuring IBM Tealeaf user accounts	71
Preparing the domain	72
Preparing the network environment	73
Configuring the network interface cards (NICs)	75
Checking server requirements	75
File permission considerations	76
User account considerations	76

Chapter 4. Installing IBM Tealeaf. 79

Installing on a single server	79
Obtaining installation files	79
Installing IBM Tealeaf CX.	80
Tealeaf CX silent installation, uninstallation, and upgrade	85

Installing CX silently	90
Upgrading silently	90
Enabling your IBM Tealeaf solution	91
Uninstalling CX silently	95
Post-installation tasks	95
Post-installation database tasks	95
Verifying system startup	96
Backing up the IBM Tealeaf event model	103
Rebuilding the canister	103
Uninstalling IBM Tealeaf	103

Chapter 5. Configuring IBM Tealeaf 105

Configuring IBM Tealeaf components	105
Configuring Portal settings	106
Configuring the IBM Tealeaf system time zone	108
Configuring IBM Tealeaf data objects	108
Accessing the Event Manager to configure data objects.	109
Configuring Tealeaf dimensions	110
Creating deviations	112
Configuring sessionization settings	112
Configuring session timeout settings.	112
Changing top-level domains	113
Configuring session close events	113

Chapter 6. Administering and managing IBM Tealeaf CX 115

Configuring IBM Tealeaf cycle services	115
Backing up and restoring IBM Tealeaf session data	116
Creating a shutdown script in Windows	116
Resetting IBM Tealeaf system service stop times	117
Updating user agent files	117
Installing the most up-to-date version of the fixed user agent file (BrowsCap.csv)	118
Installing the most up-to-date version of the mobile user agent file (WURFL.csv)	118
Managing and installing IBM Tealeaf patches.	119
How IBM Tealeaf patches work	119
Managing IBM Tealeaf patches	119
Installing IBM Tealeaf patches	120

Chapter 7. Troubleshooting 123

Using the IBM Tealeaf web application installation utility	123
How the IBM Tealeaf web application installation utility works.	123
Installing or re-installing IBM Tealeaf CX Portal components	128
Using the TLTMaint script to maintain Canister stability	128
Troubleshooting privacy performance in the PCA server	129

**Chapter 8. IBM Tealeaf documentation
and help 131**

Notices 133

Trademarks 134

Privacy Policy Considerations 135

Chapter 1. Overview

The IBM Tealeaf CX Installation Manual provides information to help IBM Tealeaf customers and Tealeaf professional services plan for, install, configure, and administer IBM Tealeaf CX.

This chapter provides an overview of the installation process, a product overview, which describes the IBM Tealeaf CX platform and the IBM Tealeaf products, and a data processing overview.

How to use the information

The way to install and use IBM Tealeaf CX can vary depending on your business requirements and your IT infrastructure. Because of this, there is no "one size fits all" installation. That said, the information in this guide demonstrates the sequence of activities that all customers need to complete to successfully plan for, install, configure, and maintain an IBM Tealeaf CX installation.

In addition to the IBM Tealeaf CX Installation Manual, you need access to the following publications:

- *IBM Tealeaf CX: Tealeaf Troubleshooting Guide*
- *IBM Tealeaf CX: Tealeaf Databases Guide*
- *IBM Tealeaf CX: Configuration Manual*
- *IBM Tealeaf CX Version 9 Release 0 Upgrade Manual*
- *IBM Tealeaf Event Manager Manual*

To access IBM Tealeaf publications:

1. Sign in to the IBM Client Success portal .
2. Sign in to the IBM Tealeaf portal.
3. Select the version of the product you are working with from the Product Documentation window.
4. Locate and download the PDF.

Related concepts:

"Installation roadmap" on page 2

"IBM Tealeaf overview" on page 3

Overview of the installation process

Installing IBM Tealeaf CX involves planning your installation, setting up your computing and network environment to accept IBM Tealeaf CX, configuring and registering Windows server features and functions, and then running the installer with the proper configuration settings.

The installer is multi-panel wizard for gathering configuration data and for implementing the IBM Tealeaf CX configuration in your environment.

After installing IBM Tealeaf CX, you can install and configure the IBM Tealeaf products that you have licensed. You can also modify your IBM Tealeaf CX configuration if you did not complete entries on all of the panels at install time.

The *Tealeaf CX Installation Manual* provides instructions for installing, configuring, and maintaining IBM Tealeaf CX. It includes references to other IBM Tealeaf publications for instructions on installing and configuring the CX features that you have licensed.

See the *Installation roadmap* topic in this guide for the sequence of tasks to install and configure IBM Tealeaf CX.

Installation roadmap

The installation roadmap presents the task-flow for preparing your environment and installing and configuring IBM Tealeaf CX.

IBM Tealeaf CX is the platform on which all IBM Tealeaf solutions run. In a typical installation, you install IBM Tealeaf CX and then you enable the CX features and functions needed to address your customer management business needs.

Table 1 lists the sequence of steps for installing IBM Tealeaf CX and provides links to the topics that contain the detailed instructions.

Table 1. Roadmap of activities for installing and configuring the software

Step	Description	Link
1	Plan your installation topology. <i>An installation topology</i> is the physical and logical layout of the servers in your IBM Tealeaf CX installation.	"Installation topologies for IBM Tealeaf CX" on page 43
2	Identify your assets to make the best use of those resources already available.	"Identifying available resources" on page 65
3	Consider High-availability. For a Tealeaf solution, high availability pertains to a configuration that ensures uninterrupted service of the capture and processing servers and can include strategies for redirecting traffic among multiple PCA devices on a network.	"High-availability considerations" on page 53
4	Address all the software and hardware requirements needed to support your IBM Tealeaf CX installation.	"System requirements for installing IBM Tealeaf" on page 47
5	Perform all pre-installation assessments of your environment, such as: <ul style="list-style-type: none"> Assessing your current IT resources Determining and cataloging the IBM Tealeaf CX component-to-server assignments 	Chapter 3, "Preparing to install the software," on page 65

Table 1. Roadmap of activities for installing and configuring the software (continued)

Step	Description	Link
6	<p>Configuring your environment and installing the prerequisite software, which includes the following activities:</p> <ul style="list-style-type: none"> • Configuring Windows Server by registering the features and functions that IBM Tealeaf CX requires, such as: <ul style="list-style-type: none"> – Installing .NET – Enabling .NET and ASP.NET functions in the Application Development role in Windows Server. – Installing SQL Server • Preparing your environment for IBM Tealeaf, which includes the following tasks: <ul style="list-style-type: none"> – Creating user accounts – Preparing the domain – Preparing the network environment – Configuring your network interface cards – Checking server requirements 	<ul style="list-style-type: none"> • “Installing and configuring prerequisite software” on page 66 • “Preparing the installation environment” on page 71
7	Install IBM Tealeaf CX.	“Installing IBM Tealeaf CX” on page 80
8	Enable IBM Tealeaf CX features and functions	“Enabling your IBM Tealeaf solution” on page 91
9	Perform required post-installation tasks	Chapter 5, “Configuring IBM Tealeaf,” on page 105

Related concepts:

Chapter 1, “Overview,” on page 1

IBM Tealeaf overview

IBM Tealeaf is a digital customer experience management and customer behavior analysis solution.

It consists of IBM Tealeaf CX, an industry-leading robust datastore of online customer information (and the engine behind all IBM Tealeaf products) and suite of products for customer behavior analysis, customer service optimization, and integrations.

Figure 1 on page 4 illustrates the IBM Tealeaf platform and the suite of products that it supports.

Each product suite (Customer Behavior Analysis, Customer Service Optimization, and Integrations) has its own set of products.




Figure 1. IBM Tealeaf customer experience solutions

Related concepts:

Chapter 1, “Overview,” on page 1

“IBM Tealeaf CX”

“IBM Tealeaf CX component architecture” on page 5

“IBM Tealeaf products” on page 8

Related tasks:

Chapter 4, “Installing IBM Tealeaf,” on page 79

IBM Tealeaf CX

IBM Tealeaf CX is the datastore of online customer information and the engine behind all IBM Tealeaf products.

It provides the framework on which all IBM Tealeaf solutions run.

IBM Tealeaf CX consists of several independent components that can either coexist on the same server or reside on separate servers. From an installation perspective, you must install and configure IBM Tealeaf CX and its components in order to use any IBM Tealeaf product or solution.

Related concepts:

“IBM Tealeaf overview” on page 3

“IBM Tealeaf products” on page 8

“IBM Tealeaf CX component architecture” on page 5

Related tasks:

“Installing IBM Tealeaf CX” on page 80

IBM Tealeaf CX component architecture

IBM Tealeaf solutions are distributed into three core components as follows:

- Passive Capture Application
- Tealeaf Server environment, which includes:
 - Processing server
 - Storage server, which includes
 - A Reporting component
 - An Archiving component
- IBM Tealeaf User clients, which can include:
 - IBM Tealeaf Portal
 - IBM TealeafCX RealiTea Viewer
 - A stand-alone application that can be deployed to Tealeaf users.

Depending the topology of your IBM Tealeaf installation, these components can coexist on the same server or reside on separate servers.

The following sections describe how the core components function.

Passive Capture component

The IBM Tealeaf Passive Capture Application (PCA) component is a network sniffer that captures the HTTP(S) request and response data passively by "sniffing" TCP/IP packets from the network via an existing spanning port or network tap. Example types of captured data include:

- HTTP headers
- URLs
- URL form fields (both get and post)
- URL Referrers
- user login names
- cookies
- server host names
- application
- client IPs
- page sizes,
- The HTML source code of the page returned to the customer and other diagnostics included in the HTTP(S) interaction between the customer's web browser and the web application.

The network sniffer provides a non-intrusive, real-time method of capturing what each customer is doing and seeing across the entire session with your Web application. It does not introduce any overhead, latency, and risk of disrupting traffic to the application.

The capture process records the total round trip, Internet delivery, and page generation times for each page requested in the session, page cancellations, and the connection-type experienced by the customer (e.g., T1, DSL, or dial-up) as measured by the page delivery rate.

The PCA server streams captured data to the IBM Tealeaf CX Server environment where it is received by Transport component of the Processing Server.

For additional security, IBM Tealeaf supports encrypting this data transport. For more information, see the "Passive Capture Application Manual" in the *IBM Tealeaf Passive Capture Application Manual*.

Tealeaf server environment

The Tealeaf Server is a highly scalable, real-time, distributed platform that processes, analyzes, indexes and archives the recorded data. The Tealeaf Server also serves as the platform for the IBM Tealeaf Portal and RealTea Viewer clients and the IBM Tealeaf family of products.

Typically deployed behind the DMZ in a trusted network segment, the Tealeaf Server environment operates on the Windows 2003 and 2008 (64-bit) Server or Advanced Server platform.

The Tealeaf server environment includes two server components:

- Processing server component
- Storage server component

The Processing server is a real-time data processing environment that receives streamed captured hits (data) from the Passive Capture Application server. The data is received through the Transport component of the Processing server. The Transport component reads the data into the Tealeaf pipeline, which consists of a series of agents that perform specific filtering and manipulation functions against the data stream to normalize and secure the data.

The processes performed by the agents include:

- Streaming the data into the Short Term Canister, an in-memory database.
- The Short Term Canister organizes by session the sequence of hits into the order experienced by the visitor.
- The Short Term Canister inspects, aggregates, and evaluates the hit data.

Data inspection is conducted by the event engine, which analyzes the captured data in real-time against a set of user-defined rules to provide insight into business process health, customer activity, application errors, page errors, and response time problems.

- Events can be defined against the occurrence or absence of a single value, text pattern, page attribute, session attribute, or compound conditions appearing in a page or across the entire session.

See "TEM Events Tab" in the *IBM Tealeaf Event Manager Manual*.

- Observed events are used for real-time alerting, reporting, and searching.
 - Alerting functions are processed by an alerting service, which evaluates event counts against user-defined thresholds and alerting actions. See "Configuring the Alert Service" in the *IBM Tealeaf CX Configuration Manual*.

- Aggregate event counts are populated by a data collector service into the Report Server reporting database on the Storage Server. See "Data Aggregation and Retention" in the *IBM Tealeaf cxImpact Administration Manual*.
- Event markers are also embedded into the sessions to enable indexing and session retrieval.

A key filtering function in the transport component removes or encrypts sensitive data to avoid unauthorized access from unauthorized personnel. Other common pipeline operations include the following:

- Data Removal (e.g., removing hits from an unwanted source as keep-alive hits)
- Data Queuing
- Data Inflation (e.g. inflating hits with compressed responses)
- Privacy
- Normalization of some fields (e.g., application, hostname, path, server, browser. OS)
- Routing

The Processing server component also indexes completed sessions to enable retrieval using both free-text and parametric search. The recorded HTML session data and associated indexes to the sessions are written to local disk in an embedded flat-file database referred to as the **Long Term Canister**.

For information about indexing, see "Configuring CX Indexing" in the *IBM Tealeaf CX Configuration Manual*.

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

The Storage server includes the *Reporting* and *Archiving* components.

The Archiving component provides a data storage environment for long-term storage of collected data. The archiving function lets you archive data for extended periods of time for problem resolution customer behavior analysis, dispute resolution, and other web archiving needs. IBM Tealeaf CX intelligently determines which sessions are of interest and should be archived based on user-defined data management rules. For example, you can archive customer sessions that include a certain class of errors or business events, a random sampling, or 100 percent of sessions captured. Sessions not marked for archiving are discarded.

By providing a distributed, dedicated storage environment with minimal disk activity, the Storage Server improves overall search response time and provides the ability to conduct data backup without taking real-time processing off-line.

The archiving component copies and maintains saved sessions and their associated indexes. At the end of each operational day, it copies the saved sessions and indexes to specified drive locations on the Storage Server.

Upon expiration, sessions and indexes can either be automatically deleted or saved off to cold storage.

The Reporting component of the storage server consists of Report Data Collector services, which collect the Processing component aggregated values and populate the reporting database, and the Portal Web Application, the user interface.

Normally, the Archiving and Reporting sub-components are kept together but for optimum efficiency they may be placed on separate machines for high-volume sites.

User clients

The IBM Tealeaf cxImpact solution includes two user clients:

- IBM Tealeaf Portal
- IBM Tealeaf RealiTea Viewer

A stand-alone application that can be deployed to Tealeaf users.

The IBM Tealeaf Portal is a real-time, web-based console that provides a centralized workspace for the production support team to identify, size and diagnose issues impacting business-critical web applications, as well as a tool for administrators to manage the overall health of IBM Tealeaf solutions.

It offers a real-time view of user activity, searching capabilities and reporting functionality to provide rapid awareness and problem resolution.

The RealiTea Viewer is a win32 client used by the production support team and other users of IBM Tealeaf solutions to recreate problems or other issues by visually replaying the real user's interaction with the web application. The user session is replayed step-by-step, as it was recorded at the time the real end-user conducted the session. In addition, the RealiTea Viewer also provides advanced search and correlation functionality for advanced causal factor isolation and problem diagnosis. For more information about the IBM Tealeaf RealiTea Viewer, see "RealiTea Viewer (RTV)" in the *IBM Tealeaf RealiTea Viewer User Manual*.

The IBM Tealeaf Portal web application is served from the Tealeaf Server. The IBM Tealeaf Portal can be served from over HTTP (port 80) or HTTPS (port 443). End-user access to the IBM Tealeaf Portal is controlled via its native authentication system or by using Windows NT Authentication to administer groups and users on the Tealeaf Server machine. Database access via the RealiTea Viewer can also be controlled using Windows NT Authentication (see Access controls using Windows Domain Authentication). In addition, data query, end-user and administrative actions conducted via the IBM Tealeaf Portal and Viewer clients is logged to enable auditing of which end-user requested and accessed what data (see Access and change auditing).

Related concepts:

"IBM Tealeaf overview" on page 3

"IBM Tealeaf CX" on page 4

"Topology variations" on page 43

Related tasks:

"Installing IBM Tealeaf CX" on page 80

IBM Tealeaf products

An IBM Tealeaf solution can consists of various IBM Tealeaf products.

IBM Tealeaf provides three distinct suites of products that address the different aspects of online Customer Experience Management. All of these products are powered by the IBM Tealeaf CX platform.

Customer Behavior Analyses suite of products

The IBM Tealeaf Customer Behavior Analysis products help businesses with their online customer experience.

The Tealeaf products in this suite are used for observing and analyzing the online behavior of customers as they interact with your web site.

Typically, the customer behavior analysis products answer the why questions about your customer-facing website. For example:

- Why do more customers abandon the credit card application on the second step rather than the first step?
- Why are customers searching for products multiple times and still not adding items to the shopping cart?

Products in the Customer Behavior Analyses suite include:

- Tealeaf cxImpact
- Tealeaf cxView
- Tealeaf cxOverstat

IBM Tealeaf cxImpact transforms the dataset captured by IBM Tealeaf CX into visually re-playable and completely searchable customer sessions. IBM Tealeaf cxImpact not only captures the page-by-page, browser-level recording of each customer session, but also offers one-click access to all of the supporting HTTP(S) request-and-response information for further technical analysis. For more information, see Tealeaf cxImpact .

IBM Tealeaf cxView gathers the rich, customer experience dataset of IBM Tealeaf cxImpact and puts it into executive-level dashboards, scorecards, and reports. Additionally, IBM Tealeaf cxView includes a powerful early warning system that makes use of algorithmic discovery to automatically determine which areas of your site represent struggle sources for visitors. For more information, see Tealeaf cxView .

IBM Tealeaf cxOverstat allows your company to:

- Identify problematic hotspots
- Optimize page content
- Improve form conversion rates

IBM Tealeaf cxOverstat is a scalable enterprise solution capable of supporting massive, high-volume sites and is an add-on product to the Tealeaf CX platform and cxImpact. For more information, see Tealeaf cxOverstat.

Customer Service Optimization suite of products

IBM Tealeaf Customer Service Optimization products improve communication between a company's call center and its web operations for more effective customer service in multichannel environments.

You can use IBM Tealeaf Customer Service Optimization products to help customer service representatives understand the full context of a customer's online sessions by preserving online interactions.

Products in the Customer Service Optimization suite include:

- Tealeaf cxReveal
- Tealeaf cxVerify

IBM Tealeaf cxReveal makes use of the core functionality of the Tealeaf CX platform — session replay and search — and packages it into a simple user interface for efficient use by customer service representatives, sales associates and marketing personnel. Using one-click retrieval, users have instant access to both live and historical customer sessions from any existing CRM solution. For more information, see Tealeaf cxReveal.

IBM Tealeaf cxVerify preserves a record of online customer interactions on a website or mobile device. It helps online businesses maintain highly reliable records for dispute resolution, fraud investigations and audit and compliance purposes. IBM Tealeaf cxVerify provides a thorough, accurate snapshot of entire customer sessions to maintain a permanent record of customer online interactions and transactions. For more information, see Tealeaf cxVerify.

Customer Experience Integration suite products

The IBM Tealeaf Customer Experience Integration products provide seamless integration with other business applications including business intelligence, web analytics, and voice of customer solutions. As the only solution capable of capturing all the data about online customers, Tealeaf's rich customer experience dataset is critical to successful cross-channel analysis and web site optimization.

Products in the Customer Experience Integration suite include:

- Tealeaf cxConnect for Data Analysis
- Tealeaf cxConnect for Voice of Customer
- Tealeaf cxConnect for Web Analytics

IBM Tealeaf cxConnect for Data Analysis allows users to integrate Tealeaf's rich customer experience dataset seamlessly with any business intelligence or reporting application to create a multi-channel view of the customer for ongoing analysis. For more information, see Tealeaf cxConnect for Data Analysis.

IBM Tealeaf cxConnect for Voice of Customer allows you to seamlessly integrate Tealeaf's rich customer experience dataset with any VOC application. Businesses use VOC technologies to understand the perspectives of and gather feedback from their online customers. When trying to uncover why customers are providing specific feedback, companies are able to use the integration with Tealeaf to review the actual experience of the individual users who provided feedback. This visibility provides full context about the user's experience and their feedback to ensure that businesses can effectively refine their efforts. For more information, see Tealeaf cxConnect for Voice of Customer.

IBM Tealeaf cxConnect for Web Analytics allows you to seamlessly integrate Tealeaf's rich customer experience data set with any web analytics application to help uncover trends or anomalies such as a difference in conversion rates between two time periods or customer segments. When integrated, cxConnect for Web Analytics enable web analysts to see the actual customers behind these trends — in order to investigate why each individual customer is either succeeding or failing online. Tealeaf cxConnect requires the Tealeaf CX platform and cxImpact. For more information, see Tealeaf cxConnect for Web Analytics.

The following illustration is an example of an IBM Tealeaf solution. It shows the IBM Tealeaf CX platform and its components stored on two servers and the IBM Tealeaf products available to the employees.




Figure 2. IBM Tealeaf CX architecture

Related concepts:

“IBM Tealeaf overview” on page 3

“IBM Tealeaf CX” on page 4

Data processing overview

IBM Tealeaf data processing includes managing services, handling user authorization, providing for data privacy, implementing data sessionization and supporting internationalization standards.

Information in the following sections provide explanations about IBM Tealeaf data processing.

IBM Tealeaf services

IBM Tealeaf provides a services for various types of functionality.

Review information in the following sections to learn about these services.

IBM Tealeaf services common to all machines

Avoid using the Windows Services Control Panel to stop and start IBM Tealeaf services. Whenever possible, use the shortcuts that are inserted into the Start menu.

Table 2. IBM Tealeaf services common to all machines

Service Name	Description
Tealeaf Management Server (TLMgmtsrv.exe)	Provides management for Tealeaf components for that machine.
Tealeaf Search Server (TlSrchsrv.exe)	Provides session search capabilities through a TCP/IP network.

IBM Tealeaf processor machine services

Processor machine services include Canister Manager, Canister Server, Session Indexer, and Transport Service.

Table 3. IBM Tealeaf Process Machine Services

Service Name	Description
Tealeaf Canister Manager (TlCanMgr.exe)	<p>This service manages and monitors a set of sub-processes. If one of the sub-processes quits unexpectedly, the Canister Manager attempts to restart the service.</p> <p>Note: To stop and restart the Canister Manager, use the Start menu shortcuts under the IBM Tealeaf CX Portal heading. The Start menu shortcut executes CanSvcs.exe, which manages the flow of hits into the canister and spooling operations. Using the Windows Services Control Panel to stop and start the Canister Manager may cause data loss.</p> <p>If the Canister Manager Service is modified to run under a specific local user other than system, all sub-processes also run under this user.</p> <ul style="list-style-type: none"> • TLStatColl.exe - Reads session statistics from the Session Evaluator and Event Reporter, totals these statistics, and publishes them to the NDLT and NDLY tables in the Short Term Canister. • TLPerfColl.exe - Polls the Short and Long Term Canisters for Canister statistics and makes the statistics available to the Data Service for inclusion into IBM Tealeaf cxImpact and Portal Status for system monitoring. • TLsesnRec.exe - Writes sessions marked for archiving from the Short Term Canister to the Long Term Canister. • TLsesnEval.exe - Closes sessions for archival upon inactivity timeout. It also evaluates end of session for events. • TLEvtRep.exe - Evaluates every hit of every session for events and defined page-level events.
Tealeaf Canister Server (ctreesql.exe)	Provides access to the STC and LTC database services.
Tealeaf Session Indexer (IndexProgram.exe)	Creates indexes of captured data so you can perform faster, more effective searches. This program can be run as a service or from the command line.
Tealeaf Transport Service (TeaLeafCaptureSocket.exe)	Receives hit data and transports it to the Short Term Canister. The service may also perform pipeline operations if this machine is also the Transport machine.

IBM Tealeaf data collector services

IBM Tealeaf includes services for collecting data for reporting.

Table 4. IBM Tealeaf data collector services

Service Name	Description
Tealeaf Reporting Server (TeaLeafReportingService.exe)	Collects Tealeaf Canister data, aggregates it, and prepares Tealeaf Reports.

Table 4. IBM Tealeaf data collector services (continued)

Service Name	Description
Tealeaf RSE Service (TeaLeafRSEService.exe)	Generates resultset data for reporting through the use of predefined resultsets or search templates.
Tealeaf Alert Service (TLAlertSrv.exe)	Provides real-time event alert detection and messaging.
Tealeaf Scheduling Service (TLSchedulerSvc.exe)	Manages scheduling of Tealeaf utilities (Tealeaf Status report and more).
Tealeaf Transport Service (TeaLeafCaptureSocket.exe)	Receives statistics hit data and inserts into the reporting statistics database.

IBM Tealeaf web tier service

IBM Tealeaf web tier service includes the ability to generate session data for replay within the Portal.

Table 5. IBM Tealeaf web tier service

Service Name	Description
Tealeaf Replay Server (TLReplaySrv.exe)	Generates session data for replay within the Portal.

IBM Tealeaf transport machine services

IBM Tealeaf transport machine service receives hit data from the capture device and processes the data.

Table 6. IBM Tealeaf transport machine service

Service Name	Description
Tealeaf Replay Server (TLReplaySrv.exe)	Receives hit data from the capture device, does common hit processing (data deletion, inflating compressed data, reference, and more).

IBM Tealeaf user accounts

IBM Tealeaf user accounts include account permissions, NT authentication, and Passive Capture Application accounts. The user accounts are described in the following sections.

Account permissions

Accounts used by a Windows-based Tealeaf service should be given Local System permissions on the server.

Note: IBM Tealeaf does not support using non-local system accounts for running IBM Tealeaf software. Running IBM Tealeaf services using permissions other than Local System permissions might cause problems, such as failures to connect, write errors, and unexpected time-outs.

NT authentication

When NT authentication is enabled for the IBM Tealeaf system, IBM Tealeaf services must run under an NT domain account with appropriate privileges. See "Authentication" in the *IBM Tealeaf cxImpact Administration Manual*.

Additional configuration might be required. See "Database Manager SQL Server Access and Permission Requirements" in the *IBM Tealeaf Databases Guide*.

Passive Capture Application accounts

The IBM Tealeaf Passive Capture Application runs on the Linux operating system. During PCA installation, the ctccap account is created, which is used for running all PCA services. See "Installation" in the *IBM Tealeaf Passive Capture Application Manual*.

Data privacy and IBM Tealeaf CX

As IBM Tealeaf CX captures and processes data, the private or personal information of the people who visit your web site can be changed, masked, or removed.

IBM Tealeaf has three points of privacy management where it can change, mask, or remove private or personal information. Each point of privacy management has distinct features, uses, and best practices associated with it.

Information in the following sections describe how privacy is handled in each privacy management area of the IBM Tealeaf system.

Note: IBM Tealeaf and the IBM Tealeaf CX system do not host or otherwise manage the personal data of visitors to your web application. During implementation and usage of IBM Tealeaf, you must determine how to comply with any local privacy laws that apply. Seek guidance from your own counsel on privacy law compliance.

Data privacy and Web application development

Data privacy in IBM Tealeaf can have implications on web application development at your company.

IBM Tealeaf can block or encrypt data in the request or response. However, there are a number of good practices in web application development that may facilitate management of data privacy.

Fields that have been encrypted using privacy rules in the IBM Tealeaf Passive Capture Application or Windows pipelines cannot be decrypted in the Portal. These encrypted fields can be decrypted only during replay.

As an alternative, you can leave the configured fields in unencrypted state in the session data and then define privacy rules specifically to be applied during session replay, permitting the display of the unencrypted data in the Portal, as needed.

Related concepts:

"Web application development considerations" on page 58

UI capture privacy

If the IBM Tealeaf UI Capture for AJAX has been implemented in your web application, you can manage the transmission of HTML form data from within the visitor's browser via JavaScript provided by IBM Tealeaf.

UI Capture privacy masks or blocks information that is collected from HTML forms and other browser elements by the IBM Tealeaf UI Capture for AJAX. This

JavaScript library periodically sends data back to your web server which enables capture and processing by the Tealeaf system.

The IBM Tealeaf UI Capture for AJAX requires additional installation and implementation in your web application. See "UI Capture for AJAX Guide" in the *IBM Tealeaf UI Capture for AJAX Guide*.

Because this client-side information can contain sensitive or personal visitor data, you might need to use IBM Tealeaf UI Capture for AJAX to cleanse the data before it leaves the browser. Cleansing data also distributes the act of blocking visitor data across all visitors' computers, which lowers the processing overhead for the Tealeaf system.

UI capture privacy cleanses data through by masking or blocking it. It can only mask or block data collected via the IBM Tealeaf UI Capture for AJAX. This library does not provide access to data contained on the page, which is not managed by the library. The following types of data cannot be made private by UI Capture:

- A visitor ID embedded in the HTML of the page.
- A static element not captured by the library, which contains the visitor's account balance or Social Security number, or similar.

Implementation considerations

You cannot use UI Capture Privacy to cleanse data if you do not have the IBM Tealeaf UI Capture for AJAX implemented in your Tealeaf system.

In the IBM Tealeaf UI Capture for AJAX, the `TealeafClientCfg.js` includes the `tlfieldblock` configuration object where you can specify the fields whose values you wish to mask. You can replace values with a specified string, which enables searching for fields that have been masked through the Tealeaf Portal, without revealing the data. See "Data Privacy in UI Capture" in the *IBM Tealeaf UI Capture for AJAX Guide*.

Passive capture privacy

After data is passed to the IBM Tealeaf system, the first place in which data can be cleansed is on the IBM Tealeaf Passive Capture Application (PCA) server.

Cleansing data at the PCA server is known as *passive capture privacy*.

The PCA server is the optimal place for blocking internal employee access to customer-sensitive information.

Types of personal information that you might want to block at the PCA server include:

- Visitor account number
- Social Security numbers
- Credit card information

In addition to preventing access to customer-sensitive information, you can use PCA privacy management tools enable to encrypt customer-sensitive data, which can then be decrypted at a later time in the process. Blocking or encrypting data on the PCA using privacy rules ensures that the rest of the IBM Tealeaf system does not ever have plain-text access to customer-sensitive data that belongs to visitors of your web site.

How to decide if implementing privacy on the PCA server is the right thing to do

Implement PCA privacy if you want to:

- Block private data
- Encrypt private data that must be recalled later
- Drop traffic that does not need to be captured
- Drop the response or "body" of unnecessary hits

Note: By default, the PCA is configured to sanitize all user input data from form fields for security purposes. If needed, you can disable this setting with the PCA Web Console and manage data privacy for fields on an individual field basis. This sanitization is managed through the TextBlockURLFields action configured by default in Rule 1. See "PCA Web Console - Rules Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

Performance implications of PCA privacy

Because the PCA processes each hit in real-time, consider limiting privacy processing as much as possible. If the PCA is burdened with many privacy rules, then it may not be able to handle the traffic in real-time, which might result in dropped or missing hits. As you configure your privacy management rules for the PCA, you should move privacy processing that is not required to the downstream IBM Tealeaf Privacy session agent in the Windows pipeline.

Note: At this time, the PCA does not support multithreaded privacy. Privacy actions can only utilize a single core. Running the PCA Privacy on a multi-CPU system does not alleviate processing load problems presented by excessive privacy actions.

Through the PCA Web Console, you can configure privacy rules to find beginning and ending text patterns, between which the text can be blocked. Privacy also enables regular expression pattern matching, so complicated text patterns can be applied. Using regular expressions is expensive to process. For more information on the web console, see "Passive Capture Configuration via Web Console" in the *IBM Tealeaf Passive Capture Application Manual*.

Pre-Built privacy rules IBM Tealeaf:

IBM Tealeaf provides prebuilt rules that enable the blocking or encryption of data fields that have been marked in the response with the appropriate comments.

When you are building your web application, you can mark content in the response with HTML comments. These comments are not visible to the user. However, IBM Tealeaf can scan the response for them and then blocks or encrypts accordingly.

For example, suppose the HTML response from your page includes the following user-readable text:

```
Your Social Security No. is 123-12-1234
```

This information may be necessary for the visitor to your site, yet for security reasons it should be blocked from the data captured by IBM Tealeaf or encrypted so that only those IBM Tealeaf users that have been authorized can see it.

Note: To manage these operations in IBM Tealeaf, it is a recommended practice that you bracket data to block or encrypt with a specific set of HTML comments.

Example of blocking data

To block the Social Security number data (123-12-1234) from IBM Tealeaf, you should bracket the data in the response using the following HTML comments:

```
Your Social Security No. is <!-- TLTIBH--> 123-12-1234<!-- TLTIE-->
```

IBM Tealeaf Privacy can be configured to recognize these specific tags and then to block the data between them.

It is possible to configure IBM Tealeaf to recognize any set of bracketing comments. However, when the provided privacy rule is enabled, the above bracketing comments are automatically recognized by IBM Tealeaf, and the data is blocked. The resulting output sent to IBM Tealeaf is the following:

```
Your Social Security No. is <!-- TLTIBH-->XXXXXXXXXX<!-- TLTIE-->
```

To enable this privacy rule, additional configuration is required. See "Privacy Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Example of encrypting data

you can configure the HTML output to enable IBM Tealeaf to encrypt the data. In our example above, the generated HTML might look like the following:

```
Your Social Security No. is <!--TLTENB-->123-12-1234<!--TLTENE-->
```

For IBM Tealeaf users who do not have the appropriate permissions to view this data, the output looks like the following:

```
Your Social Security No. is <!--TLTENB-->@@@@@@@@@@<!--TLTENE-->
```

To enable this privacy rule, additional configuration is required. See "Privacy Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Enabling prebuilt rules

In the privacy configuration, IBM Tealeaf provides a number of pre-configured privacy rules, which can be modified and enabled to meet the requirements of your application

For more information modifying prebuilt rules, see "Privacy Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Related concepts:

"Web application development considerations" on page 58

Managing multiple levels of data security:

If multiple levels of encryption are required, you can add more bracketing comments, with each set representing a different security level.

For example:

```
Your Social Security No. is
<!--TLTTrustStandardB--><!--TLTTrustSilverB-->
  <!--TLTTrustGoldB-->123-12-1234<!--TLTTrustGoldE--><!--TLTTrustSilverE-->
  <!--TLTTrustStandardE-->
```

If the examples of blocking and encrypting data shown in “Pre-Built privacy rules IBM Tealeaf” on page 16 are required, then you can enable both blocking and encrypting in the response. The response data sent from your web application should be the following:

```
Your Social Security No. is
<!--TLTTrustStandardB--><!--TLTTrustSilverB--123-12-1234<!--TLTTrustSilverE-->
  <!--TLTTrustStandardE-->
```

For each of the security levels shown above, you should create a separate user group. Taking the example above, you might create the following groups, whose privacy rules will reference the listed bracketing tags:

Table 7. Groups to create and tags to reference

Group to create	Bracketing tags to reference
TealeafTrustStandard	<!--TLTTrustStandardB-->...<!--TLTTrustStandardE-->
TealeafTrustSilver	<!--TLTTrustSilverB-->...<!--TLTTrustSilverE-->
TealeafTrustGold	<!--TLTTrustStandardB-->...<!--TLTTrustStandardE-->

In implementation, you can determine what each of these groups is permitted to see.

TealeafTrustGold example

If only the Gold group is permitted to view the social security number, the example bracketing would be the following:

```
Your Social Security No. is
<!--TLTTrustGoldB-->123-12-1234<!--TLTTrustGoldE-->
```

In the TealeafTrustGold example, when the response was viewed by a TealeafTrustGold user, the user would see the actual social security number. For all non-TealeafTrustGold users, the number appears as several @ signs. Output:

```
Your Social Security No. is
<!--TLTTrustGoldB-->@@@@@@@@@<!--TLTTrustGoldE-->
```

TealeafTrustSilver example

If only the Gold and Silver groups are permitted to see the email address, then following the example, Gold users must also be included in the Silver security group.

Example HTML:

```
Your Email Address is
<!--TLTTrustSilverB-->test@example.com<!--TLTTrustSilverE-->
```

When the response is viewed by a Gold or Silver user, the user sees the actual email. All non-Gold or Silver users see the email address as several @ signs. Output:

```
Your Email Address is
<!--TLTTrustSilverB-->@@@@@@@@@@@@@@@@@<!--TLTTrustSilverE-->
```

Steps for enabling multiple levels of security

The following is a general approach to enabling multiple levels of data security in IBM Tealeaf, after applying the above standards in your web application. You are likely to need to modify these steps to meet your enterprise requirements.

1. Create IBM Tealeaf user groups for each type of security to employ. In the above example, you might create the following groups.
 - TealeafTrustStandard
 - TealeafTrustSilver
 - TealeafTrustGold

Note: If you are using tiered levels of security clearance, then users with the highest level of clearance must be included in all of the lower levels.

2. Bracket all sensitive data with the appropriate HTML comments, as shown in the examples.
3. Add IBM Tealeaf users to the appropriate groups.
4. Create and enable the privacy rules.
5. Deploy Privacy rules accordingly.
 - For IBM Tealeaf Passive Capture Application:
You can deploy privacy rules to be applied at the point of capture, which ensures that the sensitive data is never in an unprotected state in IBM Tealeaf. However, PCA privacy rules can impede performance of real-time processing.
 - For Windows Pipeline:
You can deploy privacy rules in the Windows pipeline on the Processing Server(s). However, using the Privacy session agent to block or encrypt data means that the data has not been secured until it passes through the PCA and the Transport Service to the Processing Server. IBM Tealeaf does employ application and database security methods to protect general access.
6. Test the availability of sensitive data for one user from each privacy group.

Related concepts:

“Designing privacy rules”

Designing privacy rules:

Several types of privacy rules inherently require high CPU/time requirements and can affect PCA performance

These types of rules include the following:

- Rules that search the entire response of the hit
- Rules that use regular expressions, instead of plain text to match patterns
- Rules that use both start and end pattern regular expressions
- Rules that encrypt data

Actions based on regular expressions are the most time/CPU intensive types of privacy actions. These rules should be avoided unless absolutely necessary.

Note: A single regular expression rule can be created, which causes the PCA to begin dropping traffic. Use them cautiously.

Some regular expressions must examine each character in the data, as well as a predefined preceding number of characters to match the pattern. The net result is the following computation for number of searches:

of characters in data * # of preceding characters to match

Including several rules with these types of regular expressions can raise the processing requirement for each hit to dangerously high levels. Whenever possible, use plain text to identify the start and end pattern of what you wish to block or encrypt.

In your rule configuration, you may also use the Stop Processing flag, which stops evaluating the data if you need to block only the first occurrence of the data. This flag allows privacy to stop searching after the rule has been triggered. For example, if you know that a customer's login ID is displayed only once on any page, then this flag is useful for limiting the processing required to block it. For text sequences occurring at the beginning of the response, the Stop Processing flag can make a considerable difference in processing requirements.

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

Related concepts:

"Managing multiple levels of data security" on page 17

Testing privacy:

PCA privacy management also enables testing to identify and eliminate potential processing issues.

For example, if you have a single page that requires encryption of a specific text sequence, you can create a test to trigger that rule only when the page's URL appears in the URL section of the request buffer. Privacy passes on evaluating any page that does not contain the URL, resulting in a dramatic decrease in processing overhead.

The test for the request value is far less intensive than the processing of the entire response of the hit for a text sequence. See "PCA Web Console - Rules Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

Windows pipeline privacy

You can use session agents in the Windows pipeline to manage data privacy.

The Windows pipeline enables multi-threaded processing, multi-instance processing, and sequenced processing that can accelerate the testing and application of privacy rules. However, any data that is not cleansed in the Windows pipeline is passed through the IBM Tealeaf server system and can be monitored by system administrators and users.

By design, the Windows pipeline executes individual session agents, such as the Privacy session agent, to manipulate session data via individual threads, which enables better management of the processing. In addition, the Health-Based Routing session agent enables the distribution of the pipeline data across multiple servers for superior load balancing.

The Windows pipeline does not restrict the number of instances of a session agent, so you can create sequenced processing. One instance of the Privacy session agent can execute an action to manipulate a value in the hit, and then another session agent downstream can act on the result of the first privacy rule. Sequenced processing enables sophisticated post-processing of hit data.

The following types of privacy actions should be handled in the Windows pipeline:

- Moving data values into the request buffer from the response buffer for event triggering.
- Concatenating existing values.
- Postprocessing request or response values.
- Blocking data that IBM Tealeaf users should not see, while enabling IBM Tealeaf administrators to access it.
- General listing, manipulation, or other alterations of data values in the hit.

For information about CX Pipeline Configuration, designing privacy in the Privacy Session Agent, and the Extended Privacy Session Agent, see the *IBM Tealeaf CX Configuration Manual*

Potential security risks of Windows pipeline privacy

When hits arrive faster than the Windows pipeline can process them, those hits are spooled to temporary files on the disk drive. Spooling hits to a disk drive might present a security risk, as potentially sensitive information in those hits becomes accessible to employees. Although spooled data is deleted after the hit is processed, there is a small window of time in which the data is exposed and accessible. Some company security policies might not permit this type of spooling storage or might require additional deletion or data scrubbing of the storage area to ensure that no sensitive data remains.

Dealing with performance implications of Windows pipeline privacy

When your privacy configuration results in too much CPU processing, you can try the following to reduce the load:

- Split your privacy rules between two or more instances of the privacy session agent by creating a child pipeline.
See "Session Router Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
- Split your privacy rules across more than one server by putting a privacy session agent into the pipeline on each server with a portion of your privacy actions in each.
See "Health-Based Routing (HBR) Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Applying privacy rules to Windows pipeline session agents

IBM Tealeaf provides a utility for testing and iterating on the development of privacy rules. While these rules are applied using the Windows pipeline session agents, the same rules can be applied through PCA privacy management. For more information, see "Privacy Tester Utility" in the *IBM Tealeaf CX Configuration Manual*.

Applying Privacy to Binary Response Formats

For release 8.2 or later, IBM Tealeaf supports the ability to apply privacy rules to the responses for the MSBIN1 format.

In the Windows pipeline, this data format can be decoded into clear text using the Inflate session agent. Later in the same pipeline, the Privacy session agent can be inserted, and privacy rules applied to the clear text version. To perform privacy, you create privacy rules in the same manner as any other text/xml post or text/xml response format.

For more information, see "Privacy Session Agent" and "Inflate Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Related tasks:

"Troubleshooting privacy performance in the PCA server" on page 129

On-demand privacy

On-demand privacy enables you to apply privacy to data that you wish to save for indexing for search and reporting purposes, yet you do not want IBM Tealeaf users to see the sensitive data in clear text during replay.

Using the same methods of configuration, privacy rules can be applied to data delivered through Search Server to IBM Tealeaf users requesting sessions for replay. These privacy rules use the same mechanisms as the PCA and Windows privacy methods.

Privacy rules for replay can be deployed through the Search Server configuration. For more information, see "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

RTV Privacy Tester

The IBM Tealeaf RealiTea Viewer includes an integrated Privacy Tester utility that enables you to apply a selected privacy configuration file to one or more sessions and then view the results through RTV. Privacy is applied to request, response, and replay views in RTV. See "RealiTea Viewer - Privacy Tester" in the *IBM Tealeaf RealiTea Viewer User Manual*.

How data sessionization works

Sessionization is analytic for measuring user behavior.

A sessionization operation consists of a group of interactions that take place on your web site during a certain time frame. A sessionization operation identifies your web site visitors browsing sessions by storing recorded events and grouping them from each user, based on the time-intervals between each and every event.

IBM Tealeaf can utilize one of many different methods for tracking sessions in the captured web traffic. Depending on the method deployed for your web application, IBM Tealeaf provides an appropriate means of identifying individual sessions.

For purposes of rebuilding accurate and complete sessions, this section lists the possible methods that can be deployed within IBM Tealeaf for sessionization.

The methods listed are in the order of desirability from an IBM Tealeaf perspective.

Data sessionization with IBM Tealeaf Cookie Injector

IBM Tealeaf Cookie Injector is a lightweight utility installed on your web server or application server.

It issues HTTP cookies containing sequential, unique identifiers.

IBM Tealeaf Cookie Injector is the preferred method for tracking sessions in IBM Tealeaf.

How IBM Tealeaf Cookie Injector works

For each request submitted to each hosting server, the Cookie Injector examines the request for the cookie or cookies that it is configured to issue. If the cookie is not present, a set-Cookie header is added to the response.

Things you should know about IBM Tealeaf Cookie Injector:

- IBM Tealeaf Cookie Injector supports the injection of multiple cookies.
- IBM Tealeaf Cookie Injector can be configured to inject cookies into the request header for visibility by downstream IBM Tealeaf components.
- IBM Tealeaf Cookie Injector can also add a response header that identifies the name of the server so that this information is available in hits captured by a IBM Tealeaf Passive Capture Application server.

Because the Cookie Injector is designed specifically for the IBM Tealeaf CX platform, it provides the best method for sessionization within IBM Tealeaf. However, since it requires installation of software within your web infrastructure, it is an optional component of the platform.

Data sessionization in the CX Passive Capture Application server

If for some reason you cannot install the IBM Tealeaf Cookie Injector, the next best method for applying data sessionization is through the IBM Tealeaf Passive Capture Application.

As a best practice, it is recommended to apply sessionization to captured data at the earliest possible juncture and before it reaches the Windows pipeline, where processing loads may already be significant.

At the point of capture, the CX Passive Capture Application can sessionize on a preexisting session cookie like JSESSIONID. The session cookie is transformed into a 32-byte string that is then injected into the IBM Tealeaf request buffer as a TLTSID.

Note: For a PCA-based sessionization solution to work, the specified cookie must be persistent for the entirety of the session.

By default, the downstream components of the IBM Tealeaf platform are configured to use the TLTSID for identifying sessions.

Passive Capture can sessionize on multiple fields in the request.

For more information, see "PCA Web Console - Pipeline Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

Data sessionization in the Windows pipeline

If you cannot apply sessionization through the PCA or have special requirements for sessionization, you can apply identifiers to the IBM Tealeaf request and manage sessionized traffic using the following Windows pipeline session agents.

How data sessionization in the Windows pipeline works

After data is captured by the IBM Tealeaf Passive Capture Application, the data is forwarded to the Processing Server, where it is passed through the Windows pipeline, a configured series of processing agents that review and modify the captured data based on user-defined rules.

For more information, see "CX Pipeline Session Agents" in the *IBM Tealeaf CX Configuration Manual*.

Windows pipelines are configured through the Pipeline Editor. See "TMS Pipeline Editor" in the *IBM Tealeaf cxImpact Administration Manual*.

Sessioning session agent

The Sessioning session agent can be added to your Windows pipeline to create a session identifier. When this session agent is invoked, it scans specified request field values for a matching session identifier. For example, in the session agent configuration, if you specify to look for JSESSIONID, KSESSIONID, and LSESSIONID as the sessioning parameters, the session agent first looks for JSESSIONID and then KSESSIONID and so on.

If a match is found in the request variables, the session agent applies a hash to create a 32-byte value that is inserted as the TLTSID value in the request. This method of sessionization is considered inferior to sessionization on the PCA because it burdens the pipeline and occurs later in the data evaluation process.

For more information, see "Sessioning Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Tealeaf sessioning session agent

For special situations, the Tealeaf Sessioning session agent can be deployed in your Windows pipeline to sessionize off specific user data in the request.

The data on which to sessionize is not limited to an HTTP Cookie. For example, you can perform sessionization based on a table lookup. This lookup examines one of 65,536 available buckets where hashes are stored and applies the hashes stored in the bucket until the session identifier is retrieved.

Hash values must be evenly distributed across the buckets, or hash search times may be impeded.

Depending on the traffic to your web site, you may need to store more than 10,000 session identifiers per day, which can impact the lookup performance. For more information, see "Tealeaf Sessioning Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Related tasks:

"Configuring session timeout settings" on page 112

Data sessionization through CX UI Capture

In web applications where no suitable cookie exists, you can use the IBM Tealeaf UI Capture for AJAX library create a unique cookie on the client. This unique cookie can then be used for sessionization.

If you have installed and deployed the IBM Tealeaf UI Capture for AJAX library, you can configure it to create a unique cookie in the visitor's web browser. This cookie is posted with each request or user interface event for capture by the IBM Tealeaf Passive Capture Application. Within IBM Tealeaf, you may use the captured cookie value to identify individual sessions.

For information about IBM Tealeaf UI Capture for AJAX library, see:

- "UI Capture FAQ" in the *IBM Tealeaf UI Capture for AJAX FAQ*
- "UI Capture for AJAX Reference" in the *IBM Tealeaf UI Capture for AJAX Guide*.

Internationalization support

IBM Tealeaf supports internationalization (the use of native non-English character sets) in monitored applications.

IBM Tealeaf supports ISO-8859-1 and UTF-8 encoding only. If needed, you can configure the PCA and the Windows pipeline to handle traffic that is in ISO-8859-1 and UTF-8 encoding.

IBM Tealeaf provides limited support for double byte encoding. For more information, see International character set support.

Prerequisites and recommendations for enabling internationalization support

IBM Tealeaf release and build requirements:

- Your IBM Tealeaf solution must using a 33xx build of the Passive Capture Installation.
- Your version of the IBM Tealeaf CX platform is Release 7.0 or later.

Operating system and SQL server version recommendations:

- Install the operating system and SQL Server versions that are designed for the appropriate native locale.
- Avoid installing US versions of these software packages and then manually applying settings for your locale

If you install using a native locale, you might have to perform additional configuration steps. For more information, see "Configuring Locale for the Portal".

International character set support

IBM Tealeaf provides limited support for UTF-8 and ISO-8859-1 character sets.

- Monitored web sites must use either ISO-8859-1 or UTF-8 character encoding and may not mix them.
- Some searching and filtering operations (such as privacy and filtering rules) may not correctly process non-ASCII data (characters outside of A-Z, a-z, 0-9, and common punctuation).

- Some user interface components do not respect cultural preferences for dates, times, calendars, and numbers.
- Some components of Tealeaf CX (including PCA, HBR, and RTV) do not properly process user data that cannot be represented in the native character encoding of the operating system.

For example, these components would not process Arabic data correctly on a French Windows system. Customers should match the native character encoding of their operating system with the encoding of the data being processed.

International character set support by feature provides information about international character set support available in IBM Tealeaf component features.

Table 8. International character set support by feature

IBM Tealeaf feature	Support notes
Passive Capture Application	<p>For IBM Tealeaf release 7.x or later and when a IBM Tealeaf Passive Capture Application build 33xx or later is deployed</p> <ul style="list-style-type: none"> • During capture, the PCA transforms the request body into UTF-8. The response body is unmodified. • There is a known issue in which the PCA fails to properly recognize UTF-8 encoding in data submitted from client frameworks, and the data may be mangled in the stored session, causing issues in eventing and search. <p>For more information, see (Release 8.5) https://community.tealeaf.com/display/tealeaf85/8.5+Release+Notes.</p> <p>For information about configuring international character set support, see</p>
Windows pipeline	<p>In the Windows pipeline, the Inflate session agent automatically scans for buffer variables that are inserted by the IBM Tealeaf Passive Capture Application to support internationalization. If these variables are not present, the Inflate session agent adds them. This feature requires no additional configuration.</p> <p>For information about the Inflate session agent, see "Inflate Session Agent" in the <i>IBM Tealeaf CX Configuration Manual</i>.</p>
Indexing	Indexes are generated in UTF-8 encoding.
Canister Storage	<p>Response bodies are converted to UTF-8 for indexing. You can search for text in the response.</p> <p>Since the request data has already been converted to UTF-8 encoding, it can be searched through standard search mechanisms.</p> <p>All IBM Tealeaf-internal data associated with a session is stored in UTF-8 encoding.</p>

Table 8. International character set support by feature (continued)

IBM Tealeaf feature	Support notes
Portal	<p>The Portal can be localized in UTF-8 to any supported character set.</p> <p>The Portal inherits its number formatting from the operating system settings of the machine hosting the Portal application.</p> <p>For more information, see "Configuring Locale Settings" in the <i>IBM Tealeaf cxImpact Administration Manual</i>.</p>
Search	Through the Portal, you can search for international characters in both the request and the response.
Event Manager	<p>In the definition for a hit attribute, you can select the encoding to expect in the request data. This value should match the value configured for the PCA.</p> <p>For more information, see "TEM Hit Attributes Tab" in the <i>IBM Tealeaf Event Manager Manual</i>.</p>
Browser Based Replay	The IBM Tealeaf Browser Based Replay support the UTF-8 character set.
RTV Replay	<p>RTV is not a UNICODE application and therefore does not support search for multi-byte characters.</p> <p>The IBM Tealeaf RealiTea Viewer can replay the Response page in whatever code format is specified in the Response HTTP Content-type/encoding directive.</p> <p>The request buffer continues to be replayed in English UTF-8 format.</p>
cxConnect for Data Analysis	IBM Tealeaf cxConnect for Data Analysis can be configured to extract extended characters in UTF-8 format.

Internationalization request variables

The following variables are inserted into the [env] section of the request:

```
REQ_BUFFER_ENCODING=UTF-8
REQ_BUFFER_ORIG_ENCODING=ISO-8859-1
RESP_BODY_ENCODING=utf-8
```

REQ_BUFFER_ENCODING

The encoding of the request. This value is always be set to UTF-8.

REQ_BUFFER_ORIG_ENCODING

Identifies the original encoding of the request.

RESP_BODY_ENCODING

Identifies the encoding in use in the body of the response.

Affected Systems

The following IBM Tealeaf systems use the variables described in "Internationalization request variables" to identify international encoding schemes and to interpret captured data accordingly.

- RealiTea Viewer

For more information, see "RealiTea Viewer - Replay View" in the *IBM Tealeaf RealiTea Viewer User Manual*.

- Tealeaf Event Manager

For more information, see "Tealeaf Event Manager" in the *IBM Tealeaf Event Manager Manual*.

- CX Browser Based Replay

For more information, see "CX Browser Based Replay" in the *IBM Tealeaf cxImpact User Manual*.

- Extended Privacy Session Agent

For more information, see "Extended Privacy Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

- Event evaluation

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

Configuring internationalization support for the IBM Tealeaf Passive capture application

You can configure encoding types and enable international request variable support for IBM Tealeaf Passive capture application.

Configuring encoding types for IBM Tealeaf Passive capture application:

You can configure encoding types for IBM Tealeaf Passive capture application.

Make sure you have access to the *IBM Tealeaf Passive Capture Application Manual*.

By default, the PCA is configured to assume that all captured traffic is encoded in ISO-8859-1 format.

To enable a character encoding other than the default:

1. Open the Web Console.

For more information, see "Passive Capture Configuration via Web Console" in the *IBM Tealeaf Passive Capture Application Manual*.

2. Click the Pipeline tab

For more information, see "PCA Web Console - Pipeline Tab in the *IBM Tealeaf Passive Capture Application Manual*.

3. In the Pipeline screen, select the proper encoding values for the following settings:

- Default request encoding
- Default response encoding

Note: To set the encoding to ISO-8859-1, select the value None.

4. Click **Save Changes**.

Enabling internationalization request variables for IBM Tealeaf Passive capture application:

You can enable internationalization request variables for IBM Tealeaf Passive capture application.

Make sure you have access to the *IBM Tealeaf Passive Capture Application Manual*.

To enable internationalization request variables:

1. Open the Web Console.
For more information, see "Passive Capture Configuration via Web Console" in the *IBM Tealeaf Passive Capture Application Manual*.
2. Click the Pipeline tab
For more information, see "PCA Web Console - Pipeline Tab in the *IBM Tealeaf Passive Capture Application Manual*.
3. In the Pipeline screen, click the **Enable I18N** check box
4. Click **Save Changes**.
All traffic subsequently captured by the IBM Tealeaf Passive capture application has several variables inserted into the request. For more information, see "Internationalization Request Variables" in "Internationalization support" on page 25.

If the IBM Tealeaf Passive Capture Application has not been configured to insert the request variables, an instance of Inflate session agent deployed in the Windows pipeline can perform the insertions.

The Inflate session agent is included in the default Windows pipeline and guarantees the data to be inflated for subsequent processing by the remainder of the pipeline and other Tealeaf components. When the Inflate session agent is deployed in the pipeline, no further configuration is required.

For more information, see "Inflate Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

For information about deploying session agents in the pipeline, see "TMS Pipeline Editor" in the *IBM Tealeaf cxImpact Administration Manual*.

Internationalization support for Portal

Through the Portal, you can configure locale settings for individual users and groups for language, time zone, and currency formatting.

For more information, see "Configuring Locale Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

Table 9 lists the portal pages that do not support native locale languages.

Table 9. Portal pages not supported for translation

Page Name	Menu	Description
Portal Management	Tealeaf > Portal Management	Page for configuring IBM Tealeaf servers, user administration, and administration-level settings. For more information, see "Managing Tealeaf Servers" in the <i>IBM Tealeaf cxImpact Administration Manual</i> .
Tealeaf Management System	Tealeaf > TMS	Centralized facility for managing configurations for IBM Tealeaf servers and components. For more information, see "Tealeaf Management System" in the <i>IBM Tealeaf cxImpact Administration Manual</i> .

Table 9. Portal pages not supported for translation (continued)

Page Name	Menu	Description
Event Manager	Configure > Event Manager	Portal facility for creating and configuring events and other event-related data objects. For more information, see "Tealeaf Event Manager" in the <i>IBM Tealeaf Event Manager Manual</i> .
IBM Tealeaf cxConnect for Data Analysis	Tealeaf > IBM Tealeaf cxConnect for Data Analysis	Tealeaf product to enable extraction of Tealeaf session and event data for third-party systems. For more information, see "cxConnect for Data Analysis Administration Manual" in the <i>IBM Tealeaf cxConnect for Data Analysis Administration Manual</i> .
IBM Tealeaf cxVerify	Tealeaf > IBM Tealeaf cxVerify	IBM Tealeaf product to enable extraction of Tealeaf sessions for storage in the enterprise data warehouse. Note: IBM Tealeaf cxVerify is a separately licensable component of the IBM Tealeaf CX system. For more information, see "cxVerify Administration Manual" in the <i>IBM Tealeaf cxVerify Administration Manual</i> .

Support for IPv6

IBM Tealeaf supports the capture, processing, and storage of IPv6, which results in these addresses being available for search, replay, and reporting.

Information in the following sections provide links to information about configuring IBM Tealeaf support for IPv6 throughout the solution.

Note: Hosting IBM Tealeaf servers using IPv6 addresses is not supported at this time.

IPv6 overview

Internet Protocol Version 6 (IPv6) is the next-generation method for specifying internet protocol addresses. IPv4, the previous version, enabled 32-bit IP addresses, which permitted the specification of 2³² addresses. All IPv4 address blocks have been assigned.

IPv6 enables the specification of 128-bit IP addresses, which supports the specification of 2¹²⁸ addresses. This expanded specification allows the use of device-specific IP addresses for the ever-growing set of connected devices. Other features:

- extra flexibility in allocating addresses
- efficiency for routing traffic
- eliminates the primary need for network address translation (NAT)

Note: The IBM Tealeaf Passive Capture Application can be configured to capture IPv6 addresses, mixed IPv6 and IPv4, and to translate IPv4 to IPv6 addresses.

Note: IPv4 addresses translated to IPv6 format cannot be inserted into the PCA Web Console, but you can insert these values in the ctc-conf.xml file. The PCA is

able to consume these addresses. For information about the methods for capturing and translating IP addresses, see Capturing and normalizing IP addresses.

IPv4 Format

The Internet Protocol specification originally formatted IP addresses in the following manner. This format was in universal use through 2009.

In the following example, each three-digit set of values is called an *octet*.
AAA.BBB.CCC.DDD:EEEE

The value EEEE represents a port number and is preceded by a colon (:).

IPv6 Format

An IPv6 address is represented as a sequence of eight groups of four hexadecimal digits. The groups are separated by colons (:).

The IPv6 format provides a much larger range of potential addresses than its predecessor, the IPv4.

IPv6, which is becoming more prevalent on the Internet. It is specified in the following format:

```
2001:0db8:85a3:0000:0000:8a2e:0370:7334(8080)
```

Hexadecimal digits are case-insensitive but should be represented in lower case for consistency.

Port numbers

Since the specification utilizes the colon (:) as an separator, the colon cannot be used as the port number marker, as in IPv4:
`https://1angle:19000`

Instead, the parentheses notation is used, as in the following example:
`2001:0db8:85a3:0000:0000:8a2e:0370:7334(8080)`

Note: The port number is included in parentheses (8080). For IPv6 addresses, searches using port numbers are not supported.

Simplifications

The full representation of eight 4-digit groups may be simplified by several techniques, eliminating parts of the representation.

Leading zeroes

Leading zeroes in a group can be omitted, but each group must contain at least one hexadecimal digit. In this way, the address can be simplified. For example, rather than using this port number:

```
2001:0db8:85a3:0000:0000:8a2e:0370:7334(8080)
```

You can use this port number:
`2001:db8:85a3:0:0:8a2e:370:7334`

Note: Notice the removal of two sets of leading zeroes and two sets of octets composed of zeros.

Groups of zeroes

One or more consecutive groups of zero values may be replaced with a single empty group using two consecutive colons (::).

- Substitution can only be applied once in an address, as multiple occurrences create an ambiguous representation.
- If more than one such substitution could be applied, the substitution that replaces the most groups should be used. If the number of groups is equal, then the leftmost substitution should be used.

With these rules, the example address is simplified even further, to this:

2001:db8:85a3::8a2e:370:7334

Special addresses

Table 10. Special addresses

Address name	Raw address	Shortened address
The localhost (loopback) address	0:0:0:0:0:0:0:1	::1
The IPv6 unspecified address	0:0:0:0:0:0:0:0	::

Source: <http://www.wikipedia.org>

Supported uses of IPv6

The following uses of IPv6 addresses are supported by IBM Tealeaf at this time:

- URLs in web application
- Enterprise-internal IPv6 addresses for servers

Restrictions for using IPv6

IBM Tealeaf components cannot be hosted on networks that use IPv6 addresses at this time.

Capturing and normalizing IP addresses

To make IPv6 addresses available for search, addresses of either IPv4 or IPv6 format must be captured and normalized to a format that is known to the IBM Tealeaf indexing and search processes.

IBM Tealeaf supports two methods of capturing and translating addresses:

- PCA
- Inflate session agent

PCA support for IPv6:

If you have deployed PCA Build 3501 or later, you can enable the capability to capture IPv6 addresses.

IPv4 addresses can be translated into an IPv6 format for indexing and search.

Note: If you cannot upgrade to PCA Build 3501 or later at this time, you must deploy the Inflate session agent in every Windows processing pipeline in order to support indexing and search of IPv6 addresses.

IPv6 addresses in the Web Console

Beginning in PCA Build 3600, you can configure the Web Console to accept IPv6 addresses by default.

- See "Passive Capture Configuration via Web Console" in the *IBM Tealeaf Passive Capture Application Manual*.

Prior to PCA Build 3600, IPv6 addresses could not be entered through the PCA Web Console.

Data insertions into the request - IPv6 format

When IPv6 capture is enabled and IPv6 addresses are detected in the capture stream, the following variables are inserted into the [env] section of the request:

```
...
IPV6_XLAT=False
IPV6=True
...
REMOTE_ADDR=fe80::20b:dbff:fe93:a462
LOCAL_ADDR=fe80::213:72ff:fe67:ed26
SERVER_NAME=fe80::213:72ff:fe67:ed26
IPV6_REMOTE_ADDR=FE80:0000:0000:0000:020B:DBFF:FE93:A462
IPV6_LOCAL_ADDR=FE80:0000:0000:0000:0213:72FF:FE67:ED26
IPV6_SERVER_NAME= fe80::213:72ff:fe67:ed26
...
```

IPV6_XLAT

When IPv6 is set to True, this option, if True, indicates whether IP addresses inserted into the request contain IPv4 addresses and should be translated.

IPV6 Indicates if captured traffic is IPv6, if True.

REMOTE_ADDR

The raw IP address, as captured, for the remote address may be in IPv6 or IPv4 format.

This value may be inserted by the PCA.

LOCAL_ADDR

The raw IP address, as captured, for the local address may be in IPv6 or IPv4 format.

This value may be inserted by the PCA.

Note: This value may be compressed for IPv6 format.

SERVER_NAME

Existing field name can now accept IPv6 data.

SERVER_NAME is not indexed.

IPV6_REMOTE_ADDR

The REMOTE_ADDR value rendered in IPv6 uncompressed format.

This value may be inserted by the PCA.

IPV6_LOCAL_ADDR

The LOCAL_ADDR value rendered in IPv6 uncompressed format.

This value may be inserted by the PCA.

IPV6_SERVER_NAME

New field name is used to store **SERVER_NAME** value in uncompressed IPv6 format

IPv6 Translate mode

In IPv6 Translate mode, the PCA translates IPv4-native addresses into a format that is readable using components on the IBM Tealeaf Windows servers. The PCA inserts the following fields in the request. In addition to the above fields, the original values for the following are inserted:

- IPV6_REMOTE_ADDR_ORIG
- IPV6_LOCAL_ADDR_ORIG
- IPV6_SERVER_NAME_ORIG

Example:

```
IPV6_XLAT=True
IPV6=True
REMOTE_ADDR=254.147.164.98
LOCAL_ADDR=254.103.237.38
SERVER_NAME=254.103.237.38
?
IPV6_REMOTE_ADDR=0000:0000:0000:0000:0000:FFFF:FE93:A462
IPV6_LOCAL_ADDR=0000:0000:0000:0000:0000:FFFF:FE67:ED26
IPV6_SERVER_NAME=0000:0000:0000:0000:0000:FFFF:FE67:ED26
?
IPV6_REMOTE_ADDR_ORIG=FE80:0000:0000:0000:020B:DBFF:FE93:A462
IPV6_LOCAL_ADDR_ORIG=FE80:0000:0000:0000:0213:72FF:FE67:ED26
IPV6_SERVER_NAME_ORIG=FE80:0000:0000:0000:0213:72FF:FE67:ED26
```

IPV6_REMOTE_ADDR_ORIG

Contains the original IPv6 address for the **REMOTE_ADDR** before it is translated.

IPV6_LOCAL_ADDR_ORIG

Contains the original IPv6 address for the **LOCAL_ADDR** before it is translated

IPV6_SERVER_NAME_ORIG

Contains the original IPv6 address for the **SERVER_NAME** before it is translated.

Data insertions into the request - IPv4 format

If the PCA detects IPv4 addresses, the following fields are inserted in the request.

```
IPV6_XLAT=False
IPV6=False
REMOTE_ADDR=10.10.20.105
LOCAL_ADDR=152.163.17.33
SERVER_NAME=152.163.17.33
IPV6_REMOTE_ADDR=0000:0000:0000:0000:0000:FFFF:0A0A:1469
IPV6_LOCAL_ADDR=0000:0000:0000:0000:0000:FFFF:98A3:1121
IPV6_SERVER_NAME=0000:0000:0000:0000:0000:FFFF:98A3:1121
```

For indexing purposes, the IPv4 source addresses are converted into an IPv6 format and inserted into the following destination variables in the request:

Table 11. Data insertions into the request

IPv4 source	Source example	IPv6 destination	Destination example
REMOTE_ADDR	10.10.20.105	IPV6_REMOTE_ADDR	0000:0000:0000:0000:0000:FFFF:0A0A:1469
LOCAL_ADDR	152.163.17.33	IPV6_LOCAL_ADDR	0000:0000:0000:0000:0000:FFFF:98A3:1121
SERVER_NAME	152.163.17.33	IPV6_SERVER_NAME	0000:0000:0000:0000:0000:FFFF:98A3:1121

For descriptions of these fields, see "Data insertions into the request - IPv6 format" on page 33.

For an example of translating an IPv4 address to an IPv6 address, see *IPv4 translated to IPv6 address example* in Indexing for IPv6.

Related concepts:

"Inflate session agent support for IPv6"

Inflate session agent support for IPv6:

If the PCA cannot be upgraded to a IPv6-supported build at this time, you must deploy the Inflate session agent to insert the appropriate values in the request for indexing and search of IPv6 addresses.

When hits are passed through pipelines containing the Inflate session agent, the following IPv6-compatible fields are inserted into the request, if they are not already present:

- IPV6_REMOTE_ADDR
- IPV6_LOCAL_ADDR

For information about insertions performed by the Inflate session agent, see "Data insertions into the request - IPv6 format" in the topic *PCA support for IPv6*.

Also, for information about configuring the Inflate session agent, see "Inflate Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Related concepts:

"PCA support for IPv6" on page 32

Processing server

Session data containing IPv6 addresses is passed through the Processing Server without interruption. These hits are processed and stored transparently.

Indexing and search services have been updated to support IPv6, as described in the following sections.

Indexing for IPv6

During the indexing process, most values in the [env] section of the request are indexed, including the IP address values.

For more information on indexing in general, see "Configuring CX Indexing" in the *IBM Tealeaf CX Configuration Manual*.

Indexing for IPv4

During the indexing process, the indexer reviews and normalizes the remote_addr, local_addr, referrer, and T1tStsIPAddr values that are IP address. Normalization for IPv4 involves:

- Zero-padding all octets so that each is three digits.
- Removing the dots between octets, so that the remaining value is a string of twelve digits.

For example, an IPv4 address of 1.12.123.4 is normalized as 001012123004 for indexing purposes.

Note: Values in the [appdata] section of the request that contain IP addresses are indexed as text values, instead of IP addresses.

Searching for IPv4 addresses through the Portal requires entering the IP address in the above normalized format. For more information, see Search through the Portal.

Indexing for IPv6

If a hit is an IPv4 hit, the indexer populates the [env] IPv6 fields with IPv4 equivalents. In the IPv6 address, the last two groups are the hexadecimal equivalent of the IPv4 address, and the first 6 groups are all set to be zeroes, as shown in the following example:

```
REMOTE_ADDR=152.163.17.33
IPV6_LOCAL_ADDR=0000:0000:0000:0000:FFFF:98A3:1121
```

Using this example, the indexer adds the following IPv6-related value as a search keyword:

```
ipv6_remote_addr
```

To retain backward compatibility and for legacy purposes, the fields remote_addr, local_addr and tltsipaddr are indexed if the addresses are IPv4 values.

Additionally, IPv4 addresses are spread across more groups, as shown in the following example:

**Source Item
Data**

```
IPV6_LOCAL_ADDR request variable  
0000:0000:0000:0000:FFFF:98A3:1121
```

```
ipv6_remote_addr index variable  
0000:0000:0000:0000:0098:00A3:0011:0021
```

By splitting the data in the final two groups into four different data values, it is possible to enable searching across data ranges through the Portal and RTV. See Search for IPv6.

IPv4 translated to IPv6 address example

For indexing purposes, the PCA inserts IPv4 addresses in a new, standards-based format that is compatible with IPv6. This data insertion in the request enables IBM Tealeaf indexing and search to operate with minimal changes while supporting both IP formats.

Note: These address formats are used internally by IBM Tealeaf to support search and indexing.

In the following example, the PCA is configured to capture IPv4 addresses only.

- REMOTE_ADDR and LOCAL_ADDR addresses are written in source IPv4 format.
- The new fields:
 - IPv6_REMOTE_ADDR
 - IPv6_LOCAL_ADDR

- IPv6_SERVER_NAME
- In this mode, the source addresses are written in IPv6 format using the last eight digits in hexadecimal format, with the leading five groups written as zeroes.
- Addresses are formatted with a prefix FFFF group for the concluding two groups, which represent an IPv4 address in IPv6 hexadecimal notation.
This new format is used to identify an IPv4 address translated in the IPv6 format.

[env] values

```

IPV6_XLAT=False
IPV6=False
REMOTE_ADDR=10.10.20.105
LOCAL_ADDR=152.163.17.33
SERVER_NAME=152.163.17.33
IPV6_REMOTE_ADDR=0000:0000:0000:0000:0000:FFFF:0A0A:1469
IPV6_LOCAL_ADDR=0000:0000:0000:0000:0000:FFFF:98A3:1121
IPV6_SERVER_NAME=0000:0000:0000:0000:0000:FFFF:98A3:1121

```

When the above values are indexed, IPv4 values are zero-padded so that all IP addresses are of the same length.

For IPv4 addresses, the last 4 groups are indexes to the last 4 words of the IPv6 index variable, which enables range search on the IP address. For more information about example searches, see Example Advanced Searches.

Table 12. Examples of indexed values - IPv4 translated to IPv6 address

Index variable	Indexed value
remote_addr	010010020105
local_addr	152163017033
tlststipaddr	010010020105
ipv6_remote_addr	0000:0000:0000:0000:000A:000A:0014:0069
ipv6_local_addr	0000:0000:0000:0000:0098:00A3:0011:0021

IPv6-only compressed example

In the following example, the IP addresses are submitted as IPv6 only. The legacy fields generated by the indexer (remote_addr, local_addr, and tlststipaddr) are therefore blank.

[env] values

```

IPV6_XLAT=False
IPV6=True
REMOTE_ADDR=fe80::20b:dbff:fe93:a462
LOCAL_ADDR=fe80::213:72ff:fe67:ed26
IPV6_REMOTE_ADDR=FE80:0000:0000:0000:020B:DBFF:FE93:A462
IPV6_LOCAL_ADDR=FE80:0000:0000:0000:0213:72FF:FE67:ED26

```

Table 13. Examples of indexed values - IPv6-only compressed

Index variable	Indexed value
remote_addr	empty
local_addr	empty
tlststipaddr	empty
ipv6_remote_addr	FE80:0000:0000:0000:020B:DBFF:FE93:A462

Table 13. Examples of indexed values - IPv6-only compressed (continued)

Index variable	Indexed value
ipv6_local_addr	FE80:0000:0000:0000:0213:72FF:FE67:ED26

IPv6-only uncompressed example

In this example, all IP addresses are in uncompressed IPv6 format. The legacy index fields are thus blank.

[env] values

```
IPV6_XLAT=False
IPV6=True
REMOTE_ADDR=1234:5678:90ab:cdef:2123:4321:fd da:12cd
LOCAL_ADDR=ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
IPV6_REMOTE_ADDR=1234:5678:90AB:CDEF:2123:4321:FD DA:12CD
IPV6_LOCAL_ADDR=FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF
```

Table 14. Examples of indexed values - IPv6-only uncompressed

Index variable	Indexed value
remote_addr	empty
local_addr	empty
t1tstsipaddr	empty
ipv6_remote_addr	1234:5678:90AB:CDEF:2123:4321:FD DA:12CD
ipv6_local_addr	FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF

IPv4 translated into IPv6 address example

In this example, the PCA has been configured to translate IP addresses from IPv4 format to IPv6 format. The configuration settings are both true:

```
IPV6_XLAT=True
IPV6=True
```

The PCA translates the captured address to IPv4 format and inserts it into the REMOTE_ADDR and LOCAL_ADDR values. These values are inserted as the last two groups in the IPv6 values in the [[env] section: IPV6_REMOTE_ADDR and IPV6_LOCALADDR.

[env] values

```
IPV6_XLAT=True
IPV6=True
REMOTE_ADDR=254.147.164.98
LOCAL_ADDR=254.103.237.38
IPV6_REMOTE_ADDR=0000:0000:0000:0000:0000:FFFF:FE93:A462
IPV6_LOCAL_ADDR=0000:0000:0000:0000:0000:FFFF:FE67:ED26
```

Table 15. Examples of indexed values - IPv4 translated into IPv6 address

Index variable	Indexed value
remote_addr	254147164098
local_addr	254103237038
t1tstsipaddr	254103237038
ipv6_remote_addr	0000:0000:0000:0000:00FE:0093:00A4:0062

Table 15. Examples of indexed values - IPv4 translated into IPv6 address (continued)

Index variable	Indexed value
ipv6_local_addr	0000:0000:0000:0000:00FE:0067:00ED:0026

Search for IPv6

All search fields that are configured to search for IP addresses, such as Client IP or Server IP Address, can be used to search for addresses in IPv4 or IPv6 format.

Search through the Portal:

When searching through the portal for IP addresses, specific search keyword fields are used.

The following table lists the search fields and the search keyword associated with the search field.

Table 16. Search fields and associated key words

Search field	Associated key word
Client IP address	ipv6_remote_addr
Server IP address	ipv6_local_addr

Note: If you have upgraded your IBM Tealeaf solution from a pre-Release 8.4 build, the Portal automatically checks the upgrade data and adjusts searching for the above search fields to use the appropriate search keyword, using the legacy versions if the search dates include days before the upgrade was completed.

For more information on searching for IPv6 addresses through the Portal, see "Searching for IP addresses" in the *IBM Tealeaf cxImpact User Manual*.

Search in RTV:

Through the RTV Search Builder, you can construct searches for IP address fields using IPv4 or IPv6 format.

Configuring RTV search fields

Before you search for IP addresses through RTV, you must add the appropriate fields to your search templates to enable searching for addresses. See "IP address field configuration" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Searching for IP addresses through RTV

See "Searching for IP addresses" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Example Advanced Searches

The following searches can be configured using advanced search methods.

- In the Portal, you can specify these types of searches in the All Text search field.
- In RTV, these searches are specified through the Advanced Search Builder tab.

Table 17. Examples of advanced searches for IPv4

To look for IPv4 address	Enter the following in an IPv4 address field
10.10.20.100	ipv6_remote_addr contains: ?0000 0000 0000 0000 000A 000A 0014 0064?
IP?s between 10.10.20.100 thru 10.10.20.200	ipv6_remote_addr contains: ?0000 0000 0000 0000 000A 000A 0010 0010 0014 0064~00C8?
IP?s between 10.10.20.100 thru 10.50.20.200	ipv6_remote_addr contains: ?0000 0000 0000 0000 000A 000A~0032 0020 0064~00C8?

Table 18. Examples of advanced searches for IPv6

To look for IPv6 address	Enter the following in an IPv6 address field
FE80:0000:0000:0000:020B:DBFF:FE93:A462	ipv6_remote_addr contains: ?FE80 0000 0000 0000 020B DBFF FE93 A462?
IP?s between :FE80:0000:0000:0000:020B:DBFF:FE93:A462 thru :FE80:0000:0000:0000:020B:DBFF:FE93:FFFF	ipv6_remote_addr contains: ?FE80 0000 0000 0000 020B DBFF FE93 A462~FFFF?
IP?s between :FE80:0000:0000:0000:020B:DBFF:FE93:A462 thru :FE80:0000:0000:0000:FFFF:DBFF:FE93:FFFF	ipv6_remote_addr contains: FE80 0000 0000 0000 020B DBFF~FFFF FE93 A462~FFFF?

Search Keywords

The following table lists and describes the search keywords.

Table 19. Search fields and associated key words

Search Keyword	Description
tltsipaddr	Legacy IPv4 client IP address.
remote_addr	Legacy IPv4 client IP address.
local_addr	Legacy IPv4 server IP address.
ipv6_remote_addr	IPv6 client IP address.
ipv6_local_addr	IPv6 server IP address.

These search keywords are available for constructing advanced search queries through the Portal and RTV, as well as the Portal Search API. For more information:

- See "Search Keywords" in the *IBM Tealeaf cxImpact User Manual*.
- See "Search through the Portal" , for more information about searching for sessions through the Portal
- See "cxReveal Web Services API" in the *IBM Tealeaf cxReveal API Guide*, for more information about searching using the Portal API.

Search Templates

Search templates that reference any of the above search keywords enable the searching of IPv4, IPv6, or mixed addresses. See "Configuring Search Templates" in the *IBM Tealeaf cxImpact Administration Manual*.

Reporting for IPv6

IPv6 data is stored in the same databases used by IPv4.

Reporting objects may need to be reviewed and modified to accommodate the IP address configuration in your network environment.

Report data:

After IPv6 data capture has been deployed, values in the reporting data may be changed, depending on your configuration options.

Dimensions

The Server dimension provided by IBM Tealeaf is defined to capture values from the REMOTE_ADDR request variable. If you have enabled the X-FORWARDING feature in the PCA to pull data from a different request field containing IPv6-formatted addresses, the content of your Server dimension is changed. See "X-forwarding" in the *IBM Tealeaf Passive Capture Application Manual*.

Until the dimension data is purged, it may be difficult to produce consistent reporting data if IPv4 addresses are being translated into IPv6 format.

Note: After you have enabled capture of IPv6 data and have enabled IPv4 translation into IPv6, do the following:

- If possible, consider purging any dimensions that capture IP addresses of their data. See "Purging Dimension Data" in the *IBM Tealeaf Event Manager Manual*.
- If dimension purge is not acceptable, all report users should be informed of the date in which the IPv6 switch was enabled. Dimension data before and after the switch will differ for the same values.

Other IP-related event objects

Any objects that you have created to capture IP addresses should be reviewed to verify that they are sourced from the proper hit data. See "Tealeaf Event Manager" in the *IBM Tealeaf Event Manager Manual*.

Database

IPv6 and IPv4 (if present) addresses are stored in the appropriate IBM Tealeaf database fields. Access to the data is transparent to the user.

IP data is stored in the following databases:

- Reports
- System (if your internal network uses IPv6)
- Visitor Staging
- Visitor Reports
- RSE

Replay for IPv6

Through Browser Based Replay (BBR) and the desktop IBM Tealeaf RealiTea Viewer (RTV) application, the replay of sessions containing IPv6 addresses is unaffected and transparent to the user.

- See "CX Browser Based Replay" in the *IBM Tealeaf cxImpact User Manual*.

- See "RealiTea Viewer - Replay View" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Client-Side capture

IBM Tealeaf provides a DLL plugin for use with Fiddler to capture sessions for replay through the local client. This tool is provided for proof-of-concept work in environments where IBM Tealeaf is not installed or available.

Client-side capture supports the mapping of detected addresses to IPv6 format.

See "Using Client-Side Capture for Fiddler" in the *IBM Tealeaf Client-Side Capture Manual*.

cxConnect for IPv6

IBM Tealeaf cxConnect for Data Analysis supports the transparent export of IPv6 addresses for use in third-party systems.

For more information, see the *IBM Tealeaf cxConnect for Data Analysis Administration Manual*.

Chapter 2. Planning the installation

Planning for IBM Tealeaf before introducing its software into your enterprise information system helps ensure that the system you implement meets your needs.

This section describes how to plan for IBM Tealeaf, including assessing your current environment and your business requirements.

Planning activities include:

- Learning about IBM Tealeaf installation topologies.

Many factors go into determining how to use the servers, including traffic (both hits/sec and average page size), count and type of events, and amount of retained data.

By learning about the supported IBM Tealeaf installation topologies, you will have a better understanding of which topology is best suited for your intended use of the product.

- Making sure your environment meets the hardware and software requirements required for your IBM Tealeaf solution.
- Determining your how to implement high-availability for your IBM Tealeaf installation.

Installation topologies for IBM Tealeaf CX

An installation topology is the physical and logical layout of interconnected servers on which IBM Tealeaf CX components are installed.

You can install IBM Tealeaf CX on a single server or across multiple servers. The decision on which topology to implement depends your existing IT infrastructure, the volume of data to be captured, and your IBM Tealeaf design and usage model.

If you are unsure about the topology to use for your IBM Tealeaf installation, consult with IBM Tealeaf professional services.

After you have installed and configured IBM Tealeaf CX and IBM Tealeaf products, you can manage the deployment environment for the IBM Tealeaf CX components a single server or as a collection of servers.

Related tasks:

“Identifying available resources” on page 65

Topology variations

The physical architecture of your IBM Tealeaf installation environment can be distributed across multiple servers.

Figure 3 on page 44 illustrates a three server configuration for IBM Tealeaf. The illustration shows the various processing components on the servers and illustrates the functional relationship between the servers at run time.




Figure 3. IBM Tealeaf CX components processing

Allocating IBM Tealeaf functions in a multi-server environment

The distributed architecture of the IBM Tealeaf solution provides you with options for distributing component functionality to meet varying data volume, network architecture, and business requirements.

For small-volume deployments, you can deploy the Storage Server on the same physical machine as the Processing Server. While deploying the Processing and Storage Servers on a single machine does not provide the benefits of two separate machines, it provides greater advantages for backup and search performance than running an entire Tealeaf system on a single Processing Server.

For high-volume and multi-data center environments, the Tealeaf architecture can be scaled horizontally and distributed to meet real-time processing loads while still providing a central location for session and aggregate data storage and a single IBM Tealeaf user experience.

The following scenario examines a high-volume, geographically distributed Web application where the production support team is concentrated in a fourth, separate operations location that contains a SQL Server farm and SAN. This example illustrates how the IBM Tealeaf CX architecture can be flexibly deployed while still delivering optimal performance and data security:

- Data Center 1: PCA server and Processing Server
- Data Center 2: PCA server and Processing Server
- Data Center 3: PCA server and Processing Server
- Operations Center:
 - Storage Server is configured to store session and index archives on SAN.
 - The Portal Web application is also deployed on the Storage Server.

- The Report Server is deployed on the SQL Server farm.
- Additional Processing Servers can be deployed if the data volumes exceed the capacity of a single Processing Server.

Each Processing Server is configured to store sessions for the current day. The aggregate reporting data is stored in the Report Server deployed on the SQL farm. Using the Portal and IBM Tealeaf RealiTea Viewer, Tealeaf users are presented a single, unified view of live activity, monitoring, search, session replay, and reporting across all Processing and Storage server instances.




Figure 4. Deployment Variations

The distributed architecture of the Tealeaf solution also permits the prescribed architecture to be modified to meet varying data volume, network architecture, and business requirements.

Related concepts:

“IBM Tealeaf CX component architecture” on page 5

Related tasks:

“Verifying system startup” on page 96

“Verifying system start up for single server topology” on page 96

“Verifying system start up for a multiserver installation environment” on page 97

Support for virtualized server environments

IBM Tealeaf supports the installation of a subset of Tealeaf components into a server virtualization solution, such as those offered by IBM, VMWare, Citrix, or Microsoft.

For all virtualized server environments, it is recommended to use a dedicated environment due to the unpredictable and volatile nature of the system resource requirements necessary to support data processing in the IBM Tealeaf system.

Resources must be allocated to meet the minimum system requirements specified by the Tealeaf sizing recommendation provided by Tealeaf Professional Services.

Support restrictions

Tealeaf is not responsible for supporting the identification or resolution of performance issues when Tealeaf is deployed in a shared VM environment.

For any Tealeaf component that is deployed in a virtual server environment, Tealeaf software in a non-shared virtual server environment or non-virtual server environment.

Tealeaf components that support virtualization

Table 20. Tealeaf components that work in a virtualized environment

Component	Virtual machine support (Yes or No?)	Notes
Passive Capture Appliance	No	The IBM Tealeaf Passive Capture Application is a high-performance real-time processing engine that requires a dedicated server.
Processing Server	Yes	Due to high processing and I/O loads, at least 8 cores and 16 GB of RAM must be dedicated to the instance that runs the Tealeaf software.
Other Windows-based Tealeaf Servers	Yes	Tealeaf supports Microsoft SQL Server 2008 R2, 2012, and 2014 in virtual server environments that adhere to Microsoft support policies, guidelines and restrictions.

System requirements for installing IBM Tealeaf

To install IBM Tealeaf, make sure the operating system, hardware, and software requirements are met.

Operating system requirements

Make sure that the machines on which you install IBM Tealeaf CX meet the operating system requirements documented here.

Operating system requirements for the Report server

The time zone for the Windows machine hosting the IBM Tealeaf Report Server must be configured to match the IBM Tealeaf system time zone.

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

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

Operating system requirements for running IBM Tealeaf IBM Tealeaf installer

One of the following operating systems must be installed prior to running the IBM Tealeaf IBM Tealeaf installer:

- Microsoft Windows Server 2008 64-bit
IBM Tealeaf supports SQL Server 2008 Standard or Enterprise (32-bit or 64-bit versions are supported. The R2 version is recommended.), SQL Server 2012, and SQL Server 2014.

Note: For Windows 2008 (non-R2), all Windows updates must be applied.

- Microsoft Windows Server 2012 64-bit

If the system has more than 32GB of RAM, then you must install an Enterprise version of one of the supported operating system versions.

For optimal performance, Tealeaf software should be installed on dedicated hardware. Tealeaf does provide limited support for virtual machine (VM) environments.

.NET Framework requirements

Tealeaf requires that Microsoft .NET Framework be installed on all IBM Tealeaf IBM Tealeaf servers, including the SQL Server system.

- General requirements
 - Tealeaf supports .NET Framework V 4.5
- Database Server
 - The following .NET Framework requirements pertain to specific versions of SQL Server:
 - SQL Server 2014
 - SQL Server 2012: V.35
 - SQL Server 2008: V.35
- Windows Server

- The following versions of Windows Server are compatible with this version of Tealeaf:
 - Windows Server 2008
 - Windows Server 2012

Note: The default Windows Server installation also installs the .NET Framework. Please verify that the installed version is one of the versions supported by Tealeaf.

When you configure Windows Server, you must register .Net Framework with IIS to enable it to accept the IBM Tealeaf.

Role Services requirements for IIS 7 on Windows Server 2008

If you using IIS 7 for Tealeaf on a Windows Server 2008 system, please verify that the following Role Services have been enabled for IIS7 through the operating system:

- Static Content
- Default Document
- ASP.NET
- .NET Extensibility
- ISAPI Extensions
- ISAPI Filters
- IIS Management Console

Supported browsers

On each desktop that will access the IBM Tealeaf Portal, one of the following Web browser versions must be installed:

- Microsoft Internet Explorer Web browser version 7 or later
- Mozilla Firefox 9 or later
- Google Chrome 17 or later

IBM Tealeaf cxOverstat requires web browser support for HTML5.

IBM Tealeaf Portal does not support access from mobile browsers. See "Logging in to the Tealeaf Portal" in the *IBM Tealeaf cxImpact User Manual*.

If you have licensed IBM Tealeaf cxOverstat , the list of supported browser versions is smaller, as HTML5 support is required. See "cxOverstat Installation and Configuration" in the *IBM Tealeaf cxOverstat User Manual*.

Keeping the operating system and software add-ons up to date

After you install IIS, .Net Framework, and Microsoft SQL Server install the latest patches for each.

Some guidelines include:

- Before installing IBM Tealeaf CX and enabling IBM Tealeaf products, run Windows Update to make sure all of the latest patches have been applied. Run Windows Update until there are no more critical updates to apply.
- Turn off Windows Update while installing and configuring IBM Tealeaf CX and the IBM Tealeaf products.

Windows updates might restart your server automatically and that might cause problems during the installation.

Consult with your IT team before turning off Windows Update to make sure it is OK.

Related tasks:

“Identifying available resources” on page 65

“Installing .NET framework” on page 67

Software requirements

Make sure your installation environment meets the software requirements documented here.

Required software components

Install the following software components in the sequence in which they are listed prior to running IBM Tealeaf CX installer:

Note: If you have already installed some or all of these components, please install the remaining components and run the command listed below these items.

1. Internet Information Services (IIS) World Wide Web Server version 6 or version 7

If you are installing IIS into Windows Server 2008, you must assign a minimum set of Role Services to IIS. For more information see *Role Services requirements for IIS 7 on Windows Server 2008* in “Operating system requirements” on page 47.

If you are running the IBM Tealeaf CX solution on a multi-CPU server hosting multiple web applications, create a new application domain that is assigned to IBM Tealeaf only. When it is assigned to a single domain, the application can use an entire CPU or core for its purposes, thus enhancing performance.

The IBM Tealeaf CX solution can be installed in any virtual directory, regardless of the website to which it may be attached, as long as the application domain is running .NET applications only. This deployment option is useful if the Portal is accessed using `http://TealeafPortal.MyCompanyDomain.com/Portal` or a similar domain-dependent construction. Please note the `/Portal` at the end. If you must reference the Portal using `http://TealeafPortal.MyCompanyDomain.com`, please contact Tealeaf Support for guidance.

2. Microsoft .NET Framework

Microsoft .NET Framework must be installed on all servers used by IBM Tealeaf. See *.NET Framework Requirements* “Operating system requirements” on page 47.

3. Microsoft SQL Server 2008 or SQL Server 2012

The database server may be the same machine as the Portal Server or a remote machine that can be accessed by the Portal Server. The database server and SQL Server software must be set up and configured prior to running the IBM Tealeaf CX Installer.

If you are installing IBM Tealeaf cxConnect for Data Analysis, a local SQL instance is strongly recommended for the IBM Tealeaf cxConnect for Data Analysis database. See “cxConnect Installation” in the *IBM Tealeaf cxConnect for Data Analysis Administration Manual*.

R2 is the recommended version for SQL Server 2008.

As of Release 8.1, Tealeaf requires 64-bit machines for all Windows-based servers hosting Tealeaf software. If your Release 8.1 or later installation of Tealeaf databases is hosted on a machine that also hosts other Tealeaf software, such as the Reporting Server, you may use a supported 32-bit version of SQL Server. However, Tealeaf recommends 64-bit versions of SQL Server.

If you are upgrading from Tealeaf 6.2 or earlier, you must upgrade to SQL Server 2008 or later. See "Upgrading SQL Server for Tealeaf" in the IBM Tealeaf Databases Guide.

Software requirements for the Passive Capture Application server

The software requirements for the PCA can be found in the *IBM Tealeaf PCA Manual*.

Software requirements for the Processing server

The server used for processing capture data must meet the following software requirements:

- Microsoft Windows Server 2008 (64bit)
- Administrator rights on the server
- .Net 4.5 installed and registered
- IIS Installed

Software-related actions:

- Turn off Anti-virus
- Turn off Windows Firewall

Software requirements for the SQL server

The SQL server must meet the following software requirements:

- Microsoft SQL Server 2008, 2012, or 2014

Software-related actions:

- Do not configure virtual ports
- Must be configured for mixed-mode security
- Turn off Anti-virus
- Turn off Windows Firewall

Related tasks:

"Identifying available resources" on page 65

"Installing and configuring prerequisite software" on page 66

Software requirements for the data preagggregator

The following software is required to install and operate the data preagggregator.

- Microsoft Windows Server 2008 (64-bit)

Note: Administrator rights to the server are required.

- .NET 4.5

Hardware requirements

For optimal performance, install IBM Tealeaf on dedicated hardware.

Processor requirements

As of release 8.1 of IBM Tealeaf, 64-bit processors are required for all Windows-based servers hosting the IBM Tealeaf Processing Server, Report Server or Archiving component.

For the IBM Tealeaf Passive Capture Application (which runs on Linux only) and the IBM Tealeaf CX RealiTea Viewer replay client (installed on desktop systems) both 32-bit processors and 64-bit processors are supported.

Note: If your Release 8.1 or later installation of IBM Tealeaf databases is hosted on a machine that also hosts other IBM Tealeaf software, such as the Reporting Server, you may use a supported 32-bit version of SQL Server. However, a 64-bit version of SQL Server is recommended. See "SQL Server Administration" in the *IBM Tealeaf Databases Guide*.

Hardware requirements for the Passive Capture Application server machine

The server used for passive capture must meet the following hardware requirements:

- Dual processor, quad core:
Intel quad-core Xeon processor at 2.8 GHz or better for total of 8 cores minimum.
- 3 network interface cards (NIC), each card with 1-Gigabit.
- Hard drive space according to sizing
- Hard drive speed: SAS or SCSI hard drive with 15ms access time and 7200rpm drive speed.

Related tasks:

"Identifying available resources" on page 65

"Recording machine identities" on page 66

"Checking server requirements" on page 75

Data preaggregator hardware sizing information

Every customer environment is unique and it is important that you size your IBM Tealeaf installation for your environment. Use the following information to help you identify the hardware that is needed to install and operate a data preaggregator.

The hardware requirements for a data preaggregator can vary based on the size of your IBM Tealeaf environment. The minimum requirements assume that two canisters are deployed. Additional hardware is required if additional canisters are deployed. Use the following table to configure the hardware and software before installing a data preaggregator server.

Table 21. Data preaggregator hardware sizing information

Hardware component	Minimum configuration	Recommended configuration
Processor	<p>Processor: 64-bit processor with 4 cores.</p> <p>An additional processor core is required for every two additional canisters that connect to the data preaggregator. For example, if you have 4 canisters, 5 dedicated 64-bit processor cores are required.</p>	<p>Processor: 64-bit processor with 4 cores.</p> <p>An additional processor core for each additional canister that connects to the data preaggregator. For example, if you have 4 canisters, it is recommended to have 6 dedicated 64-bit processor cores.</p>
Memory	<p>Memory: 4 GB</p> <p>An additional 1 GB of memory is required for every two additional canisters that connects to the data preaggregator. For example, if you have 4 canisters, 5 GB of memory is required.</p>	<p>Memory: 4 GB</p> <p>An additional 1 GB of memory is recommended for each additional canisters that connects to the data preaggregator. For example, if you have 4 canisters, it is recommended to install 6 GB of memory.</p>
Storage capacity	500 GB	1 TB

User account requirements

IBM Tealeaf user accounts include account permissions, NT authentication, and Passive Capture Application accounts.

Account permissions

As a recommended best practice, give any account used by a Windows-based IBM Tealeaf service Local System permissions on the server.

Using non-local system accounts for running IBM Tealeaf software is not supported. Running IBM Tealeaf services with permissions other than Local System permissions might cause problems, such as failures to connect, write errors, and unexpected time-outs.

NT authentication

When NT authentication is enabled for the IBM Tealeaf system, IBM Tealeaf services must run under an NT domain account with appropriate privileges. See "Authentication" in the *IBM Tealeaf cxImpact Administration Manual*.

Additional configuration may be required. See "Database Manager SQL Server Access and Permission Requirements" in the *IBM Tealeaf Databases Guide*.

Passive Capture Application (PCA) accounts

The IBM Tealeaf PCA runs on the Linux operating system. During PCA installation, the ctccap account is created, which is used for running all PCA services. "Installation" in the *IBM Tealeaf Passive Capture Application Manual*.

Related tasks:

High-availability considerations

Before you install IBM Tealeaf, consider the implications of making the system highly available.

In order to provide uninterrupted service, Tealeaf allows highly available configurations of its capture and processing platforms.

The main objective of making Tealeaf highly available is to protect against data loss or the inability to capture data, which results in data loss.

Note: For IBM Tealeaf, protecting against the temporary unavailability of already-captured data is not a primary objective of a high-availability.

The high availability features of IBM Tealeaf cxImpact may be considered from the perspective of the main IBM Tealeaf cxImpact functional components.

The functional components will be considered in the order that captured data passes through them. Each section addresses one of the main functional units of a IBM Tealeaf cxImpact installation. The order of data flow through these components is as follows: **Passive Capture Application server > HBR server > Processing servers > Reporting server**

The relative positions of these components can be described as being "upstream" or "downstream" of each other. In this sense the IBM Tealeaf Passive Capture Application server is upstream of the HBR Server. The Reporting Server is downstream of the Processing Servers.

Typically, an upstream component is responsible for monitoring the health of the component immediately downstream from itself. This document assumes that the reader is already familiar with the architecture of a IBM Tealeaf cxImpact installation.

High availability: CX Passive Capture Application server

The Tealeaf IBM Tealeaf Passive Capture Application Server (PCA server) is responsible for the extraction of HTTP requests and responses from raw TCP/IP network data.

An HTTP request/response pair is combined to form a hit, and hits are sequenced into sessions based on defined criteria, such as the value of a session cookie.

The PCA server can also decrypt encrypted data and obscure or destroy sensitive data such as credit card numbers.

In a IBM Tealeaf cxImpact solution, there is at least one PCA server. In high-volume solutions, additional PCA servers may be deployed.

PCA device failover

The Tealeaf Passive Capture software supports failover and failback between primary and secondary Passive Capture devices. If a Tealeaf PCA server fails, the data being captured by that server is lost.

To protect against data loss resulting from such a failure, additional PCA server systems can be configured to run in a master/slave configuration. A heartbeat check from the secondary device published NIC polls the primary device published NIC at a pre-configured interval; if the polling is not successful, the secondary device begins writing data to the Processing Server.

- The master server in a pair is the normally active server.
- The slave server continually monitors the status of the master through a heartbeat check.

In the event of a failure in the master, the slave assumes responsibility for the capture of data from the master server. The slave server can be configured to fall back to the master server when the master becomes operational again.

Requirements - To enable PCA device failover, please verify the following requirements in your Tealeaf environment:

- PCA configurations on each device must be identical except for the failover settings.
Changes to PCA configurations in one device must be applied to the other.
- A second capture point (SPAN port or tap) must be active and connected to the secondary device.
See "PCA Web Console - Interface Tab" in the *IBM Tealeaf Passive Capture Application Manual*.
- Both PCAs must be receiving identical traffic feeds.

Delivery failover mode

The IBM Tealeaf Passive Capture Application supports two methods of delivery failover management, depending on the version in use:

- PCA Build 3500 or later:
In PCA 3500 or later, the default failover method is to use even distribution, which automatically redistributes traffic from a failed delivery peer evenly across the remaining delivery peers in the environment.
Set the Delivery Mode to **Even Distribution**
- PCA Build 34xx or earlier:
The PCA can be configured to recognize failures in delivery targets and then to failover to secondary targets. Set the Delivery Mode to **Failover**.

Note: This method is supported for legacy purposes and may be deprecated in a future release.

For either method, each PCA must have at least two delivery targets, a primary and secondary. If the connection to the primary target is lost, the PCA begins to send traffic to the secondary peer in Failover mode or to all remaining peers in Even Distribution mode.

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

High availability: Processing servers

A IBM Tealeaf cxImpact installation includes one or more Processing Servers, depending on the volume of data to be captured.

Among other functions, the Tealeaf Processing Server is responsible for the following:

- event processing
- alerting
- indexing of captured data
- storage of captured data

To protect against failure of one or more of these Processing Server functions, additional capacity should be configured. Additional capacity can be in the form of additional servers or additional resources within each server.

If one of the processing servers fails, the additional capacity configured allows the remaining servers to assume the load of the failed server.

For multi- Processing Server environments, an HBR Server can be deployed to monitor Processing Server health and load balancing. For more information, see HBR Server.

In a single Processing Server environment, excess capacity should be available within the Processing Server.

Each Processing Server monitors the health of its own Canister, which stores session data for Active and Completed sessions. If the Canister is temporarily unavailable or if the server is falling behind in processing, data may be spooled locally until the situation is resolved.

Depending on whether you are using Health-Based Routing, failover of the Processing Servers is handled in one of two ways.

HBR Server

Tealeaf supports Health-Based Routing within Tealeaf software for data transports to multiple Processing Servers. HBR enables load routing based on Canister health and failover in multi-Processing Server configurations.

The Health-Based Routing (HBR) server distributes incoming captured data across multiple Processing Servers. The HBR server monitors the health of the Processing Servers in the system, so it is recommended that an HBR server be configured in any Tealeaf installation with more than one Processing Server that must be highly available.

Smaller Tealeaf installations might not require an HBR Server. Smaller installations typically have low data volumes requiring a single Processing Server only.

The HBR Server manages the distribution of incoming session data among the available Processing Servers. HBR monitors the health of the Processing Servers and, if one of them becomes unavailable, stops sending data to that server and redistributes the incoming data across the remaining servers. When the unavailable server becomes operational again, the HBR server resumes sending data to it.

Functionally, the HBR Server is a Tealeaf Processing Server without a local canister. A Windows machine is dedicated to running the Tealeaf Transport Service. It polls each Processing Server for availability and spooling status. The HBR Server runs a Tealeaf pipeline that includes the HBR session agent. This agent performs the following functions:

- Monitor the availability of downstream processing servers.
If a machine is unavailable or is spooling data, then HBR reallocates its traffic to other available Processing Servers.
- Distribute the incoming captured data amongst the available processing servers.
- Spool incoming data in the event that no processing servers are available or the available servers cannot deal with the volume of incoming data between them.

Since the HBR Server performs a central data management and distribution role, it is a potential single point of failure. To protect against this, a hot standby HBR Server may be configured. The standby server should be identical in capabilities and configuration to the primary HBR Server.

The PCA server can be configured with a primary and secondary delivery peer, where the primary peer is the active health-based router, and the secondary peer is the failover HBR machine.

See "Health-Based Routing (HBR) Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Non-HBR configurations with multiple Processing Servers

You can achieve higher availability by having two PCA servers feeding two Processing Servers. For configurations with multiple Processing Servers and no Health-Based Routing that use the PCA failover feature, the supported solution is to use an "active/active" model for failover management. In this model, the PCA is configured to send half its data to each Processing Server. If the connection to a Processing Server is lost, the PCA sends all its data to the remaining server.

An alternative approach, using one active Processing Server with a standby Processing Server for failover, is not supported.

Delivery to multiple Processing Servers can be configured through the PCA Web Console. For information about using the Web console, see "PCA Web Console - Delivery Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

High availability: Portal web application

In the event of a failure of the Portal Web application, a failover switching script can be executed to resume functioning of the Web application on a separate platform.

Typically, a secondary machine with identical configuration is deployed as a failover, with the suite of Tealeaf Data Services stopped. When this machine becomes active, a script starts the services, which then assume collection of session data from the Processing Servers.

To start the secondary machine, login to the machine and select the following shortcut from the Windows Start menu: **Start > > All Programs > Tealeaf Technology > Start Tealeaf Services**.

High availability: Reporting server

The Tealeaf Reporting Server hosts the Report Server and Portal Web application, which are supporting components of the Tealeaf Web Portal application.

These components store data in a SQL Server database that can be installed locally on the Reporting Server or installed remotely on a separate server.

User access to Tealeaf functionality is via a combination of the Tealeaf Web Portal and the IBM Tealeaf RealiTea Viewer. Both access methods depend on the Reporting Server being available.

At any time, a Tealeaf installation should only have a single instance of the Tealeaf Data Service, a component of the Reporting Server, since this data is collected from Processing Servers, aggregated, and then removed from subsequent collection. As a result, high availability strategies for the Reporting Server require a cold standby system or rapid rebuild/replacement of an unavailable Reporting Server.

Note: When deploying a failover Report Server, you must also deploy a mirrored version of all Tealeaf databases. Tealeaf cannot natively update two sets of databases in real-time, so you must schedule with your database administrator daily or otherwise periodic updates of your mirrored set of databases.

Cold standby

In the cold standby scenario, a second identical Reporting Server is inactive until needed. If the active Reporting Server becomes unavailable, the standby server can be quickly activated to assume the responsibilities of the failed server.

This method has the benefit of minimizing the period of time during which Tealeaf reporting data is unavailable to Tealeaf users. However, it incurs the overhead of maintaining a second Reporting Server that is not being used most of the time.

Note: Care should be taken to ensure that changes to the configuration of the active server are also made to the standby server.

Rebuild/Replace

In the rebuild/replace scenario, a new server is provisioned to replace the failed Reporting Server. This scenario requires installation and configuration of the hardware, operating system, and Tealeaf components.

This method has the benefit of incurring the overhead of a second server only as needed. However, Tealeaf reporting data is unavailable to Tealeaf users for a longer period than in the cold standby method. If a new server can be made available within an acceptable period of time, this method may be appropriate.

Note: Statistical data is retained on the Processing Server for a period of 72 hours before it is discarded. As long as a replacement server is provisioned within this period, no data is lost.

Example: High availability configuration

Example Resilient Tealeaf Installation below depicts a Tealeaf installation with the following characteristics:

- Two primary capture servers, each with a failover slave server

- A primary and standby HBR server
- Three processing servers
- An active and a cold standby reporting server




Figure 5. Example Tealeaf failover configuration

Web application development considerations

There are things you should consider when building or reviewing your web applications for use with IBM Tealeaf.

Generally, Tealeaf acquires most of its data from the real-time HTTP data stream between your web servers and the client computers. Upon capture, Tealeaf parses this text data in order to:

- organize hits into the related session
- detect user-defined event patterns
- generate indices for searching
- create reports
- replay the sessions

All of these functions require clear, easily parsed request and response data.

Tealeaf does not require JavaScript tags to be implemented on your web application.

Browser support

The browser-based replay feature supports the following browsers:

- Mozilla family of browsers
- Gecko

In Release 8.0 and later, the Replay Server utilizes the Gecko browser control to render session content on the server for delivery to the browsers of Tealeaf

users. Replay of such sessions may still be supported through the IBM TealeafRealiTea Viewer, a desktop client that must be separately installed on each Tealeaf user's desktop. See "RealiTea Viewer (RTV) User Manual" in the *IBM Tealeaf RealiTea Viewer User Manual*

Data privacy

To facilitate the blocking, encryption, or hashing of sensitive data points, you should release code that minimizes the instances of sensitive data passed between web server and client.

For each potentially sensitive data point:

- Identify whether the data is required for reporting or alerting purposes. If the data is required for reporting or alerting purposes, it must be made available to some Tealeaf users in some way. Review the data with your security team.
- Verify with your security team the conditions under which the data must be blocked.
- Add a consistent tag to the unique identifier for the data point, such as `-private` so that all sensitive data can be easily identified using configured privacy rules.

As a security practice, do not display any sensitive data such as credit card details or site passwords in the HTTP response. Instead, you can echo this type of data in the response after the visitor has encountered an error, such as a credit card authorization error, to assist the visitor in completing the transaction.

Note: Some industry standards expressly limit the transfer of this data in the response. For example, the Payment Card Industry (PCI) Data Security Standard permits the presentation of the first six or the last four digits of a credit card number only.

Tealeaf can be configured to block this data. However, this configuration is expensive in terms of processing and may be complicated to successfully complete.

As a best practice, develop your Web application such that Tealeaf's blocking of sensitive data is only necessary within the HTTP request.

If sensitive data is appropriately bracketed in the HTML responses generated by your web application, Tealeaf includes some pre-built privacy rules that automatically block or encrypt the data.

Note: To utilize these pre-built rules as they are configured by default, you must bracket sensitive data in your HTML responses with specific comment tags.

Pop-up windows

Since client-side pop-ups and client-side validation do not involve any communication with the web server, they cannot be tracked using Tealeaf's IBM Tealeaf Passive Capture Application.

If you need to track these types of client-side interactions for your web application, the Tealeaf IBM Tealeaf UI Capture for AJAX library can be deployed to gather this information. This library requires additional deployment, configuration, and maintenance and may involve Tealeaf Professional Services. See "UI Capture for

AJAX Guide" in the *IBM Tealeaf UI Capture for AJAX Guide*.

Request and response data

When developing your web applications, employ the developing practices listed in Table 22.

Table 22. Considerations for web application development and request and response data

Do this	For this reason
Run W3C validation services on your web site	<p>Tealeaf products require that web sites conform to W3C standards, including having properly formed HTML.</p> <p>Making sure your web site complies to W3C standards can save time during the Tealeaf integration process.</p> <p>The processing that Tealeaf performs on web pages is significantly more complex than a simple rendering; meaning that just because web pages render correctly in a web browser, does not guarantee that they will replay with Tealeaf.</p>
Use unique and common identifiers, as appropriate	<p>General principles of uniqueness:</p> <ul style="list-style-type: none"> • Objects that must be specifically identified by Tealeaf, such as screen elements or JavaScript items, require unique identifiers. • Groups of objects that must have common actions performed on them inside of Tealeaf should have a shared tag in them for easy identification. For example, sensitive data that must have privacy applied to them could have the tag -private as part of the identifier. <p>For more information, see Table 23 on page 61.</p>
Avoid the use of encrypted or encoded field values	<p>Use of encrypted or encoded field data within the web site application can complicate and slow down the extraction process.</p> <ul style="list-style-type: none"> • Where data can be decoded, Tealeaf can extract information with some additional configuration work. • In cases where the data is truly encrypted, Tealeaf may not be able to extract that information. <p>Encryption refers to application-level encryption of specific data items such as form field or cookie values.</p>
Avoid the use of client-generated data	<p>Applications should avoid using date and timestamps generated on the client, as these values are different from the values recorded by the server at the time of capture by Tealeaf. For best results, use server-side data for generating values such as timestamps.</p>
Embed report segmentation data in requests	<p>Provide hidden fields or name and value pairs in the request for data that may be used for report segmentation. For example, when a customer logs in to the web application, submit in the request hidden fields such as Account, Company, and <i>Geographic Region</i>.</p>
Insert name/value pairs in response data	<p>In the response, inserting name/value pairs as HTML comments near the top of the response makes them efficient to find and easy to read for the analyst:</p>

Table 22. Considerations for web application development and request and response data (continued)

Do this	For this reason
Provide server errors	If possible, pass server errors in the response. Please provide some mechanism to identify application servers that are generating errors or exceptions. In Tealeaf, events can be created to monitor errors passed in the response, which enables the reporting of them.
Identify language-dependent or frequently-changing text using HTML markup	<p>Tealeaf can monitor user-visible error messages in your web site, which provides useful insight into customer behavior and web application friction points. Depending on the web application, its error messages may frequently change, as well.</p> <p>Tealeaf events can be configured to track errors by detecting the text of the error messages in the response data. However, this approach can result in a high overhead to maintain events when error messages are frequently changed. This overhead is compounded on multi-lingual sites, where the number of events affected by a change is equal to the number of changed error messages multiplied by the number of languages supported by the site.</p> <p>To minimize event maintenance, you can identify frequently-changing text such as error messages with some HTML markup, such as an ID or name attribute.</p>

Table 23. Objects in your web application that should be unique

These objects should be unique	For this reason
Name/value pairs	In general, the key to easy creation and management of accurate events is to create explicit name/value pairs in your web data. Generate unique and parsable field names.

Table 23. Objects in your web application that should be unique (continued)

These objects should be unique	For this reason
Page URL paths	<p>Intentionally or accidentally, many sites are designed with non-unique URL paths (URI stems) to identify page locations. This lack of uniqueness creates a challenge when analyzing a specific data point on a specific web page of the site.</p> <p>An example URL path might be the following: /deposits/moneymarket/accountsetupStep2.jsp</p> <p>Where possible, please observe the following guidelines:</p> <ul style="list-style-type: none"> • Unless a coding, security, or other standard requires otherwise, each viewable page should have a unique URL path, which includes a distinct application directory structure defining the application (section) of the site and specific page in that application. • If unique URL paths are not permitted or feasible, then a consistent request or response header, request field, response field, or response text on all pages can perform the same function. • The application (section) of the site should define each section of the site on which a report could be delineated. • Regardless of technique, for each viewable web page and the application of the site, a consistent unique parameter displayed in either the request or response should be created. <p>Ensuring this uniqueness should be a joint effort between the developers and content owners of your site.</p>
Form field names or IDs	<p>Tealeaf can extract and display a wealth of data from captured sessions. A large proportion of this information is obtained from data provided by the user in the form of information entered in form fields.</p> <p>The use of unique form field names across the site avoids confusion in event definition and improves the accuracy of the reported data.</p>
JavaScript identifiers	<p>To simply replay of sessions, please use unique identifiers for JavaScript elements.</p>
Step identifiers in business processes	<p>Tealeaf provides the Business Process Scorecard report, which can provide insight into the effectiveness of site business processes and factors affecting the conversion rate of those processes. The foundations of a process scorecard are step events, which track the user's progress through the steps of the process. To track this progress, configured Tealeaf events must be able to unequivocally identify each process step.</p> <p>The process of creating these events can be dramatically simplified if the web application provides unique step identifiers that do not change between web application revisions. Some examples of step identifiers are:</p> <ul style="list-style-type: none"> • a unique URL • a Web Analytics tag • an HTML comment • a breadcrumb trail indicator such as text or an image to indicate the step in the process that the user has reached

Rich internet applications

When developing a site that includes one or more Rich Internet Applications utilizing AJAX, Flex, or significant client-side script, you must deploy the IBM Tealeaf UI Capture for AJAX to capture client-side user interface events. Development of these types of applications has additional recommended practices. See "UI Capture for AJAX Best Practices" in the IBM Tealeaf UI Capture for AJAX Guide.

```
<!-- A_UNIQUE_VAR_NAME=SOME_VALUE -->
```

Session tracking and cookies

Tealeaf evaluates all of the request/response pairs (hits) between the servers and the clients to recreate the customer experience. From this interleaved data, Tealeaf requires all requests and responses of the session and a common key to join them to the session.

Typically, this identifier is a session-level cookie defined by the web application, such as ASP.NET_SessionId or JSESSIONID.

In some cases, such as landing pages, this session-tracking cookie is not set on each hit. This case can result in orphaned, standalone hits that are not assigned to a specific session. To limit these orphaned hits, verify that session tracking information is available on every web page served by the web application.

If the web application cannot provide this identifier, the Tealeaf Cookie Injector can be deployed on your web server to provide an appropriate session-tracking cookie. See "Installing and Configuring the Tealeaf Cookie Injector" in the *IBM Tealeaf Cookie Injector Manual*.

Related concepts:

"Data privacy and Web application development" on page 14

"Pre-Built privacy rules IBM Tealeaf" on page 16

Chapter 3. Preparing to install the software

Preparing to install IBM Tealeaf CX involves assessing your current IT resources, determining and cataloging the IBM Tealeaf CX component-to-server assignments, installing and configuring prerequisite software and preparing the servers and your network environment to accept IBM Tealeaf.

Identifying available resources

Identify your assets to make the best use of those resources already available.

You must be familiar with your current hardware and software.

Prepare a list of the available assets.

You are assessing your current enterprise information system to determine if you require any additional hardware or software to meet your business needs.

1. Determine the number of physical computer systems you will use and itemize each piece of physical hardware.

Note the following:

- Amount of installed memory
- Number and type of installed microprocessors
- External media
- Whether a particular unit can be upgraded

2. Itemize the currently installed software and database applications. Note:

- Function
- Breadth of use across the company
- Security requirements

3. List your current IT personnel and note whether you have available the means for installation and maintenance of IBM Tealeaf, as well as the required expertise in database management.

Make sure all involved have user IDs with the appropriate authorizations to successfully install all products and files.

4. Make sure you have access to appropriate file transfer software and remote access utilities.

To move Tealeaf installation files from your laptop to the Linux box that will be your Tealeaf Passive Capture Server and to the Windows box that will be your all-in-one CX Sever, you might need additional 3rd party software.

To access and configure LINUX server remotely, you need a command shell SSH client.

If you are not sure whether your current assets are adequate for the IBM Tealeaf solution to be deployed, IBM Tealeaf Services personnel can assess the traffic on your site and tell you whether you need to purchase additional hardware.

Related concepts:

“Software requirements” on page 49

“Hardware requirements” on page 51

“Operating system requirements” on page 47

Recording machine identities

Record the identities of the machines to be used in the IBM Tealeaf installation environment.

Bios-level machine names are acceptable. However, DOMAIN-LEVEL MACHINE NAMES ARE NOT ACCEPTABLE.

Recording the identities can help you manage the installation process and administer the machines post-installation if necessary.

Use the following table as a worksheet for recording machine identities.

Table 24. Server names and IP addresses

Server name	IP address	Machine name (optional)	Notes
Passive Capture server			None.
CX server			None.
SQL server			The SQL server can be separate box.
HBR server			The HBR server is rarely needed. Consult with Tealeaf services personnel to determine if your environment configuration requires an HBR server.
Report server			In most cases, the Report server and CX server are the same servers.
SMTP server			This is the server that IBM Tealeaf uses for alerts.

Related concepts:

“Hardware requirements” on page 51

Installing and configuring prerequisite software

Before you install and configure IBM Tealeaf CX, you must install and configure the prerequisite software.

The task flow for installing and configuring prerequisite software is:

1. Install Microsoft IIS.

The IBM Tealeaf Portal will not work without IIS installed. It might be that you already have IIS installed and you just need to verify that it is installed.

Make sure that the **Startup Type** for IIS is set to Automatic.

For information about installing IIS, see the appropriate Microsoft web sites and documentation.

2. Determine what version of .NET installed on your Windows Server
IBM Tealeaf CX requires .NET 4.5
3. If you do not have the required version of .NET installed on your Windows Server, install it from the Microsoft Developer Network.
4. Enable the required .NET framework features for the Application Development server role.
5. Register IIS.
6. Install SQL Server.

For information about installing SQL Server, see the appropriate Microsoft web sites and documentation.

7. Configure SQL Server to work with IBM Tealeaf.

See *Software requirements* in this guide for information about installing and using IIS and SQL Server with IBM Tealeaf.

Related concepts:

“Software requirements” on page 49

“Setting up SQL Server to work with Tealeaf” on page 68

Installing .NET framework

In order to run IBM Tealeaf Portal, you must install Microsoft .NET framework and ASP.NET

IBM Tealeaf requires version 4.5 of .NET framework.

Before installing .NET framework:

- Make sure you have access to Windows Server documentation and the Microsoft Developer Network web site.
Microsoft Developer Network web site provides articles that help users install, configure, administer and use Microsoft products.
- Search the Microsoft Developer Network for articles and help on how to determine what version of .NET framework is currently installed.
- Make sure you have installed IIS.

There can be variations in the installation procedure for .NET and ASP.NET based the version of Windows Server you are running. Also, Microsoft provides more than one way to install .NET and register it for use. For example, you can install and register NET framework via a command-line or by using the Server Manager interface to add the required NET and ASP.NET features. Consult your Microsoft documentation and technical resources to determine which installation procedure suits your Windows Server configuration.

This procedure provides general guidelines only for installing .NET Framework and enabling the role-based functions required by IBM Tealeaf components and your SQL Server installation.

1. Determine what version of the .NET framework is installed on Windows Server.

Search the Microsoft Developer Network web site for articles on how to determine your .NET framework version.

If the installed version of .NET framework and ASP.NET does not meet the version level required by IBM Tealeaf CX, then you must install the required version.

2. Install the .NET framework

There are multiple methods for installing the .NET framework:

- You can install .NET framework using the Server Manager Interface
- You can install .NET framework using PowerShell

Search the Microsoft Developer Network for articles to determine which installation method is best for your site.

3. Add the following features to the Application Development role:

- ASP.NET 4.5
- .Net Extensibility 4.5
- ISAPI Extensions
- ISAPI Filters

You can use the Windows Server Add Roles and Features wizard to add features to the Application Development role. For information on how to add roles and features, see the Windows Server documentation or articles on MSDN.

You now need to register IIS.

Related concepts:

“Operating system requirements” on page 47

Registering IIS

After you install the latest .Net Framework you must register IIS to ensure that the framework is ready to receive the Tealeaf installation.

To register IIS:

1. From the **Windows** icon, select **Run...**
2. Type `cmd` and press **Enter** to open the Windows Command Prompt.
3. Navigate to the directory where .Net Framework is located.

Typically, .Net Framework is located in the following directory:

`systemroot\Microsoft.NET\Framework\versionNumber.`

For example: `\windows\Microsoft.Net\Framework64\v2.0.50727.`

4. From the .Net Framework directory, enter the command `aspnet_regiis -i.`
5. Press Enter.

You can now ready to configure SQL Server to work with IBM Tealeaf.

Setting up SQL Server to work with Tealeaf

After installing SQL Server you need to configure it to work with Tealeaf.

Setting up SQL Server 2008 to work with Tealeaf

Use the information in Table 25 as a guideline for configuring SQL Server 2008 to work with IBM Tealeaf. Consult SQL Server 2008 documentation and the Microsoft Developer Network for specific instructions with regard to configuring SQL server.

The information in this topic assumes you have installed SQL Server. For information about installing SQL Server, see SQL Server documentation.

Table 25. Configuring SQL to work with Tealeaf - A reference table

Things to do	Notes						
<p>Check the amount of RAM allocated to the SQL Server.</p>	<p>For a single-server topology, no more than 2 GB of the system RAM should be allocated to Microsoft SQL.</p> <p>To check the amount of RAM</p> <ol style="list-style-type: none"> 1. Right-click the Computer icon on the desktop. 2. Select Properties <p>View basic information about your computer</p> <hr/> <p>Windows edition</p> <p>Windows Server Enterprise Copyright © 2007 Microsoft Corporation. All rights reserved. Service Pack 2</p> <hr/> <p>System</p> <table border="0"> <tr> <td>Processor</td> <td>Intel(R) Xeon(TM) CPU 3.0GHZ (2 processors)</td> </tr> <tr> <td>Memory (RAM)</td> <td>4.00 GB</td> </tr> <tr> <td>System type:</td> <td>32 bit Operating System</td> </tr> </table>	Processor	Intel(R) Xeon(TM) CPU 3.0GHZ (2 processors)	Memory (RAM)	4.00 GB	System type:	32 bit Operating System
Processor	Intel(R) Xeon(TM) CPU 3.0GHZ (2 processors)						
Memory (RAM)	4.00 GB						
System type:	32 bit Operating System						
<p>Connect to the Server</p>	<p>Login to SQL is managed via the Microsoft SQL Server Management Studio.</p> <p>A login screen is displayed that allows you so specify the Tealeaf Server name and click Connect.</p> <p>Note: If the Database Administrator completed a non standard configuration, you will need details from the administrator.</p> <p>If after clicking Connect you are not connected to the database, contact the database administrator. The database administrator can provide you with the information needed to authenticate with the database.</p>						

Table 25. Configuring SQL to work with Tealeaf - A reference table (continued)


Things to do	Notes
<p>Configure SQL Server Consumption</p>	<p>The following instructions are a guideline. Consult your Microsoft SQL Server documentation or articles from the MSDN for specific instructions.</p> <ol style="list-style-type: none"> From the SQL Server Management Studio, right-click on the top most entry and select Properties. Confirmation that the instance of SQL Server isn't set up to consume more than 50% of the server's RAM is required. Note: If you do not configure the server consumption rate, SQL Server consumes all of the allocated server memory.  <ol style="list-style-type: none"> Modify the Maximum Server Memory to be half of the System RAM available to the system per the available System Memory. Click OK Note: This does not automatically change the memory allocations. These must be applied by executing the script. To have the script appear, select the script icon in the upper left corner. This script is shown once clicking on OK. Click Execute. The messages generated by running the SQL Script indicate that the changes have been applied to the database are displayed.

Table 25. Configuring SQL to work with Tealeaf - A reference table (continued)

Things to do	Notes
Configure SQL Server for Mixed Authentication Mode	<p>The following instructions are a guideline. Consult your Microsoft SQL Server documentation or articles from the MSDN for specific instructions.</p> <ol style="list-style-type: none"> 1. From the SQL Server Management Studio, right-click on the top most entry and select Properties. 2. In the navigation pane, click Security. The security configuration information for the server is displayed in the viewing area. 3. Make sure the radio button for SQL Server And Windows Authentication Mode is selected. Note: If any other parameter is set, contact the database administrator to confirm that modifying the authentication mode is allowed. 4. Close the SQL Server Management Studio.
Configure SQL Server 2008 for TCP/IP Protocol Support	<p>The following instructions are a guideline. Consult your Microsoft SQL Server documentation or articles from the MSDN for specific instructions.</p> <ul style="list-style-type: none"> • Open the SQL Server Configuration Manager from the start menu. • Expand the SQL Server Network Configuration tab. • Make sure that TCP/IP is enabled. • Highlight Protocols • For MSSQLSERVER and ensure that TCP/IP is Enabled. If TCP/IP is not enabled, right-click on TCP/IP and change the setting to Enabled. •
Configuring SQL Server 2008 for DATABASE Port 1433	<p>Ensure that the TCP DYNAMIC PORTS is not enabled.</p> <p>If the database has been configured using dynamic ports contact will again be needed with the local database administrator to have them disabled.</p>

Related tasks:

“Installing and configuring prerequisite software” on page 66

Preparing the installation environment

You must prepare your environment to accept the IBM Tealeaf installation.

Preparing the environment involves:

- Configuring IBM Tealeaf user accounts
- Preparing the domain
- Preparing the network environment
- Configuring the network interface cards (NICs)

Configuring IBM Tealeaf user accounts

You need to assign proper authorization and permissions to individuals at your site who are to access IBM Tealeaf.

For more information on IBM Tealeaf cxImpact users and groups, see "CX User Administration" in the *IBM Tealeaf cxImpact Administration Manual*.

For more information on IBM Tealeaf cxView users and groups, see "cxView User Administration" in the *IBM Tealeaf cxImpact Administration Manual*.

For more information on IBM Tealeaf cxReveal users and groups, see "cxReveal User Administration" in the *IBM Tealeaf cxReveal Administration Manual*.

Related concepts:

"User account requirements" on page 52

Preparing the domain

You must perform several tasks to prepare the domain.

1. Set up a local user account

Set up a local user account called Tealeaf with local administrator rights on the IBM Tealeaf CX Server box.

This local Tealeaf user runs IBM Tealeaf CX processes and services.

Adding the Tealeaf user on the IBM Tealeaf CX box eliminates issues with password changes that may arise from using a domain administrator's or the main local administrator's account and password.

2. Configure virus scanning software

Exclude Tealeaf directories from Real Time File Protection. This eliminates possibility of the antivirus software interfering with the operation of Tealeaf software.

At a minimum, configure the virus scan to not include the directories in which Tealeaf creates and opens files in locked exclusive mode. These directories include the:

- Canister
- Indexes
- Spool
- Logs

directories.

The locations for these directories are determined when you install the software. To find out the actual location of these directories after an installation:

- Tealeaf install directory:

<Tealeaf_install_directory>

The value is recorded in the Windows registry under the following entry:

HKEY_LOCAL_MACHINE\SOFTWARE\TeaLeaf Technology\TeaLeafPath

- Canister: View the file

<Tealeaf_Install_Directory>\Ctree\Server\ctsrvr.cfg

The Canister directory is specified in Local Directory.

- Logs: Typically, Tealeaf components write logs to the Logs directory.
- Indexes:

To find the location of the indexes, use the "Tealeaf Management System" in the IBM Tealeaf cxImpact Administration Manual. See "Configuring CX Indexing" in the *IBM Tealeaf CX Configuration Manual*.

- Spool:

Log File directories are specified in TealeafCaptureSocket.cfg file in the Tealeaf install directory. The location is specified in the <SpoolDir> property value

- Decouple:

Log File directories are specified in TealeafCaptureSocket.cfg file in the Tealeaf install directory. The location is specified in the **<LogDir>** property value.

- If this value is commented out or unspecified, the directory defaults to the Logs directory inside the Tealeaf install directory.

- Backup:

The Backup directories are specified in the **<BackupDirectory>** property in the following file:

```
<Tealeaf_install_directory>\tools\TLBackupCFG.xml
```

3. Disable the following services on the IBM Tealeaf CX server to improve system performance:

- NNTP
- FTP
- SMTP Relay
- Disk Indexing

Preparing the network environment

The network environment involves the network infrastructure including ports, email administration, and replay server.

Before you deploy IBM Tealeaf, verify that your network infrastructure is configured to send data to the IBM Tealeaf Passive Capture Application in a method that is suitable for capture. See "Tealeaf PCA Network Capture Traffic Requirements" in the *IBM Tealeaf Passive Capture Application Manual*.

Preparing the network environment involves opening ports, configuring email requirements and configuring the replay service.

To prepare the network environment:

1. Configure firewall openings for bidirectional communication.

See Table 26 for a description of the ports to open.

Table 26. Opening Ports

Port	2 way?	Origination	Destination	Change port?	Description
22	Y	IBM Tealeaf CX server or administrator's desktop	IBM Tealeaf Passive Capture Application server (PCA server)	N	Remote Secure Shell
25	Y	IBM Tealeaf CX Server	SMTP Mail system	N	Sends mail via SMTP from components of IBM Tealeaf CX to the mail server
80 or 443	Y	Portal thin client (IE) and IBM Tealeaf RealiTea Viewer Pro (Analyst desktop)	Portal Server (IBM Tealeaf CX server)	N	Thin Client Note: To enable effective replay, the Replay Server might require access to static content stored on the origin server. Typically, this access occurs over ports 80 or 443. See "Configuring the Replay Server" in the <i>IBM Tealeaf CX Configuration Manual</i> .
1433	Y	Report Server	Remote SQL Server	Y	Needed only if Reporting DB is MSSQL and Report Server is on a platform other than MSSQL server
1966	Y	Capture Filters (Native and PCA)	IBM Tealeaf CX server	Y	Captured data stream

Table 26. Opening Ports (continued)

Port	2 way?	Origination	Destination	Change port?	Description
5597	Y	Report Server	Canister	Y	Communications between Portal and ctree session database on Canister.
8080 or 8443	Y	IBM Tealeaf CX server or administrator's desktop	PCA	Y	PCA Status and configuration program
19000	Y	Search Server	IBM Tealeaf CX Server	N	Searches and retrieves sessions via TCP/IP
20000	Y	Portal Server, Tealeaf Management System (on all servers except PCA server)	Tealeaf Management System (on all servers except PCA server)	N	Communication for retrieving and sending configurations of Tealeaf servers, components, and databases
23000	Y	Portal	Tealeaf databases	N	Tealeaf Data Service
38000	For advanced configuration of your replay server, contact your account representative or Professional Services.				

2. Configure email requirements.

Most organizations configure email systems to accept email at a particular machine, such as smtp.company.com or mail1.company.com.

To support IBM Tealeaf CX email requirements:

- a. Configure the SMTP receiving computer to accept mail for the IBM Tealeaf CX server.
- b. Verify that the IBM Tealeaf CX server can resolve the SMTP server name and that the SMTP server can be pinged from the IBM Tealeaf CX server. The SMTP port between the IBM Tealeaf CX server and the SMTP server must be opened through any intervening firewalls.

Note: It is recommended that you identify a Tealeaf administrator with a valid, reachable email address before you begin the installation process.

3. Configure the replay service

The replay server manages replay of sessions.

You can configure the replay server to access replay of sessions or disable interactive services detection.

- a. Enable Replay server access to the web application

For more information, see "Configuring the Replay Server" in the *IBM Tealeaf CX Configuration Manual*.

- b. Disable Interactive Services Detection

When the WebKit/Chrome renderer is enabled for the Replay Server, it conflicts with the Interactive Services Detection service on the hosting machine. This service must be disabled.

The Replay Server is a Tealeaf component used to render sessions for replay through Browser Based Replay. It is typically hosted on the Portal Server.

The WebKit/Chrome renderer is used to render for replay sessions initiated from mobile devices.

Note: The Interactive Services Detection service is automatically disabled by the Tealeaf Installer or Upgrader during normal operations. Afterward, the service can be disabled through the Windows Services control panel on the server hosting the Replay Server.

Note: If your environment utilizes multiple Replay Servers, this service must be disabled on each Replay Server in your environment.

Alternatively, you may disable this mobile-based renderer for individual groups through the Portal Management page.

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

See "cxReveal User Administration" in the *IBM Tealeaf cxReveal Administration Manual* section.

Related tasks:

"Enabling your IBM Tealeaf solution" on page 91

Configuring the network interface cards (NICs)

Before you install IBM Tealeaf Passive Capture Application server or IBM Tealeaf CX, configure the settings on your network adapter.

Configure the following network adapter settings in your operating system for use in your network environment:

- Host Name
- DNS Servers
- IP Address (IP address for management on the PCA and one IP address for the IBM Tealeaf All-in-one server)
- Subnet Mask
- Default Gateway

Note: The IBM Tealeaf Passive Capture Application server can have multiple NICs. For information about configuring NICs for the Passive Application Server, see the *IBM Tealeaf Passive Capture Application Manual*.

Checking server requirements

Before you install IBM Tealeaf CX and enable CX features and functions, make sure that the server meets the RAM, disk space, CPU core count, and NIC card requirements.

1. For information about disk space, select **Start > Computer** from Windows.
 - a. Click the **General** tab to display information about disk capacity the free space available.
2. For information about CPU and RAM, select **Start > Control Panel** from Windows.
 - a. Select the option for **System and Security** and then select the option for System.
3. Use the following table to record information.

Table 27. CX server space and memory information.

Element	Expected	Actual	Notes
RAM			

Table 27. CX server space and memory information. (continued)

Element	Expected	Actual	Notes
CPU / Core Count			
Disk type / space			
NIC card			

Related concepts:

“Hardware requirements” on page 51

File permission considerations

The IBM Tealeaf installer program modifies permissions as required for the Portal Service.

Install IBM Tealeaf using the default Windows Server file permissions.

Other things to consider with regard to file permissions and installing IBM Tealeaf:

- You should not modify the Windows Server file permissions prior to installing IBM Tealeaf.
- You can look into modifying or restricting file permissions after you verify the IBM Tealeaf installation.
- If you modify or restrict file permissions after installing IBM Tealeaf and you determine that your changes have caused problems with IBM Tealeaf, undo the file permission changes.

Note: If your site requires you to change the default Windows Server file permissions when installing software, contact IBM Tealeaf Professional Services so that they can assess whether the changes will impact the IBM Tealeaf installation.

User account considerations

IBM Tealeaf user accounts include account permissions, NT authentication, and Passive Capture Application accounts.

Account permissions

As a recommended best practice, give any account used by a Windows-based IBM Tealeaf service Local System permissions on the server.

Using non-local system accounts for running IBM Tealeaf software is not supported. Running IBM Tealeaf services with permissions other than Local System permissions might cause problems, such as failures to connect, write errors, and unexpected time-outs.

NT authentication

When NT authentication is enabled for the IBM Tealeaf system, IBM Tealeaf services must run under an NT domain account with appropriate privileges. See "Authentication" in the *IBM Tealeaf cxImpact Administration Manual*.

Additional configuration may be required. See "Database Manager SQL Server Access and Permission Requirements" in the *IBM Tealeaf Databases Guide*.

Passive Capture Application (PCA) accounts

The IBM Tealeaf PCA runs on the Linux operating system. During PCA installation, the ctccap account is created, which is used for running all PCA services. "Installation" in the *IBM Tealeaf Passive Capture Application Manual*.

Chapter 4. Installing IBM Tealeaf

Installing IBM Tealeaf involves obtaining the installation files for IBM Tealeaf CX and for the IBM Tealeaf products and then running the product installers.

Before installing the software for IBM Tealeaf, assess your current environment and your business requirements to ensure that the system you implement meets your needs.

Use of UNC paths in the IBM Tealeaf Installer is supported only if they reference locations on the local machine.

The information in this chapter explains how to obtain the installation file for IBM Tealeaf CX and how to run the installer. It also contains information about installing and enabling IBM Tealeaf products, pointing you to product-specific publications when necessary.

Related concepts:

“IBM Tealeaf overview” on page 3

Installing on a single server

The installation procedure for a single server topology is different than the installation procedure for a multiserver topology.

Overview of the installation procedure for a single server topology

If your IBM Tealeaf deployment environment consists of a single server, the installation steps involved are as follows:

1. Install IBM Tealeaf CX.
2. Enable additional IBM Tealeaf products on the same server.
3. On local Windows desktops, you can install IBM Tealeaf RealiTea Viewer, the desktop application for searching and replaying IBM Tealeaf sessions.

Note: It is recommended that you use Browser Based Replay, the web browser-based method of replaying sessions. This client requires no additional installation. For information about Browser-based replay, See "CX Browser Based Replay" in the *IBM Tealeaf cxImpact User Manual*.

4. Perform post-installation tasks and product configuration tasks for you single-server solution.
5. Start the IBM Tealeaf CX solution and verify that it is operational.

Obtaining installation files

You can obtain IBM Tealeaf products from Passport Advantage in either electronic (downloadable) or physical (DVD) format.

For more information on Passport Advantage, visit the web site at <http://www-01.ibm.com/software/passportadvantage/>.

Follow your site's standards and procedures for saving software product files to your environment.

Installing IBM Tealeaf CX

Install IBM Tealeaf CX before installing any other IBM Tealeaf product.

Make sure the operating system, hardware, and software requirements have been met and that you have completed all tasks to prepare for the installation.

Verify that you have access to a Windows user account with administrator permissions. The IBM Tealeaf CX installer includes panels for setting up databases. If you plan on setting up databases during the installation, make sure your user account has the authorization to set up databases.

Make sure you have obtained the installation files for IBM Tealeaf CX and that you have access to the directory in which the installation files reside.

The IBM Tealeaf CX installer gathers the configuration information for your IBM Tealeaf CX installation and creates the installation when you click **Install**.

Installing IBM Tealeaf CX provides the platform that is required for IBM Tealeaf solutions.

The installation wizard includes panels for installing databases, however, you can choose to install databases after running the installer.

The default installation for the Replay Server component is 32-bit. You can run **TLReplaySrv.exe -install** on the x64 sub-directory to switch from the default 32-bit Replay server to a 64-bit Replay server anytime after the installation. See "Switching to a 64-bit Replay server" in the *IBM Tealeaf CX Configuration Manual* for instructions.

Note: Switching to a 64-bit Replay server supports loading a larger number of sessions on a single replay server.

Note: The following installation instructions assume that you are installing IBM Tealeaf CX for the first time. This process installs the base files and creates the necessary databases. You need to provide the additional information to identify the Data Collector and Report Server if you are not performing a full installation and you are installing an instance of the Tealeaf Portal that does not use the default Data Collector and Report Server settings. The following information is required to configure the Data Collector and Report Server settings.

- SQL server name and port
- Database name for the Data Collector and Report Server
- Database path for the Data Collector and Report Server

After the Tealeaf Portal is installed, you can log in to the portal and click **Tealeaf > Portal Management > Manage Servers** to edit the server configuration for your IBM Tealeaf environment.

To install IBM Tealeaf CX.

1. In the software distribution, navigate to the following folder: cx\TealeafCX.
2. Double-click setup.exe to launch the installer.

3. In the Select the language for the installation dialog box, select the language for the installation screens and click **OK**.
4. In the Welcome page, click **Next**.
5. In the License page, review the terms of the license agreement and select **I accept the terms in the license agreement** and click **Next** to continue the installation.
6. Complete the remaining pages in the wizard to set up your IBM Tealeaf CX configuration.

The following table lists the pages by name and describes the purpose of the page.

Table 28. Tealeaf CX Installer pages

Page Name	Description
Destination folder	Use this page to set the installation directory for IBM Tealeaf CX. If you want to install IBM Tealeaf CX to a location other than the default, click Change and select another location. Note: Do not include non-ANSI characters in the installation path.
Custom Setup	Use this page to customize the installation. The options in the drop-down list next to the component name allow you to control which IBM Tealeaf CX components are installed and whether to install all of the features for the component. Note: IBM Tealeaf validates the selected components. If an invalid configuration is selected, an error message is displayed. To continue the installation, verify that the appropriate components are selected for your installation. For example, the data preaggregator must be installed to a dedicated server without any additional IBM Tealeaf components. For information about the options on this panel, click Help .
Transport Pipeline	Use this page to set the configuration for the Transport Pipeline. <ul style="list-style-type: none"> • Default Transport Pipeline (No HBR): Select this option for an environment with a single Processing Server. • HBR Transport Pipeline (w/HBR): Health-Based Routing can be installed with the Transport Service. In multi-Processing Server environments, HBR can be deployed to manage load balancing and failover between Processing Servers. <p>The Transport Service receives data to process on a machine separate from the web server to avoid burdening the web server with additional tasks.</p>

Table 28. Tealeaf CX Installer pages (continued)

Page Name	Description
Master Management Server	<p>Use this page to specify the name or IP address of the Master Management Server. In most configurations, the Master Management Server is the server on which you install the IBM Tealeaf Portal.</p> <ul style="list-style-type: none"> • If you install everything (including Portal) on the same machine, then the Master Management Server is "localhost". • In a multiserver topology, where Portal is installed on a different machine, but all of the other components are installed locally, then when prompted, you must enter the name or IP address of the machine where Portal is installed. <p>Each machine that is subordinate to the Master Management Server must know the name of the master machine from which it can acquire its configuration information.</p> <p>The installer presents a prompt asking if you want to configure the Tracking service. If you click Yes, you are presented with the Tracking Service configuration panel.</p>
Tracking Service	<p>Specify a Server and Port and indicate whether you want to install the Tracking Service.</p> <p>The server and port information you enter is applied to the common registry used by IBM Tealeaf servers.</p> <p>Note: The registry is still available inside Tracking Service component, but other services look at the same registry location.</p> <p>During session replay:</p> <ul style="list-style-type: none"> • Replay Server uses the registry keys to get information about where the Tracking service is installed and which port it is running on. • Replay Server sends the logging data to the configured tracking service host and port. • The Data Service reads the registry keys and pulls information from the same logging service. <p>You can modify the Tracking Service configuration in TMS after installing IBM Tealeaf. For information about modifying the Tracking Service configuration using TMS, see the <i>IBM Tealeaf CX Configuration Manual</i>.</p> <p>Note: If there are multiple instances of the Tracking service, you must specify the same Server and Port on all of these servers. Each of the specified servers determines if it can start with the given server and port information. If a server cannot start, it terminates itself. This way, only one instance of logging server will be available.</p>
Report Data Collector	<p>Use this page to specify the Data Service Host Name or Machine IP address.</p> <p>Specifying the Data Service Host Name or Machine IP address allows the Replay Server connect to the Data Service and retrieve Server Information.</p> <p>This information is necessary for the Replay Server to determine which Replay Server is the master server.</p>

Table 28. Tealeaf CX Installer pages (continued)

Page Name	Description
Mail Server	<p>Use this page to specify the name or IP address of the SMTP Mail server.</p> <p>Providing this information is optional. If the mail server is not known, use the default mail value.</p> <p>You can change the value that you provide here after the installation is complete.</p>
Canister Destination folder	<p>Use this page to set the directory in which the IBM Tealeaf Canister is created.</p> <p>If you want to choose another directory, click Change... and select another folder.</p> <p>At run time, the system stores sessions in the Canister. In most cases, the Canister is located on a different drive than IBM Tealeaf, mainly because this drive requires a lot of space.</p>
Report Database	<p>Use this page to set the configuration for the Report Database.</p> <p>If you are installing the report database on the local server, verify the server name and the path for the installation directory and click Next.</p> <p>If you want to install the report database on a remote server, enter the name of the machine (or IP address), and the port number.</p> <p>At the Should the installer set up the SQL Reporting database(s)? screen, click Yes to install the database or click No to install the database at a later time.</p> <p>Note: If you select Yes to set up the Reporting database at install time, but do not have the required authorization to do so, the installation finishes with errors and the Reporting database is not set up.</p>
Online Help	<p>Use this page to set up access to the online help.</p> <p>If you do not know the user name and password to enter, you can leave the fields blank and enter them later through the IBM Tealeaf Portal.</p> <p>See "CX Settings" in the <i>IBM Tealeaf cxImpact Administration Manual</i>. When these values are configured, users can automatically access IBM Tealeaf Online Help by clicking Help the Portal.</p>

Table 28. Tealeaf CX Installer pages (continued)

Page Name	Description
Point of Contact	<p>Use this page to set contact information for the IBM Tealeaf Administrator at your site.</p> <p>The Point of Contact is the individual to contact when Tealeaf users have questions about user accounts or the Tealeaf installation.</p> <p>If you do not know this information, you can leave these fields blank and enter them later through the IBM Tealeaf Portal. For more information, see "CX Settings" in the IBM Tealeaf cxImpact Administration Manual.</p> <p>Note: The email address for the Tealeaf Administrator is not required, however IBM Tealeaf uses the specified email address to deliver important updates. Updates can contain information about system issues and configuration changes, including the automated disabling of events. If you do not currently have a designated IBM Tealeaf administrator, enter a valid email address and update after the installation is completed.</p>
Processing Servers	<p>Use this page to configure the Processing Server.</p> <p>Select the server name or IP address of the Tealeaf processing server (Canister)</p> <p>If you want to add another server, enter the server name or the IP address of the server and click Add. You can also select a server from the list and click Remove to delete the server from the list.</p>
Ready to Install the Program	<p>Use this page to install IBM Tealeaf CX using the configuration from the previous settings.</p>

7. At the Ready to Install the Program page, click **Install**.
The installer installs IBM Tealeaf CX with the specified configuration settings.
The display updates the status of the installation.
8. When the Installation Complete screen is displayed, click **Finish**.

You have successfully installed IBM Tealeaf CX.

You are now ready to enable your IBM Tealeaf solution.

Enabling your IBM Tealeaf solution involves running `setup.exe` files for the products that you have licensed and installing and configuring the Passive Capture Application server. For information on how to install the Passive Capture Application server, see the *IBM Tealeaf CX Passive Capture Application - PCA Manual*

Related concepts:

"IBM Tealeaf CX component architecture" on page 5

"IBM Tealeaf CX" on page 4

Tealeaf CX silent installation, uninstallation, and upgrade

Tealeaf CX installation uses an install wizard that prompts you for information. You can use a command line command to install CX instead of the install wizard. When you use the command line, you are not prompted for any information. You can use silent install to do a fresh installation or an upgrade.

Silent installation types

There are two ways to do a silent installation, uninstallation, and upgrade:

- Command and parameters - specify all of the options and parameters on a command line. You might specify all of the parameters on a command line if you are doing a one-time installation or uninstall.
- Command and file - specify the parameters and options in a file and call the file from the command line. You might use the file option to do multiple installs that have the same options.

Silent installation command

You use the `setup.exe` command to silently install CX.

Silent uninstallation command

You use the `msiexec` command to silently uninstall CX. You must run this command as Administrator.

Silent upgrade command

You use the `setup.exe` command to silently install CX.

Parameters

All property names and parameters are case sensitive. This table lists and describes the parameters that you can set for silent installation:

Parameter	Description
TARGETPASS	The path that is assigned to the <code>INSTALLDIR</code> property, if a value is not found in the registry. For example, <code>C:\Program Files\Tealeaf</code> . This parameter is used for installation only and is ignored by uninstall.
INSTALLPORTAL	Whether to install the portal. Valid values are: <ul style="list-style-type: none">• YES• NO This parameter is a CX installation parameter.

Parameter	Description
INSTALLDATACOLLECTOR	<p>Whether to install the data collector. Valid values are:</p> <ul style="list-style-type: none"> • YES • NO <p>This parameter is a CX installation parameter.</p>
INSTALLPREAGGREGATOR	<p>Installs the data preaggregator.</p> <p>Value: Yes installs the data preaggregator.</p> <p>The data preaggregator must be installed to a dedicated server without any other tealeaf components. Use the following values to make sure that no additional components are installed:</p> <ul style="list-style-type: none"> • NSTALLPROCESSOR=NO • INSTALLPORTAL=NO • INSTALLDATACOLLECTOR=NO • INSTALLREPLAYSERVER=NO • INSTALLTRANSPORT=NO • INSTALLPREAGGREGATOR=YES <p>Additionally, you must specify the following additional information:</p> <ul style="list-style-type: none"> • PORTALSERVER=<i>host name</i>: Identifies the IBM Tealeaf portal server address. • TMSERVER=<i>host name</i>: Identifies the IBM Tealeaf TMS server address. • REPORTDBSERVER=<i>host name</i>: Identifies the IBM Tealeaf reporting database server address. • DATACOLLECTOR=<i>host name</i>: Identifies the IBM Tealeaf data collector server address. • REPORTDBPORT=<i>port number</i>: Identifies the port number that is used to transfer data to the reporting database.
INSTALLPROCESSOR	<p>Whether to install the processor. Valid values are:</p> <ul style="list-style-type: none"> • YES • NO <p>This parameter is a CX installation parameter.</p>
INSTALLREPLAYSERVER	<p>Whether to install the replay server. Valid values are:</p> <ul style="list-style-type: none"> • YES • NO <p>This parameter is a CX installation parameter.</p>

Parameter	Description
INSTALLTRANSPORT	Whether to install the transport pipeline. Valid values are: <ul style="list-style-type: none"> • YES • NO This parameter is a CX installation parameter.
HBROPTION	Whether to install the HBR transport pipeline. Valid values are: <ul style="list-style-type: none"> • 1 - do not install HBR • 2 - install HBR This parameter is a CX installation parameter.
PROCSERVERLIST	A list of servers that are separated by a pipe. For example (SERVERNAME1) or (SERVERNAME1 SERVERNAME2). <p>This parameter is a CX installation parameter.</p>
PORTALSERVER	The portal server name. <p>This parameter is a CX installation parameter.</p>
TMSERVER	The TMS server name. <p>This parameter is a CX installation parameter.</p>
MAILSERVER	The mail server name. <p>This parameter is a CX installation parameter.</p>
REPORTDBSERVER	The report database server name. <p>This parameter is a CX installation parameter.</p>
REPORTDBPORT	The report database port number. <p>This parameter is a CX installation parameter.</p>
REPDBLOCATION	The report database location. Valid values are: <ul style="list-style-type: none"> • 1 - Local • 2 - Remote This parameter is a CX installation parameter.
REPDBINSTALL	Whether to install the report database. Valid values are: <ul style="list-style-type: none"> • 0 - do not install • 1 - install This parameter is a CX installation parameter.

Parameter	Description
CANISTERDIR	The Cannister directory path. This parameter is a CX installation parameter.
SQLDATA	The SQLDATA directory path. SQLDATA directory Path must be specified (if REPDBLOCATION = 2 (installed on Remote machine)), otherwise Report Database installation will fail. This parameter is a CX installation parameter.
SPOOLDIR	The Spool directory path. This parameter is a CX installation parameter.
DATACOLLECTOR	The remote data collector host name or IP address. This parameter is a CX installation parameter.
FORCETOREMOVE	Whether to force the uninstall. . Default value is NOValid values are: <ul style="list-style-type: none"> • YES - uninstall CX even other Tealeaf products are still installed • NO (Default) - do not Uninstall CX if any Tealeaf products are still installed This parameter is a CX uninstall parameter.
REMOVEDATABASE	Whether to remove the database. Valid values are: <ul style="list-style-type: none"> • 0 - do not remove the database • 1 - remove the database This parameter is a CX uninstall parameter.
SOURCEFILEPATH	The location of *.ini file that contains the list of properties for the install. If no value is assigned, the installer ignores this property.
TRACKINGHOST	The server where the Tracking service is installed. Default is localhost. During an upgrade, you can specify this parameter only if the tracking service is not already installed,
TRACKINGPORT	The port that the Tracking service will use. Default is 9001. During an upgrade, you can specify this parameter only if the tracking service is not already installed,

Parameter	Description
TRACKINGSERVICECONFIGURE	<p>Enable or disable the Tracking services. Valid values are:</p> <ul style="list-style-type: none"> • 0 - disable tracking service • 1 - default, enable tracking service <p>During an upgrade, you can specify this parameter only if the tracking service is not already installed,</p>
DBUPGRADEOPTION	<p>Whether to upgrade the database. Valid values are:</p> <ul style="list-style-type: none"> • 1 - upgrade the database • 2 - do not upgrade the database <p>During an upgrade, all other parameters are ignored.</p>

Silent installation or uninstall file

You can set the parameters for installation and uninstall in a file and call that file from a command line to do silent installation or uninstall. You create a .ini file for the values that you want to configure. You might use the file option to do multiple installs that have the same options.

This example shows the contents of a .ini file that installs transport but not data collector, portal, processor, or replay server, uses the HBR transport and defines two servers:

```
[COMMON]
TARGETPASS=C:\Program Files\Tealeaf

[CX]
INSTALLTRANSPORT=YES
INSTALLDATACollector=NO
INSTALLPORTAL=NO
INSTALLPROCESSOR=NO
INSTALLREPLAYSERVER=NO
HBROPTION=2
PROCSERVERLIST=server1|server2

[CXUNINSTALL]
FORCETOREMOVE=NO
REMOVEDATABASE=1
```

Tealeaf project GUIDs for unistall

When you uninstall a Tealeaf product, you must specify the project GUID. The GUIDs for the Tealeaf projects are:

- CX - {249E5938-69C9-4017-9C39-B7E53867F2E5}
- Upgrader - {173E74FF-DFFA-49F3-8E4B-AED9A4A492B8}
- RTV - {464592FF-2576-4A64-B2CB-A0F57095DC5D}
- cxView - {9DBE60C5-02E0-4601-A7A7-0C8A1C2DCC9E}
- cxVerify - {D1490BBD-1DDE-4772-936D-3640BFF43B67}
- cxReveal - {72D988CA-4F53-460A-8632-9543D8F10C92}
- cxOverstat - {7471B3BF-CC11-42BE-AC40-0C4E8455E0B0}

- cxMobile - {EFF86135-C1F6-481E-BB6B-EDF3BBF723DB}
- cxImpact - {0962511C-1FBF-492F-9A6A-BD732AAA1D7A}
- cxConnectWA - {B97DCE20-E3C9-443D-AA21-E33E18980A5A}
- cxConnectVOC - {D33F5260-B354-487B-979D-9F420D10296A}
- cxConnectDA - {9846E3C9-3457-4986-8C4F-9585C4FD7F3A}

Installing CX silently

You can silently install CX from the command line by specifying the parameters or by calling a file that has parameter values in it. You silently install when you do not want to go through the installation prompts, for example if you want to automate nightly installations.

Before you do a silent installation, you must identify the parameters and values that you want to specify on the command line or in the .ini file that you want to use for installation.

1. Log in to the server where you want to install CX.
2. Optional: To silently install CX from the command line without a .ini file, enter the setup.exe installation command and the parameters on the command line. For example, to specify a new Tealeaf installation folder, and install transport but not data collector, portal, processor, or replay server, use the HBR transport and define two servers, you might use this command:

```
"C:\src\current\Installations\TealeafCX_MSI\PROJECT_ASSISTANT\TealeafCX\
DiskImages\DISK1\setup.exe" /s /v"/qn TARGETPASS="C:\Program Files\
Tealeaf\" INSTALLTRANSPORT=YES INSTALLDATACOLLECTOR=NO INSTALLPORTAL=NO
INSTALLPROCESSOR=NO INSTALLREPLAYSERVER=NO HBROPTION=2
PROCSERVERLIST=server1|server2"
```

After the installation is complete, a completion message displays on the screen.

3. Optional: To silently install CX with a .ini file, enter the setup.exe installation command and specify the .ini file. For example, you created a silent.ini file in C: that specifies a Tealeaf installation folder, and install transport but not data collector, portal, processor, or replay server, use the HBR transport and define two servers. You might use this command to call the .ini file and install CX:

```
"C:\src\current\Installations\TealeafCX_MSI\PROJECT_ASSISTANT\TealeafCX\
DiskImages\DISK1\setup.exe" /s /v"/qn SOURCEFILEPATH="C:\Silent.ini\""
```

After the installation is complete, a completion message displays on the screen.

Upgrading silently

You can use the silent installation process to upgrade an existing installation.

Before you do a silent installation, you must identify the parameters and values that you want to specify on the command line or in the .ini file that you want to use for installation.

The TRACKINGHOST, TRACKINGPORT, TRACKINGSERVICECONFIGURE parameters can be used during Upgrade only if the Tracking Service is not installed. Otherwise, these parameters are ignored.

1. Log in to the server where you want to upgrade CX.
2. Optional: To run an upgrade and upgrade the database from the command line, enter the setup.exe installation command and the parameters on the command line. You might use this command:

```
"C:\src\current\Installations\Upgrade_MSI\PROJECT_ASSISTANT\Upgrader_MSI\
DiskImages\DISK1\setup.exe" /s /v"/qn"
```

3. Optional: To run an upgrade CX without upgrading the database, enter the setup.exe installation command and the parameters that you want to use. You might use this command:

```
"C:\src\current\Installations\Upgrade_MSI\PROJECT_ASSISTANT\Upgrader_MSI\
DiskImages\DISK1\setup.exe" /s /v"/qn DBUPGRADEOPTION=2"
```

Enabling your IBM Tealeaf solution

After you install IBM Tealeaf CX you can enable it for use by running setup.exe files.

Running a setup.exe file unlocks features and functions (in most cases through a license mechanism) specific to the IBM Tealeaf solution that you will use manage your customer's experience.

Some IBM Tealeaf solutions can be enabled only if IBM Tealeaf Portal and IBM Tealeaf CX components are installed, while other IBM Tealeaf products can be enabled regardless of whether Portal is installed.

The following IBM Tealeaf solutions can be enabled only if you have installed IBM Tealeaf Portal:

- IBM Tealeaf cxImpact
- IBM Tealeaf cxConnectVOC
- IBM Tealeaf cxConnectWA
- IBM Tealeaf cxMobile cxOverstat cxView.

Running the setup.exe files for these products registers them and enables them for use.

You must run the setup.exe files on the server hosting the IBM Tealeaf Portal.

The following IBM Tealeaf solutions can be enabled regardless of whether you have installed IBM Tealeaf Portal:

- IBM Tealeaf cxConnectDA
- IBM Tealeaf cxReveal
- IBM Tealeaf cxVerify

You can run the setup.exe files for these products locally (on the same machine as IBM Tealeaf CX) or on a different machine.

- If you run the setup.exe files for these products locally and you have installed IBM Tealeaf Portal, you can then use Portal to register and enable the products for use.

If you have not installed Portal, you can still enable these products for use by running the setup.exe files locally. Enabling products for use in this manner involves changing configuration files manually.

- If you are not running the setup.exe files for these products locally (meaning you are not enabling them for use on the same machine on which IBM Tealeaf CX is installed), then you can enable them by running the setup.exe on a separate machine.

The following table lists the IBM Tealeaf solutions and the publications needed to install and enable the solutions for use.

Table 29. Publications with installation instructions

To enable this IBM Tealeaf solution	Go to this publication for instructions	Notes
IBM Tealeaf cxReveal	<i>IBM Tealeaf cxReveal Administration Manual</i>	IBM TealeafcxReveal enables features and functions that use core functionality of the Tealeaf CX platform — session replay and search — and packages it into a simple user interface for efficient use by customer service representatives, sales associates and marketing personnel.
IBM Tealeaf cxConnect for Data Analysis	<i>IBM Tealeaf cxConnect for Data Analysis Administration Manual</i>	IBM Tealeaf cxConnect for Data Analysis integrates IBM Tealeaf customer experience dataset with any business intelligence or reporting application to create a multichannel view of the customer for ongoing analysis.
IBM Tealeaf CX RealiTea Viewer	<i>IBM Tealeaf RealiTea Viewer User Manual</i>	For servers that are hosting Tealeaf software, you may wish to install RTV for troubleshooting purposes. Note: It is recommended that you use IBM Tealeaf cxImpact Browser Based Replay (BBR) for most replay tasks. You can access BBR through the Portal web application. IBM Tealeaf CX RealiTea Viewer requires administrator level access to the machine on which it is being installed. It can be installed on the local Windows desktop of Tealeaf users. For more information on installation, see "RealiTea Viewer Overview" in the <i>IBM Tealeaf RealiTea Viewer User Manual</i> .
IBM Tealeaf CX Passive Capture Application	<i>IBM Tealeaf CX Passive Capture Application Manual</i>	IBM Tealeaf Passive Capture Application provides the capability to capture requests and responses of your website's traffic. IBM Tealeaf CX Passive Capture Application runs on Linux only. See Table 30 for a list of commands that are useful when installing and configuring IBM Tealeaf CX Passive Capture Application.

Table 30. Useful Linux commands

To do this. . .	Use this command
Change directory	cd [desired directory]
Move a file	mv [file name] [destination directory]
Rename a file	mv [current file name] [new file name]

Table 30. Useful Linux commands (continued)

To do this. . .	Use this command
Remove a file	<code>rm [file name]</code>
List files in the current directory	<code>ls</code> or <code>ls -la</code>
Display current location	<code>pwd</code>
Check for traffic	<code>tcpdump -n -i eth[specify interface]</code>
Check for traffic on a port	<code>tcpdump -n -i eth[specify interface] port [specify port -optional] host [specify host ip -optional]</code>
Check for traffic excluding ports and/or hosts	<code>tcpdump -n -i eth[specify interface] and not port [specify port -optional] and not host [specify host ip -optional]</code>
Display interface details	<code>ifconfig</code>
Install Tealeaf RPM file	<code>rpm -ivh [tealeaf rpm file]</code>
Update the Tealeaf RPM file	<code>rpm -uvh [tealeaf rpm file]</code>
Use Sudo Root	<code>sudo su-</code>
Change Ethernet settings (full duplex, etc.)	<code>ethtool eth[specify interface]</code>

If you are installing the Portal Server, the Installer automatically disables the Interactive Services Detection service, which causes conflicts with the BBR rendering instances. See *Preparing the network environment* for information about the Interactive Services Detection service.

The steps for enabling a product for use by running its `setup.exe` are same regardless of the product. To enable an IBM Tealeaf solution:

1. In the software distribution, navigate to the solution you are enabling.

Table 31. Tealeaf solution and software distribution path

If you are enabling this solution	Navigate to	Notes
IBM TealeafcxImpact	Enablers\cxImpact	IBM Tealeaf cxImpact enables IBM Tealeaf features and functions that transform the dataset captured by IBM Tealeaf CX into visually re-playable and completely searchable customer sessions.
IBM TealeafcxConnect for Voice of Customer	Enablers\cxConnectVOC	IBM Tealeaf cxConnect for Voice of Customer provides integration to the IBM Tealeaf customer experience dataset with any Voice of Customer (VOC) application.
IBM TealeafcxConnect for Web Analytics	Enablers\cxConnectWA	IBM Tealeaf cxConnect for Web Analytics provides you with a product that integrates the IBM Tealeaf customer experience dataset with any web analytics application to help uncover trends or anomalies such as a difference in conversion rates between two time periods or customer segments.

Table 31. Tealeaf solution and software distribution path (continued)

If you are enabling this solution	Navigate to	Notes
IBM TealeafCX Mobile	Enablers\cxMobile	IBM TealeafCX Mobile modules capture, replay, and analyze the experiences of visitors those are accessing your web application from mobile devices.
IBM TealeafcxOverstat	Enablers\cxOverstat	IBM Tealeaf cxOverstat provides a product that identifies problematic hotspots, optimizes page count and improves form conversion rates.
IBM TealeafcxVerify	Enablers\cxVerify	IBM Tealeaf cxVerify preserves a record of online customer interactions on a website or mobile device. It helps online businesses maintain highly reliable records for dispute resolution, fraud investigations and audit and compliance purposes.
IBM TealeafcxView	Enablers\cxView	IBM Tealeaf cxView gathers the customer experience dataset of IBM Tealeaf cxImpact and puts it into executive-level dashboards, scorecards and reports.

2. Double-click setup.exe.
3. Select the language that you want to use for the installation wizard pages and click **OK**.
4. In the License screen, review the license terms. If you agree to the license terms, select **I accept the terms** and click **Next**.
5. When the **Begin Installation** page is displayed, click **Install** to start the installation process.
6. When the Installation Complete page is displayed, click **Finish**.

You have successfully enabled the IBM Tealeaf solution(s).

The following native application frameworks are installed on the Portal Server as a result of installing IBM Tealeaf CX Mobile:

- Tealeaf Android Logging
- Tealeaf iOS Logging

These logging frameworks are installed as ZIP files in the following location on the Portal Server:

<Tealeaf_install_directory>\NativeMobile SDK

See "Overview of CX Mobile Configuration" in the *IBM Tealeaf CX Mobile Administration Manual*.

Related tasks:

"Preparing the network environment" on page 73

Uninstalling CX silently

The uninstall wizard prompts you for information on the parts of CX that you want to uninstall. You can use the `msiexec` to uninstall CX and not use the uninstall wizard.

Before you uninstall CX silently, you must identify the component that you want to uninstall and have the GUID for the component.

You must run the commands in this task as Administrator.

1. Log in to the server where you want to uninstall CX.
2. Optional: To silently uninstall CX from the command line without a `.ini` file, enter the `msiexec` installation command and the parameters on the command line. For example, to uninstall CX and remove the database, you might use this command:

```
msiexec /X{249E5938-69C9-4017-9C39-B7E53867F2E5} /qn REMOVEDATABASE=1
```

After the components are uninstalled, a completion message displays on the screen.

3. Optional: To silently uninstall CX with a `.ini` file, enter the `msiexec` installation command and specify the `.ini` file. For example, you created a `silent.ini` file in `C:`. You might use this command to call the `.ini` file and uninstall CX:

```
msiexec /X{249E5938-69C9-4017-9C39-B7E53867F2E5} /qn SOURCEFILEPATH="C:\TestSourceFile.ini"
```

After the components are uninstalled, a completion message displays on the screen.

Post-installation tasks

After you complete your installation, specify a backup strategy, recovery model, and log file compression settings.

Post-installation database tasks

Depending on the database environment required for your installation, you might need to perform some post-installation database tasks.

Decide which database tasks you need to perform.

Table 32. Database task reference

Task	Description	For more information see . . .
Installing the database	If you elected to not install the databases through the Tealeaf Installer, you should install them as soon as you can.	"Using Tealeaf Database Manager" in the <i>IBM Tealeaf Databases Guide</i> .
Loading the event definitions	If you have installed databases individually and one of them includes the System database, you must install or reinstall the event definitions through the Tealeaf Database Manager.	"Installing Tealeaf Databases" in the <i>IBM Tealeaf Databases Guide</i> .
Implementing a database backup strategy	The Tealeaf databases should be regularly backed up. During installation, backup scripts are automatically generated for use in the Tealeaf installation tree.	"Database Backup Strategy" in the <i>IBM Tealeaf Databases Guide</i> .

Table 32. Database task reference (continued)

Task	Description	For more information see . . .
Setting recovery model and logging levels	Tealeaf recommends using the Simple database recovery model and compressing log files for the Tealeaf databases.	"Database Sizing "in the <i>IBM Tealeaf Databases Guide</i>
Setting the file path on a remote SQL Server	<p>If you installed the SQL Server data files on a remote server, you must validate the file path to these files:</p> <p>All data files are in a single location on one server: In this case, you must specify the path after installation completes. In the registry, update the following value with the correct location of the files:</p> <p>HKLM->Software->Tealeaf Technology->Datastore->ReportServer->MSSQL Data</p> <p>All data files are in multiple locations on the remote server: During database upgrades that manipulate these files, you must run the Tealeaf Database Installer in Advanced mode. For each database, the file paths must be verified and specified on the Filegroups tab.</p>	

Verifying system startup

After you complete the installation process and configure the components on the respective servers, you can start IBM Tealeaf CX.

The procedure for verifying system startup vary depending on the server topology of your IBM Tealeaf installation environment.

This section provides instructions for:

- Verifying system start up for single server topology.
- Verifying system start up for a multi-server topology.

Related concepts:

“Topology variations” on page 43

Verifying system start up for single server topology

After you complete the installation process and configure the components on a single server, you can start IBM Tealeaf.

A single server installation results in a topology in which all IBM Tealeaf services are installed on a single machine.

To verify the system startup for a single server topology:

1. Login to the machine and select the following Windows Start menu shortcut:
All Programs > Tealeaf Technology > Start Tealeaf Services.
 The services for startup are displayed in a window.
 When the services are all started and the window closes.
2. Open a browser and navigate to the following URL:
<http://localhost/portal>

The IBM Tealeaf Portal login screen is displayed.

3. Login using the administrator account information provided to you by IBM Tealeaf.

The Portal Administration page is displayed.

Note: If you are unable to login or you encounter permission issues related to IIS, you might be able to rectify the problem by reinstalling the IBM Tealeaf CX Portal application or web service. See *Using the IBM Tealeaf web application installation utility* in the Troubleshooting chapter.

You have successfully logged into the Portal, which means that the core web application is operational.

Start the IBM Tealeaf CX Passive Capture Application. See "Installation" in the *IBM Tealeaf Passive Capture Application Manual*.

When the IBM Tealeaf CX Passive Capture Application started, verify that session data is being captured and processed through the **Pipeline Status** utility in the Tealeaf Management System. See "TMS Pipeline Status Tab" and "Tealeaf Management System" in the *IBM Tealeaf cxImpact Administration Manual*.

Related concepts:

"Topology variations" on page 43

Related tasks:

"Using the IBM Tealeaf web application installation utility" on page 123

Verifying system start up for a multiserver installation environment

After you complete the installation process and configure the components on multiple servers, you can start IBM Tealeaf CX.

A multiserver installation results in a deployment environment in which Tealeaf services are installed on different servers.

To verify the system startup of a multiserver Tealeaf installation environment:

1. Start the Processing Server(s).
Login to each machine, and start all Tealeaf services.
If a Processing Server is a TMS slave server, you may see event log messages warning that the slave server is unable to connect to the TMS master. These can be ignored for the moment.
2. Start the TMS Master.
Login to the machine, and start all Tealeaf services. The TMS master performs some initial configuration after the first startup.
3. Start all remaining servers.
Login to the machine, and start all Tealeaf services.
Wait until all Tealeaf services have started and the window that lists the services closes.
4. Open a browser and navigate to the following URL:
`http://localhost/portal`
The Tealeaf Portal login screen is displayed.
5. Login using the administrator account information provided to you by Tealeaf.

The Portal Administration page is displayed.

You have successfully logged into the Portal, which means that the core web application is operational. See "Logging in as Admin" in the *IBM Tealeaf cxImpact Administration Manual*.

Start the IBM Tealeaf Passive Capture Application. See "Installation" in the *IBM Tealeaf Passive Capture Application Manual*.

When the PCA has been started, you can verify that session data is being captured and processed through the Pipeline Status utility in the Tealeaf Management System.

See "TMS Pipeline Status Tab" in the *IBM Tealeaf cxImpact Administration Manual*.

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

Related concepts:

"Topology variations" on page 43

Traffic requirements

The IBM Tealeaf CX Passive Capture Application requires a minimum of a bidirectional traffic stream or two unidirectional traffic streams containing all HTTP request and response traffic between your web application(s) and the visitor's browsers that interact with your web applications.

If the data stream coming into the IBM Tealeaf CX Passive Capture Application is not complete (or "clean") you might get the erroneous impression that Tealeaf is not working correctly.

Basic traffic requirements

To capture traffic, the PCA needs to see the start of all TCP connections.

To allow monitoring of a complete HTTP(S) conversation, the PCA requires that the mirrored network traffic be of very high integrity and quality. Any loss of critical network TCP packets can prevent the PCA from reassembling the TCP traffic into HTTP hits.

Lost TCP packets may result in Tealeaf sessions with missing pages, partial pages or both. In a worst-case scenario, the entire session may be unusable.

Check with your IT team to confirm if HTTP persistent connections have been enabled in the IT infrastructure.

Individual HTTP persistent connections may be used by multiple visitors to your web application and may be deployed by a load balancer such as an F5 network device, a front-end proxy such as an Akamai server or the web server itself.

HTTP persistent connections, which can also be called HTTP keep-alive, or HTTP connection reuse, are the idea of using the same TCP connection to send and receive multiple HTTP requests/responses, as opposed to opening a new one for every single request/response pair.

The Passive Capture Application requires that it sees the start of all HTTP/TCP connections. If HTTP persistent connections are enabled then the PCA will not be able to reassemble hits from in-progress connections.

- **Traffic stream:**

The PCA requires a minimum of one bidirectional traffic stream or two unidirectional traffic streams containing all HTTP request and response traffic between the web application and the visitor's browser interacting with it.

- **No errors or dropped packets:**

No errors, dropped packets, or overrun packets at operating system network interface card and network level.

An **ifconfig ethx** command on the capture server should display a constant number of dropped packets or errors. The X in "ethx" will be the number of the NIC card, e.g. ifconfig eth0.

If the number is increasing at a high rate, there may be problems with the fidelity of the traffic sent to the PCA, inadequate sizing of your PCA hardware for your traffic volume, or both.

- **Real visitor IPs:**

The capture point can see the real visitor IPs or host address of visitor's IP.

Access to the real IP address of your visitors is a valuable resource for troubleshooting purposes. For customers using load balancers, this requirement may not be possible.

- **Filtered traffic:**

Spanned traffic is filtered down to the essential traffic only.

It is recommended that you filter out as much unnecessary traffic as possible at the network level before it is delivered to the PCA. Filtering off-loads processing resources that the PCA has to use to filter out traffic.

- **TCP persistent connections disabled:**

If HTTP persistent connections are enabled then the PCA will not be able to reassemble hits from in-progress connections.

TCP Connections Sources of Traffic

SPAN PORT

A SPAN port is also known as "Port Mirroring".

Port Mirroring is used on a network switch to send a copy of network packets seen on one switch port (or an entire VLAN) to a network monitoring connection on another switch port.

This is commonly used for network appliances that require monitoring of network traffic, such as an intrusion-detection system.

Port mirroring on a Cisco Systems switch is generally referred to as Switched Port Analyzer (SPAN); some other vendors have other names for it e.g. Roving Analysis Port (RAP) on 3Com switches.

NETWORK TAP

A network tap is a hardware device which provides a way to access the data flowing across a computer network.

In many cases, it is desirable for a third party to monitor the traffic between two points in the network. If the network between points A and B consists of a physical cable, a "network tap" may be the best way to accomplish this monitoring.

The network tap has (at least) three ports: an A port, a B port, and a monitor port. A tap inserted between points A and B passes all traffic between A and B through unimpeded but also copies that same data to its monitor port. This enables a third party to listen.

Network taps are commonly used for Tealeaf, network intrusion detection systems, VoIP recording, network probes, RMON probes, packet sniffers, other monitoring and collection devices and software that require access to a network segment. Taps are used in security applications because they are non-obtrusive, are not detectable on the network (having no physical or logical address), can deal with full-duplex and non-shared networks and will usually pass through traffic even if the tap stops working or loses power.

SPAN PORT AGGREGATOR

SPAN Port Aggregation is a technology that combines a bidirectional full duplex data transmission into one single stream of data.

Additionally, aggregation can allow for the combination of data transmitted from multiple networks or SPAN ports.

An identical copy of this single stream of data can then be sent to any connected monitoring device. The connected monitoring device can receive the entire full duplex conversation or aggregate data from multiple networks with a single network interface card (NIC) also without having to reassemble the traffic Port Requirements

Port Definition and Configuration

Your IT team might be required open various TCP ports to enable communication with IBM Tealeaf.

The types of communications can include:

- Source Request
- Destination server to handle request
- Port Number: The port number which needs to be opened.

Table 33. Port details

Source	Destination	TCP Port
Capture Server (Linux)	CX Processing Server	1966
Tealeaf End User desktop	CX Processing Server (for session retrieval)	19000
Tealeaf End User desktop	CX Reporting/Portal (for the Tealeaf Web Portal)	80
Tealeaf Administrator desktop	to Processing and Portal/Reporting Servers (for Terminal Services access to the server)	3389
Tealeaf Administrator desktop	Capture Server (for web interface setup and administration of the Capture Server)	8080 and/or 8443
Tealeaf Administrator desktop	Capture Server (for SSH setup and administration of the Capture Server)	22
CX Portal/Reporting Server	SMTP server (for Scorecard, Alerts e-mails)	25
CX Portal/Reporting Server	X Processing Server (for Tealeaf Management Server)	20000

Starting and verifying the Passive Capture Application server is working

After installing IBM Tealeaf CX Passive Capture Application you can start it and verify that it is functioning properly.

Before you starting and verify that the Passive Capture Application server is working, verify that IBM Tealeaf CX Passive Capture Application installed correctly. See "Validate PCA Install" in the *IBM Tealeaf CX PCA Manual*.

The *IBM Tealeaf CX PCA Manual* includes additional information about starting the PCA and configuring it for use.

To start and verify the IBM Tealeaf CX Passive Capture Application:

1. Start the IBM Tealeaf CX Passive Capture Application
 - a. Open the passive capture web console by going to `http://www.pca_host_ip.com:8080/`.
 - b. Click **Console** tab.
 - c. Click **Start** to start the IBM Tealeaf CX Passive Capture Application server and to begin sending hits to the IBM Tealeaf appliance.
2. Verify the IBM Tealeaf CX Passive Capture Application is handling traffic.
 - a. Click the **Summary** tab.
 - b. Scroll and check **Filtered traffic**.
 - c. Verify that **kbytes/sec** is greater than zero.

0	22 kbytes/sec	Filtered traffic kbytes/sec
3	0 packets	If non-zero, packets are being dropped because they exceed the max size limit.

In a healthy production system, you might also see something like this:

Instance Compound Statistics

ID	Status	Description
0	0.75 %	The percentage of alien packets
0	false	If true, reasstd cannot keep up with listend.
0	0.06 %	The percentage of dropped packet connections
0	1 %	The percentage becoming unidirectional traffic
0	65 hits/sec	The rate reasstd is currently reassembling non-SSL hits.
0	false	If true, encountered Diffie Hellman SSL
0	0.02 %	The percentage of aged connections
0	0 keys/sec	Missing SSL keys/sec
0	3312 kbytes/sec	Filtered traffic kbytes/sec
0	0 packets	If non-zero, packets are being dropped because they exceed the max size limit.

Verifying system traffic

As part of system start up, verify that Tealeaf is processing traffic.

Verifying system traffic involves checking the pipeline status.

To verify system traffic:

1. If you have not already done so, log into the Tealeaf Portal.
2. From the menu bar, select **Tealeaf > TMS**.
3. Click the **Pipeline Status** tab.
4. From the **Pipeline Status** confirm that traffic is being received.

Look for non zero entries in the hit/sec column that go all the way from the top line down to the canister line.

The screenshot shows the IBM Tealeaf CX interface. The top navigation bar includes 'Dashboards', 'Active', 'Search', 'Analyze', 'Configure', 'Tealeaf', and 'Admin'. Below this, there are tabs for 'WorldView', 'Jobs', 'Pipeline Status' (which is highlighted with a red box), and 'Advanced'. The 'Server' is set to 'QABASE070' and the 'Update Interval (secs)' is set to '2'. The main content area is divided into two sections: 'Connections' on the left and a data table on the right. The 'Connections' table has columns for Slot #, Machine, Pg/Sec, Char/Sec, and Protocol. The data table on the right has columns for Active, Char Count, Char/Sec, Hit Count, Hit/Sec, Queued, and Max. The 'Canister' row is highlighted in yellow, and the 'Char/Sec', 'Hit Count', and 'Hit/Sec' columns for all rows are highlighted in red.

Active	Char Count	Char/Sec	Hit Count	Hit/Sec	Queued	Max
CSS_1966	150,944,983,024	1,261,726	2,373,074	55	0	0
DataDrop	150,903,084,480	1,261,726	2,371,637	55	0	0
DecoupleEx	151,000,953,040	1,261,734	2,396,966	55	0	494
Inflate	151,003,850,584	1,261,734	2,396,966	55	0	0
PrivacyEx	151,003,850,584	1,261,734	2,396,966	55	0	0
TLTRef	152,376,745,800	1,261,845	2,396,966	55	0	0
SessionRouter	152,269,070,304	1,261,845	2,371,637	55	0	0
Canister	19,625,285,155	1,261,845	2,371,637	55	0	0
Null	0	0	0	0	0	0

Figure 6. Verifying traffic from the Pipeline Status tab

If there are only a few hits every minute that stop part way down, these are statistics hits from the PCA server and are not “real” HTTP traffic.

If there are zero’s in all columns, then no traffic is being received from the PCA.

Verifying system status

You can use the Tealeaf portal to monitor and verify system status.

To verify system status:

1. If you have not already done so, log into the Tealeaf Portal.
2. From the menu bar, select **Tealeaf > System Status**.
3. Select the type of systems status to check. You can perform the following types of checks:
 - Status Summary
 - Canister
 - DecoupleEX
 - Storage
 - Health Based Routing
 - Database Filegroup Size
 - Database Table Size
 - Data Collector Statistics
 - Fact Collection Rate
 - Collection Status

Backing up the IBM Tealeaf event model

Through the Portal, you can define tasks to automatically create backups of critical event model data. As needed, Portal backups can be restored, returning your Tealeaf® solution to a previously saved state.

For information about how to back up the IBM Tealeaf event model, see "Backing up the IBM Tealeaf event model" in the *IBM Tealeaf cxImpact Administration Manual*.

Rebuilding the canister

You can remove all Tealeaf captured session data from the canister by running the Canister Rebuild utility (CanRebuild.exe). Running the Canister Rebuild utility reverts back to an installation with no captured session data present.

Rebuilding the canister is optional.

You might want to run the Canister Rebuild utility after you have run tests to verify a proof of concept that your Tealeaf installation is working and before you go into production with your solution.

Note: Rebuilding a canister is an unrecoverable process. The results are final and cannot be undone. Before you rebuild the canister, make sure that you do not need the session data in the canister.

To rebuild a canister:

1. Go to the Tealeaf installation directory and locate the Canister Rebuild executable CanRebuild.exe.
2. Double-click CanRebuild.exe.
3. Select **Rebuild Full Canister** and **Preserve Session Data**.
4. Click **Rebuild** and click **Yes** to confirm the rebuild operation.

Uninstalling IBM Tealeaf

You can remove IBM Tealeaf CX and IBM Tealeaf products from your environment by uninstalling them.

To uninstall IBM Tealeaf CX and IBM Tealeaf products, use **Add and Remove Programs** or **Programs and Features** from the Control Panel.

Consult Table 34 for a listing of Tealeaf component publications that include instructions for uninstalling.

Table 34. Uninstalling Tealeaf Components - Documentation reference

Component to uninstall	Documentation
IBM Tealeaf CX Databases	<i>IBM Tealeaf CX Database Guide</i>
IBM Tealeaf cxReveal	<i>IBM Tealeaf cxReveal Administration Manual</i>
IBM Tealeaf CX Passive Capture Application	<i>IBM Tealeaf CX Passive Capture Application</i>
IBM Tealeaf cxVerify	<i>IBM Tealeaf cxVerify Administration Manual</i>
IBM Tealeaf cxConnect for Data Analysis	<i>IBM Tealeaf cxConnect for Data Analysis Administration Manual</i>

Chapter 5. Configuring IBM Tealeaf

After you have installed IBM Tealeaf, you must complete additional configuration tasks to fully prepare your runtime environment.

Configuring IBM Tealeaf components

Depending on which IBM Tealeaf components you enabled, you might need to configure services, features, and component functionality.

For a multi-machine implementation, configure all the machines before starting any of the IBM Tealeaf services.

You can find information about configuring IBM Tealeaf components in the *IBM Tealeaf CX Configuration Manual*.

Common configuration tasks include:

- **Configuring the System Timezone:**

IBM Tealeaf requires that a single time zone be defined across all IBM Tealeaf servers in the system. For some IBM Tealeaf operations such as searching, the time zone may change the meaning of parameters such as today or yesterday. Among other features, this system-wide time zone is used as the basis for determining when scheduled reports are executed and delivered.

- **Configuring the Alert Service:**

This service manages execution and delivery of event-based alerts. You can configure this service to send email messages depending on threshold values defined for the event or Top Mover.

- **Configuring the Transport Service:**

The Transport component is responsible for accepting hits from the Capture Server, performing a series of pipeline operations, and then delivering the hit to the Processor component

- **Configuring the CX Canister:**

After installation, you may need to perform additional configuration of the IBM Tealeaf CX Canister for single server or multi-server installations.

Configuration settings can be used to deploy the Canister and indexing functions across multiple servers, or you can install multiple Canisters on the same machine or across multiple machines.

- **Configuring CX Indexing:**

A session index is a database that stores the locations of meaningful words and fields in each session. Because an index does not contain all text from each session, it can hold a large quantity of session information in a single file.

- **Configuring the Report Server:**

The Report Server consists of the Portal Web Application, the databases, and the Data Service. You can configure the Report server timezone and change various Report server configuration settings.

- **Configuring the Search Server:**

Search Server implements several low-level functions used by the IBM Tealeaf system to retrieve session data and to monitor the systems that maintain it. You can configure Search server settings.

- **Configuring the Scheduling Service:**

The IBM Tealeaf Scheduling Service can be used to schedule repeated Tealeaf-specific jobs, which include TLI archiving, backups, and extractions. During installation, the Tealeaf Scheduling Service is configured to automatically start up, yet some default jobs are disabled. You can change these settings as you see fit.

For information about configuring the Tealeaf Event bus, see the *IBM Tealeaf cxConnect for Data Analysis Administration Manual*.

For information about configuring the Tealeaf Cookie Injector, see "Installing and Configuring the Tealeaf Cookie Injector" in the *IBM Tealeaf Cookie Injector Manual*.

Configuring Portal settings

After you log in to the Portal, you can configure it for use.

Table 35 lists the initial configuration tasks for the Portal. Instructions on how to perform the configuration tasks are documented in the *IBM Tealeaf cxImpact Administration Manual*.

Note: You can use the Tealeaf Management System (TMS) to perform many of the operations for configuring Tealeaf servers and services. Tealeaf Management System (TMS) provides users and administrators with a centralized component for administering the IBM Tealeaf system. For information about using TMS, see "Tealeaf Management System" in the *IBM Tealeaf cxImpact Administration Manual*.

Table 35. Portal configuration tasks

Task	Description
Configuring the Portal announcements	<p>Portal announcements are messages displayed to all Portal users upon login. These messages can be used to display current system status, scheduled maintenance windows, or other Tealeaf-related issues.</p> <p>With each installation or upgrade, it is recommended that you create a Portal announcement message to your users. The announcements provide information as to the current stage of the machine, its state, and whether there are any current issues that might impact Tealeaf users.</p>
Configuring miscellaneous settings for Portal	<p>In the Miscellaneous settings panel, you can define a variety of settings, including the Tealeaf administrator's contact information.</p> <p>See "CX Settings" in the <i>IBM Tealeaf cxImpact Administration Manual</i>.</p>
Configuring Portal authentications	<p>By default, the Portal is configured to use Portal authentication, which leverages the Portal's user management capabilities to control access.</p> <p>Additionally, you can configure the following authentications:</p> <ul style="list-style-type: none"> • NT (AD - Active Directory) authentication • Mixed Mode authentication • Web Services authentication • Single Sign On (SSO) Authentication

Table 35. Portal configuration tasks (continued)

Task	Description
Identifying additional servers to the Portal	<p>During the installation process, reference information for any installed Reporting Server, Replay Server, Visitor Server, and any Canister Servers(s) are inserted into the database for use by the Portal.</p> <p>If there are other Tealeaf servers in the environment, the Portal must be made aware of them. See "Managing Tealeaf Servers" in the <i>IBM Tealeaf cxImpact Administration Manual</i>.</p>
Creating user accounts for the Portal	<p>You can create user groups for the following Tealeaf products:</p> <ul style="list-style-type: none"> • cxImpact For information, see "CX User Administration" in the <i>IBM Tealeaf cxImpact Administration Manual</i>. • cxView For information, see "cxView User Administration" in the <i>IBM Tealeaf cxImpact Administration Manual</i>. • cxReveal For information, see "cxReveal User Administration" in the <i>IBM Tealeaf cxReveal Administration Manual</i>.
Configure and define reporting mechanisms for the Portal	<p>Reporting mechanisms include:</p> <ul style="list-style-type: none"> • Scorecards Scorecards are used to report in graphical or tabular format useful metrics on data captured by Tealeaf For more information on configuring scorecards, see "Configuring Scorecards" in the <i>IBM Tealeaf cxView User Manual</i>. • Dashboards Dashboards can be used to arrange multiple reporting components into a single page. For more information on configuring dashboards, see "Configuring Dashboards" in the <i>IBM Tealeaf cxView User Manual</i>. • Tealeaf Status report Tealeaf Status report delivers status information on the Portal and other components of the Tealeaf system. For more information on configuring the Tealeaf Status report, see "Tealeaf Status Report" in the <i>IBM Tealeaf cxImpact Administration Manual</i>. • Search and List Templates Search and List Templates can be configured to enable specific search fields for specified groups. For more information on configuring search templates, see "Configuring Search Templates" in the <i>IBM Tealeaf cxImpact Administration Manual</i>. <p>For more information on configuring reports, see "Report Configuration" in the <i>IBM Tealeaf cxImpact Administration Manual</i>.</p>

The following are examples or Portal announcement messages:

- Stage: Proof of Concept
Stage: Proof of Concept
State: Functional
Notes:
This system is intended to demonstrate Tealeaf capabilities only. This system should not be used for any functions other than demonstrating system functionality.

For more information, contact your Tealeaf administrator. From the Portal menu, select Help > Contact Tealeaf Administrator.
- Stage: Development
Stage: Development
State: Functional
Notes:
This system is under construction. Users may experience problems using the system.

For more information, contact your Tealeaf administrator. From the Portal menu, select Help > Contact Tealeaf Administrator
- Stage: Staging
Stage: Staging
State: Functional
Notes:
This system is currently being tested for production release. Please report any problems.

For more information, contact your Tealeaf administrator. From the Portal menu, select Help > Contact Tealeaf Administrator.
- Stage: Production

Configuring the IBM Tealeaf system time zone

In the IBM Tealeaf system, the time zone is used by all Tealeaf servers to synchronize a variety of tasks.

During the installation process, this system-wide setting must be defined and applied to each server in the environment.

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

Configuring IBM Tealeaf data objects

After you perform the initial configuration tasks for the Tealeaf Portal, you must configure the data objects provided by Tealeaf.

Data objects are configured in the Tealeaf Event Manager, a Portal-based interface for creating event and event-related objects for deployment in the stream of data captured and processed by Tealeaf.

The topics in this section explain how to:

- Access the Tealeaf Event Manager
- Configure the Login ID to be searchable
- Remap components to reference session attributes
- Configure Tealeaf dimensions
- Create deviations

After configuring the IBM Tealeaf data objects you can use them to verify aspects of system operation.

Accessing the Event Manager to configure data objects

To configure Data objects you must access the Event Manger.

To access the Event Manager:

1. Login to the Portal using the admin account.
2. From the Portal menu, select **Configure > Event Manager**.
3. The Events tab of the Tealeaf Event Manager is displayed.

From the Events tab you can:

- Configure the Login ID to be searchable
- Remap components to reference session attributes

Configuring the Login ID to be searchable

As part of the set of provided data objects, IBM Tealeaf includes a *hit* attribute and an event for detecting the Login ID displayed in your web application. You must configure data objects to detect the Login ID.

A hit attribute is used to define the patterns in request or response data that demarcate an element of data that you wish to track through an event.

An event is triggered by a condition. In this case, the condition is the presence of the Login ID hit attribute. When this hit attribute is detected, the event is fired, which stores the Login ID value as the first session attribute (Login ID).

IBM Tealeaf supports the creation of up to 64 session attributes. A session attribute is a session-level variable that can be populated and updated based on events.

To enable these data objects, you must configure them to detect the Login ID's that are published by your web application. In this End-to-End scenario, you can complete the steps required to configure the Login ID hit attribute, event, and session attribute and then surface this data for use in search and search results.

To configure the Login ID to be searchable:

1. Access the Event Manager.
2. See "Configure Hit Attribute Login ID" in the *IBM Tealeaf cxImpact User Manual*.

Re-mapping components to reference session attributes

Similar to the Login ID configuration, some IBM Tealeaf CX components are pre-configured to reference the first four session attribute slots.

These pre-configured mappings are used to support customers upgrading from Release 7.2 or earlier. In earlier versions of IBM Tealeaf, these first four session attributes were known as UserDef 1 - UserDef 4.

Table 36 shows the default display names and the internal identifiers for these session attributes, which are assigned automatically during upgrade or new installation:

Table 36. Review Session Attributes 1 - 4

Default Display Name	Internal Name	Release 7.2 or earlier name
Session Attribute 1	CustomVar 1	UserDef 1
Session Attribute 1	CustomVar 2	UserDef 2
Session Attribute 3	CustomVar 3	UserDef 3

Table 36. Review Session Attributes 1 - 4 (continued)

Default Display Name	Internal Name	Release 7.2 or earlier name
Session Attribute 4	CustomVar 4	UserDef 4

For some IBM Tealeaf CX components, you can remap the session attributes to reference this data. Table 37 lists the components that may use these session attributes and where you may be able to configure them:

Table 37. Components and session attributes

IBM Tealeaf CX Component	Description of Use	Documentation on Changing
IBM TealeafcxImpact	<p>In IBM TealeafcxImpact, these attributes can be created and defined in the Tealeaf Event Manager to track session-based information captured by events. See "TEM Events Tab" in the <i>IBM Tealeaf Event Manager Manual</i>.</p> <p>Session attribute data can be captured to dimensions for use in reporting. See "Tealeaf Report Builder" in the <i>IBM Tealeaf Reporting Guide</i>.</p>	See "TEM Session Attributes Tab" in the <i>IBM Tealeaf Event Manager Manual</i> .

Configuring Tealeaf dimensions

You can configure predefined dimensions to track contextual information about the activities of visitors to your web application.

These predefined dimensions are the:

- URL
The URL dimension identifies the URL of the hit.
- host
The host dimension identifies the host of your web application.
- application
The application dimension identifies the application name.
- server
The server dimension identifies the name of the server hosting your application

When configured properly, values for these dimensions are captured in the Tealeaf pipeline and periodically logged to the database.

Using the Tealeaf Event Manager, you can map values for these dimensions to values that are useful for reporting purposes by creating whitelists of accepted values. For example, multiple URLs may be mapped to a single URL value: search.

To initialize the predefined dimensions:

1. When the installation is complete, logging of these dimensions has been enabled. As traffic is being captured and processed by Tealeaf, applicable values for these dimensions are being logged into the Tealeaf database.

See "TEM Dimensions Tab" in the *IBM Tealeaf Event Manager Manual*.

2. These dimensions are initially configured to report from a whitelist of values, which means that the dimension reports from these whitelist values only. Upon installation, the whitelist is empty.

Before you can begin using these dimensions, you must specify the whitelist of values, which can be gathered from the logs into which detected values are inserted. After installation, please wait one hour or a suitable period of time to gather a sufficient sample of values into the logs.

3. To make changes to these dimensions, open the Tealeaf Event Manager.

- a. Login to the Tealeaf Portal.

- b. From the Portal menu, select **Configure > Event Manager**

The Tealeaf Event Manager is displayed.

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

- c. Click the **Dimensions** tab.

- d. In the left panel, click the URL/Host/App/Server report group.

The dimensions are displayed for editing.

"TEM Dimensions Tab" in the *IBM Tealeaf Event Manager Manual*.

Tealeaf also provides the following dimensions:

- Connection Type

Depending on the computed speed of traffic between visitor and web server, this dimension buckets the visitor's session into one of four connection types:

- Dial Up
- ISDN
- DSL
- T1

See "Analyzing Performance" in the *IBM Tealeaf Reporting Guide*.

- Content Type

This dimension contains the value of the HTTP header for CONTENT_TYPE captured from the hit.

If the header value begins with /text or /application, the reported value is Page. Otherwise, the value is set to unknown.

No additional configuration is required.

- Request Cancelled

If the hit was cancelled by the visitor or the server, this value is set to true.

No additional configuration is required.

- Traffic Type

This dimension identifies the type of user agent that is initiating the session. Possible values include Browser, Bot, and Mobile, among others.

Note: Detection of user agents, including the type of traffic requirements deploying the Tealeaf Reference session agent in your Windows pipeline and enabling extended user agent parsing. See "Managing User Agents" in the *IBM Tealeaf cxImpact Administration Manual*.

After a sufficient period has elapsed to capture a good sample of data for each dimension into the logs, you may specify the whitelists for each of the dimensions. See "TEM Dimensions Tab" in the *IBM Tealeaf Event Manager Manual*.

Optionally, you can configure the dimension to use a whitelist and observed values. This option allows values that are detected in the capture stream to automatically be made available for reporting purposes, in addition to any whitelist values you specify.

Note: Given the volume of URLs in the typical web application, configuring the URL dimension to whitelist + observed values is not recommended. For more information on configuring the values to record, see "TEM Dimensions Tab" in the *IBM Tealeaf Event Manager Manual*.

Creating deviations

IBM Tealeaf CX provides the capability of creating and storing standard deviation calculations for any selected event or dimension. This capability was made available in version 8.0

Standard deviation calculations, also known as *movers*, require several days of stored data (depending on configuration) before they can be used in actionable reports.

After installing IBM Tealeaf CX, you might want to create some movers off of the events that already exist in your system.

For more information on creating deviations, see "TEM Top Movers Tab" in the *IBM Tealeaf Event Manager Manual*.

For more information on deviation reports, see "Analyzing Top Movers" in the *IBM Tealeaf Reporting Guide*.

Configuring sessionization settings

Configuring sessionization involves configuring session timeout settings, changing top-level domains, and configuring session close events.

Configuring session timeout settings

By default, Tealeaf is pre-configured to time out sessions that grow too large, last too long, or are left idle for a period of time.

When these limits are exceeded, the Canister breaks the session into fragments, each of which is assigned the same **TLTSID** value.

The TLSessioning session agent continues to use the same TLTSID value for sessions that have timed out.

In the Services Controls tab of the Canister configuration, you can specify Session Size Limits, which define the maximum number of hits, size in bytes, or time per session. For information about setting size limits, see "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

Tealeaf also times out sessions that have been allowed to stand idle for a period of time. If no new pages are added to a session for a predefined period of time, the session is closed. If the visitor resumes browsing, a new session identifier is issued. See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

Related concepts:

"Data sessionization in the Windows pipeline" on page 24

Changing top-level domains

If Tealeaf is monitoring multiple top-level domains, a visitor who browses across the domains may generate a new TLTSID from the Tealeaf Cookie Injector.

For example, switching from `www.site1.com` to `www.site2.com` may cause the issuance of a new unique Tealeaf identifier, since the browser does not include `www.site1.com` in the first request to `www.site2.com`.

Cross-domain sessions may be stitched together if there is a known value present in the session when it jumps from one domain to the next. Possible solutions:

1. You could add functionality to your Web applications to forward the TLTSID cookie value into the other domains to which the TLTSID cookie itself cannot be submitted by the browser. You could use JavaScript to set TLTSID cookies into the browser's cookie cache for the other domains.
2. If a query parameter is inserted into the URL that can be searched and evaluated when the visitor exits one domain and enters the other, the session fragments may be stitched together using the TLSessioning session agent. For example, you could submit the TLTSID value in the query string of requests to other domains and modify the Web applications to consume that TLTSID query string value and to set a TLTSID cookie with that value from the Web or application servers.

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

Configuring session close events

You can configure Tealeaf events to close a session if an event condition is met.

For example, suppose a visitor is using an application to manage multiple customer accounts in a single session. You can define a session close event to close the session whenever the Close Account screen is reached for each customer that is being processed by the visitor. In this case, it is more appropriate to evaluate the dataset as separate sessions, instead of as a single visitor-centric session.

Session close events cause multiple sessions to contain the same TLTSID cookie. Sessions closed by events are not fragmented sessions and cannot be merged back together.

In Advanced Mode, session close events are configured using the JavaScript `CloseSession` function with a `SessionCloseReason` session property. See "Advanced Mode for Events" in the *IBM Tealeaf Event Manager Manual*.

Chapter 6. Administering and managing IBM Tealeaf CX

Administering IBM Tealeaf CX involves preparing, monitoring, and modifying the environment on which IBM Tealeaf solutions run.

The topics contained in this chapter describe how to administer and manage an IBM Tealeaf runtime environment.

Generally, administering an installation environment involves performing well-defined tasks periodically and regularly to ensure that IBM Tealeaf CX is running efficiently and functioning properly.

The information in this chapter pertains to administering and managing IBM Tealeaf CX specifically. Information in this chapter does not supersede your site's administrative policies and procedures for administering databases, storage, servers, and IT infrastructure in general.

Configuring IBM Tealeaf cycle services

It is recommended to cycle the services on all Processing Servers (Canisters) and Report Servers once per day during off-peak hours.

Cycling services alleviates the following issues:

- Space and memory issues.

On the Processing Server, Canister services use a lot of memory. At the end of the day, memory that is used for the Short-Term Canister can be highly fragmented and therefore less efficient. Recycling the services results in a defragmentation of the Canister memory automatically and ensures consistent performance in the Short-Term Canister.

- Residual data and information.

Cycling services flushes out any residual information about existing sessions and prepares the system for the next day

- Health checks on the Long-term Canister.

On the Processing Server, cycling services process also runs scripts through `TLTMaint.exe` to verify the integrity of the Long-Term Canister.

By default, the IBM Tealeaf Scheduling Service is configured to run a cycle services job on each IBM Tealeaf server at 12:30 AM, local time.

All IBM Tealeaf servers can be cycled at the same time. The Processing Servers and Reporting Servers are the most important servers to cycle on a daily basis.

If you deployed a Health-Based Routing server in your system, configure cycle services on the Processing Servers so that the HBR always has a Processing Server available to which to send hits. Otherwise, data may be lost. For more information about HBR, see "Health-Based Routing (HBR) Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

For more information about enabling and scheduling cycle services, see "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.

Backing up and restoring IBM Tealeaf session data

Scheduling and implementing periodic backups protects you from data loss and keeps a record of important sessions for future reference.

IBM Tealeaf provides backup and restore facilities for maintaining backups of IBM Tealeaf session and index data that is stored on Tealeaf Canisters.

Although backing up and removing session data on a daily basis is the ideal scenario, implement a backup plan that fits your production environment maintenance and backup schedule.

Depending on the volume of data you are backing up, it might make better sense to perform a partial backup instead of a full backup. A partial backup saves all canister data files that have not been previously archived.

For information about backing up and restoring session data, see "TLBackup and TLRestore" in the *IBM Tealeaf cxImpact Administration Manual*.

Creating a shutdown script in Windows

You can create a shutdown script that runs CanSvcs.exe before the computer shuts down.

When you shut down the IBM Tealeaf CX server, it is important to ensure all IBM Tealeaf services were stopped before the server shuts down. If services are not stopped in time, the IBM Tealeaf Canister can become corrupted and might lose all stored data. You can create a shutdown script that keeps the operating system running until all services are stopped.

By running CanSvcs.exe all "in-process" data is saved, which prevents the Canister from becoming corrupted.

To create a shutdown script in Windows:

1. Open a Group Policy as a stand-alone Microsoft Management Console:
 - a. Open the Microsoft Management Console (MMC) by clicking **Start > Run...** > and typing **mmc** in the **Open** field.
 - b. In the Console menu, select **Add/Remove Snap-in**.
 - c. In the Add/Remove Snap-in dialog, click **Add** and select **Group Policy** from the **Add Standalone Snap-in** dialog.
 - d. Click **Add**.
 - e. In the Select Group Policy Object dialog, click **Browse** to find the Group Policy that you want.

If you want to save the Group Policy console and be able to choose which Group Policy object opens in it from the command line, select the **Allow focus of the Group Policy Snap-in to be changed when launching from the command line** check box.
 - f. Click **Finish**.
 - g. Click **Close** in the Add Standalone Snap-in dialog and **OK** in the Add/Remove Snap-in dialog.
2. Assign a computer
 - a. In the Console tree of Management Console, select the following:

Local Computer Policy > Computer Configuration > Windows Settings > Scripts (Startup/Shutdown).

- b. In the details pane, double-click **Shutdown**.
- c. In the Shutdown Properties dialog, click **Add**.

In the Add a Script dialog:

- 1) Click **Browse** beside the Script Name field and navigate to the following path:
 \Program Files\TeaLeaf\CanSvc.exe
 - 2) In the Script Parameters field, enter **-stop**.
 - 3) Click **OK** twice.
- d. Close the MMC window and save the console. By default, the file is saved here:
 C:\Documents and Settings*<current user>*\Start Menu\
 Programs\Administrative Tools\ Console1.msc

The registry entry **HKEY_LOCAL_MACHINE > SYSTEM > CurrentControlSet > Control > > >> WaitToKillServiceTimeout** determines how long the system waits for TeaLeaf services to stop after notifying the service that the system is shutting down. This value is automatically set to 600,000 milliseconds (10 minutes) by the IBM Tealeaf CX installation program.

Resetting IBM Tealeaf system service stop times

The length of time IBM Tealeaf waits for TeaLeaf services to stop after notifying the service that the system is shutting down is set automatically by the IBM Tealeaf CX installer program.

The default value is 600,000 milliseconds (10 minutes).

You can reset the time that IBM Tealeaf waits for TeaLeaf services to stop after notifying the service that the system is shutting down by editing the registry entry **HKEY_LOCAL_MACHINE > SYSTEM > CurrentControlSet > Control WaitToKillServiceTimeout**.

Updating user agent files

IBM Tealeaf uses publicly available standards for detecting user agents.

- BrowsCap.csv

The BrowsCap.csv file detects *fixed* user agents and contains definitions for various web-capable devices.

BrowsCap.csv is widely used in commercial applications.

Note: Tealeaf administrators can use UserCap.csv to supplement the contents of the BrowsCap.csv file with user agents that are not listed in the standard.

- WURFL.csv

The WURFL standard detects mobile user agents.

Store BrowsCap.csv and WURFL.csv the in the following location:

<Tealeaf_Install_Directory>\system

Install the files on any Canister server or Health-Based Routing server where the Tealeaf Reference session agent is installed.

To enable detection of fixed and mobile user agents, the Tealeaf Reference session agent must be installed in each Windows pipeline that processes session data. See "Tealeaf Reference Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Health-Based Routing enables load balancing between multiple Canisters in a Tealeaf environment. See "Health-Based Routing (HBR) Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Installing the most up-to-date version of the fixed user agent file (BrowsCap.csv)

IBM Tealeaf uses a publicly available standard to identify your visitor's browser capabilities.

As new browsers and web devices are released or updated, the version of BrowsCap.csv that is installed on your system might become out of date.

Check the Browser Capabilities Project website periodically for the most up-to-date release information and downloads of BrowsCap.csv.

To get the version of the BrowsCap.csv file with the latest definitions:

1. Download BrowsCap.csv to the IBM Tealeaf CX server.

You can find the file at: Browser Capabilities Project website.

Note: BrowsCap is available in various formats (.ini, .xml and .csv for example). IBM Tealeaf supports the comma separated version (.csv) format only.

On startup, IBM Tealeaf verifies that BrowsCap.csv was not modified and rejects any modified versions. See "Managing User Agents" in the *IBM Tealeaf cxImpact Administration Manual*.

2. On the IBM Tealeaf CX server machine, save the file to `<Tealeaf_Install_Directory>\system`.

3. Restart the Transport Service through the Tealeaf Management System.

For information about editing, copying, assigning, and reviewing the history of all configurations managed by Tealeaf Management System, see "TMS WorldView Tab" in the *IBM Tealeaf cxImpact Administration Manual*.

Installing the most up-to-date version of the mobile user agent file (WURFL.csv)

IBM Tealeaf uses a publicly available standard to identify your visitor's mobile device capabilities.

Wireless Universal Resource FiLe (WURFL) is a Device Description Repository (DDR), i.e. a framework that enables applications to map HTTP requests to a description of the capability of the mobile device that requests the page.

The general procedure for installing the most up-to-date version of the mobile user agent file (WURFL.csv) is:

1. If you do not have the most up-to-date version of WURFL.xml, download the latest version from <http://wurfl.sourceforge.net/>.
2. Convert WURFL.xml to WURFL.csv.
3. Install WURFL.csv to `<Tealeaf_Install_Directory>\system`
4. Restart the IBM Tealeaf transport service.

For detailed instructions about how to prepare WURFL.csv for use with IBM Tealeaf, see "Maintenance for CX Mobile" in the *IBM Tealeaf CX Mobile Administration Manual*.

Managing and installing IBM Tealeaf patches

IBM Tealeaf provides automated, tested and documented software upgrades to ensure you have the most up-to-date features and functions. However, if the product development team at IBM Tealeaf discovers a problem with a product or component before a scheduled upgrade, it might release a fix to the problem in the form of a software patch.

Software patches contain fixes to problems that can be applied without having to perform a product upgrade.

Information in the topics that follow explain how to acquire, manage, and install IBM Tealeaf software patches.

Instructions in the topics that follow apply to IBM Tealeaf Portal and its components.

Instructions in the topics that follow do not apply to the following IBM Tealeaf components:

- IBM Tealeaf CX Passive Capture Application
- IBM Tealeaf CX UI Capture components
- IBM Tealeaf CX RealiTea Viewer

Contact IBM Tealeaf Technical Support for information about applying patches to IBM Tealeaf CX Passive Capture Application, IBM Tealeaf CX UI Capture components, and IBM Tealeaf CX RealiTea Viewer.

How IBM Tealeaf patches work

In scenarios where an upgrade is not the appropriate way to address an issue with the product, the IBM Tealeaf development team can provide a software patch.

Generally, software patches are tested against the most recent version of IBM Tealeaf and might not have been tested against the build you are currently running. For this reason, it is important to manage the patch process, providing for testing, rollback, and tracking of software versions in your installation environment.

Customer advisories

Periodically, IBM Tealeaf publishes Customer Advisories for known issues that may affect multiple customer installations. Customer Advisories may contain recommendations, including upgrading or patching software.

Review the list of published Customer Advisories on a regular basis to make sure you are up-to-date on all known issues. Sign in to the IBM Client Success portal to access the IBM Tealeaf support portal. You can search the portal for advisories.

Managing IBM Tealeaf patches

If you have not done so, you should create a means of storing the original files that have been patched in a location outside of the IBM Tealeaf install directory.

This location should be part of any regular backups that are performed for the IBM Tealeaf system.

Use the following procedure as a guideline to manage IBM Tealeaf patches:

1. Download the software patch from the IBM Client Success Portal for IBM Tealeaf.
You need a user ID and password to access the IBM Client Success Portal for IBM Tealeaf.
2. Create a directory on a server that is not one of the IBM Tealeaf servers and name it Tealeaf Patches.
3. For each patch for Tealeaf:
 - a. Create a sub-directory with a date stamp and the issue that is being fixed.
You might include the Customer Support Case number if applicable. e.g.:
100421 SR00010531
 - b. Within the directory, create the following sub-directories:
 - doc
 - patch
 - backup
4. Save the instructions and any case history that is applicable into the doc directory.
5. Unzip the patch into the patch directory.
6. Complete the installation instructions.

You can now install the patch or patches.

Installing IBM Tealeaf patches

The procedure for installing IBM Tealeaf patches as documented here can be applied to most patches.

Before you begin, verify that you are logged into the server using an account that has *write* authorization to the directories to which the patch files are to be applied.

When copying the patch files to the IBM Tealeaf directory, never rename files in the directory unless explicitly told to do so.

If you have not already done so, set up your patch management directories.

In some instances customers are given instructions that are specific to a particular fix. In such cases, follow those instructions rather than the instructions documented here.

If you have questions about the procedure for applying a patch, contact IBM Tealeaf Technical Support.

To install an IBM Tealeaf patch:

1. Unzip the patch into the patch directory.
2. Depending on the type of patch, you might have to stop any Tealeaf service effected by the patch.

If you do not know whether you need to stop the Tealeaf service, contact IBM Tealeaf Technical Support.

3. Create a backup (copy) of the files that are being patched from the affected sever(s) and place the files in a backup directory.

Note: Do not remove the corresponding files from inside the IBM Tealeaf install directory.

4. Copy the files in the patch directory to the appropriate destination(s) inside the IBM Tealeaf install directory.

Keep the following things in mind when copying the patch files to the IBM Tealeaf Portal directory:

- Never leave copies of files in the directory or in any sub-directories. There should be only one copy of each file in the directory tree.
- When replacing a file in the Portal directory tree, make the backup and then directly replace the file being patched.

Note: In some cases, you may be adding a new file to the directory.

- If you are installing in Windows Server 2008, you might have to unblock files before the operating system permits them to run.

Right-click the file and select **Properties...** and deselect the **Block** option.

5. Restart any services that you stopped as part of the patch process.
6. If you are patching the IBM Tealeaf Portal, you must complete the following additional steps:
 - a. Restart the Tealeaf Data Service through the Windows Services Control Panel on the Portal Server.
 - b. Perform an IIS reset at the command line of the Portal Server.

You are now ready to verify that the patch was applied successfully.

Be sure to check the logs for any error messages, if one or more IBM Tealeaf components has been restarted. See "Portal Logs" in the *IBM Tealeaf cxImpact Administration Manual*.

Chapter 7. Troubleshooting

Use the information in this section to troubleshoot the IBM Tealeaf installation and configuration.

Using the IBM Tealeaf web application installation utility

Use the IBM Tealeaf CX Web Application Installation utility to manually re-install the IBM Tealeaf CX Portal and its web service with expanded supervision of the required permissions and component installation.

In most installations, the IBM Tealeaf Portal Web Application is installed and configured automatically by the standard IBM Tealeaf CX install process. Due to issues in the installation environment or permission restrictions, the Portal components may not install properly the first time.

Do not execute the IBM Tealeaf CX Web Application Installation utility unless you are experiencing issues with the IBM Tealeaf CX Portal.

Related tasks:

“Verifying system start up for single server topology” on page 96

How the IBM Tealeaf web application installation utility works

Use the information here to learn about the IBM Tealeaf web application installation utility

To launch the IBM Tealeaf CX Web Application Installation Utility, please select the following from the Windows Start menu: **Start > Programs > Tealeaf Technology > Tealeaf CX Portal > CX Web > Application Installation Utility**




Figure 7. Application Main Screen

The areas of the user interface are described and defined below.

Components

The components section lists the components that you can install, upgrade or remove using this application. Components include:

- Portal
The Tealeaf Portal Web Application.
- Web service
The Tealeaf Portal Web Services.
Web Service requires IBM Tealeaf cxConnect for Data Analysis, IBM Tealeaf cxReveal, or both.

Options

There are two options, Standard or Extended. For most systems, only the Standard (default) options are needed. For systems that have been altered or re-purposed, the Extended option might be required to properly install the IBM Tealeaf CX Portal.

- Standard
The standard set of installation options assumes most of the default operating system permissions are still intact:

3. Directory Permissions	
.NET Framework	False
Virtual Directory	False
Temp Chart	True
Environment Temp	False
Windows\Assembly	False
IIS Temporary Compressed Files	False
IIS Help	False
Windows\InetSrv	False
IIS Log Files	False
Windows\System32	False

Figure 8. Standard options

- Extended

For systems that have been re-purposed or locked down, system settings and permissions required for an IIS Web Application to run properly may be removed by system administrators for security reasons. The Extended install options perform the same actions as the Standard set and adjust additional directory permissions.

The extended set of installation options attempts to repair as many broken file and directory permissions as possible (as required by IIS and ASP.NET).

3. Directory Permissions	
.NET Framework	True
Virtual Directory	True
Temp Chart	True
Environment Temp	True
Windows\Assembly	True
IIS Temporary Compressed Files	True
IIS Help	True
Windows\InetSrv	True
IIS Log Files	True
Windows\System32	True

Figure 9. Standard options

Some Portal installation settings cannot be changed.

Virtual Directory

The following are the configurable virtual directory specific options:

Table 38. Virtual directory option definitions.

Option	Description
Action	The action to be performed: Install, Upgrade, or Remove.

Table 38. Virtual directory option definitions (continued).

Option	Description
Name	The virtual directory name. Fore example: http://<host>/Name.
Path	The physical path of the web application or web service being installed.
Authentication	The directory security settings for the website. <ul style="list-style-type: none"> • NT for NT/Active Directory authentication • Anonymous for Database/Portal authentication. For more information on authentication, see "Authentication" in the <i>IBM Tealeaf cxImpact Administration Manual</i> .
Allow HTTP Compression	Assuming HTTP compression is enabled by the website, this setting allows HTTP compression for this virtual directory.
IIS Website	The IIS website where the virtual directory is to be installed.
ASP.NET Framework	The ASP.NET Framework associated with this virtual directory. Choosing a framework that does not match the current OS architecture and the Enable32BitAppOnWin64 setting (using 32-bit ASP.NET Framework and Enable32BitAppOnWin64=0) forces the Enable32BitAppOnWin64 flag to match the architecture of the chosen ASP.NET Framework. This change, in turn, allows the newly created web application to function properly. However, it may break existing web applications, as this flag applies to the entire website.
Application Pool	The IIS Application Pool to which this virtual directory is assigned.
Verify Script Maps	Verifies that the ASP.NET extensions (.aspx, .ascx, .asmx, and so on) are configured properly.

IIS Web Server

The following are the configurable IIS Web Server specific options:

Table 39. Virtual directory option definitions.

Option	Description
Verify Extensions and Filters	Verify that all the Web Service Extensions and ISAPI Filters physically exist on disk.
Set RTV Mime Type	Apply the RTV (.tlx - application/rtv) Mime Type to the IIS Web Server.
Enable HTTP Compression	Enable HTTP Compression for the IIS Web Service. Any websites or virtual directories configured to allow HTTP Compression automatically acquire this new setting.
HTTP Compression Level	The Compression Level to use for HTTP Compression. Standard (5) is recommended. Maximum (9) requires more processing power for an insignificant gain over Standard (5) compression.

Table 39. Virtual directory option definitions (continued).

Option	Description
Update 404 Handlers	Point all IIS 404 handlers to the 404Handler.aspx page within this virtual directory. This functionality is used by Browser Based Replay (BBR) to catch and replay to requests to the local machine that would normally result in a 404.

Directory Permissions

The following are the various directory permissions that this utility can set or update:

Table 40. Virtual directory option definitions.

Directory	Description
.NET Framework	Update directory permissions to the .NET Framework directory of this virtual directory.
Virtual Directory	Update directory permissions for this virtual directory
Temp Chart	Update directory permissions for the temp_chart directory of this virtual directory.
Environment Temp	Update directory permissions for the Windows Temp directory, as defined by the %TEMP% environment variable.
Windows\Assembly	Update directory permissions for the Windows\Assembly directory.
IIS Temporary Compressed Files	Update directory permissions for the IIS Temporary Compressed Files directory.
IIS Help	Update directory permissions for the Windows\Help\iisHelp\Common directory.
Windows\InetSrv	Update directory permissions for the Windows\InetSrv directory.
IIS Log Files	Update directory permissions for the Log Files directory of the website of this virtual directory.
Windows\System32	Update directory permissions for the Windows\System32 directory.

Registry Settings

Registry Settings The following are the various registry settings required for the web application to function properly:

Table 41. Virtual directory option definitions.

Directory	Description
FIPS Algorithm Policy	Verify that the Use FIPS compliant algorithms local policy is disabled. The Tealeaf .NET libraries cannot function properly if this policy is enabled.

Installing or re-installing IBM Tealeaf CX Portal components

You can re-install IBM Tealeaf CX Portal components.

1. Select the component to install: **Portal** or **Web Service**.
2. Click the installation mode: **Standard** or **Extended**.
3. If you wish to see the log output of the install process, select **Show Output**.
4. Modify the available options.
To reset options, click **Reset**.
5. Click **Execute** to install the component.

If you attempted to re-install the IBM Tealeaf CX Portal under Extended mode and are still experiencing issues with connecting to the Portal, contact IBM Tealeaf Customer Support.

Using the TLTMaint script to maintain Canister stability

TLTMaint checks the stability of the Canister and its files. When a corruption is detected in any of the Canister data files, an attempt is made to correct the condition by rebuilding the corrupted table.

If stability of the Canister cannot be maintained, TLTMaint does not allow the Canister services to start.

TLTMaint is normally started by the TeaLeaf Canister Server program but can also be run manually at the command line.

Note: TLTMaint rarely requires extra configuration after installation. Unless explicitly told to do so by IBM Tealeaf personnel services, you should not have to use or configure TLTMaint. Do not attempt to perform this procedure without first consulting with IBM Tealeaf technical support.

To run the TLTMaint script manually, use the following command-line options:

- -h Help
- -v Verbose output
- -CheckOnly Performs check on tables without fixing anything.
- -Noserver Recovers files without datastore (Canister/Archive) services running.

TLTMaint publishes its processing to the Windows application log. It also stores log files in the TeaLeaf\Logs directory.

On services startup, TLTMaint performs the following checks:

- Locates the installation path of TeaLeaf by checking the Registry.
- Checks for the ctsrvr.cfg file.
- Loads its own CFG file (TLTMaint.cfg).
- Creates Canister spacer file.

TLTMaint calculates spacer file size by reading the Max Ctree Bytes size from the Registry. Every time that it starts up, it can modify the spacer file size that is based on these settings.

- Checks Canister size and files to index.
- Checks database version.
- Begins checking tables that are defined in the TLTMaint.cfg.

If you want to skip the check of these tables, you can comment these out in `TLTMaint.cfg`. On bad Canister shutdowns, data files might be corrupted. An attempt is made to salvage the data file

- Runs CanTrim.
- Checks the CanTrim setting to see whether it is the right time and day to start.

CanTrim is enabled through TMS. See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

Troubleshooting privacy performance in the PCA server

Sometimes privacy actions can prevent the PCA from processing hits in real-time.

You can address performance-related issues in the PCA that are caused by privacy actions by doing the following:

1. Move all non-sensitive privacy rules into the Windows pipeline downstream.
See "Windows Pipeline Privacy" in this guide for more information.
2. Redesign your privacy rules to eliminate regular expressions.
3. Experiment with the Stop Processing flag if only one occurrence in a page requires blocking.
4. Create Tests for your CPU-intensive rules based on URL or other unique data identifier in the request so only relevant hits are processed by the privacy rule.

Related concepts:

"Windows pipeline privacy" on page 20

Chapter 8. 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 42. Getting help

To view...	Do this...
Product documentation	On the IBM Tealeaf portal, go to ? > Product Documentation .
IBM Tealeaf Knowledge Center	On the IBM Tealeaf portal, go to ? > Product Documentation and select <i>IBM Tealeaf Customer Experience in the ExperienceOne Knowledge Center</i> .
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 <i>IBM Tealeaf CX PCA Manual</i> .

Available documents for IBM Tealeaf products

The following table is a list of available documents for all IBM Tealeaf products:

Table 43. Available documentation for IBM Tealeaf products.

IBM Tealeaf products	Available documents
IBM Tealeaf CX	<ul style="list-style-type: none">• <i>IBM Tealeaf Customer Experience Overview Guide</i>• <i>IBM Tealeaf CX Client Framework Data Integration Guide</i>• <i>IBM Tealeaf CX Configuration Manual</i>• <i>IBM Tealeaf CX Cookie Injector Manual</i>• <i>IBM Tealeaf CX Databases Guide</i>• <i>IBM Tealeaf CX Event Manager Manual</i>• <i>IBM Tealeaf CX Glossary</i>• <i>IBM Tealeaf CX Installation Manual</i>• <i>IBM Tealeaf CX PCA Manual</i>• <i>IBM Tealeaf CX PCA Release Notes</i>

Table 43. Available documentation for IBM Tealeaf products (continued).

IBM Tealeaf products	Available documents
IBM Tealeaf CX	<ul style="list-style-type: none"> • <i>IBM Tealeaf CX RealTime Viewer Client Side Capture Manual</i> • <i>IBM Tealeaf CX RealTime Viewer User Manual</i> • <i>IBM Tealeaf CX Release Notes</i> • <i>IBM Tealeaf CX Release Upgrade Manual</i> • <i>IBM Tealeaf CX Support Troubleshooting FAQ</i> • <i>IBM Tealeaf CX Troubleshooting Guide</i> • <i>IBM Tealeaf CX UI Capture j2 Guide</i> • <i>IBM Tealeaf CX UI Capture j2 Release Notes</i>
IBM Tealeaf cxImpact	<ul style="list-style-type: none"> • <i>IBM Tealeaf cxImpact Administration Manual</i> • <i>IBM Tealeaf cxImpact User Manual</i> • <i>IBM Tealeaf cxImpact Reporting Guide</i>
IBM Tealeaf cxConnect	<ul style="list-style-type: none"> • <i>IBM Tealeaf cxConnect for Data Analysis Administration Manual</i> • <i>IBM Tealeaf cxConnect for Voice of Customer Administration Manual</i> • <i>IBM Tealeaf cxConnect for Web Analytics Administration Manual</i>
IBM Tealeaf cxOverstat	<i>IBM Tealeaf cxOverstat User Manual</i>
IBM Tealeaf cxReveal	<ul style="list-style-type: none"> • <i>IBM Tealeaf cxReveal Administration Manual</i> • <i>IBM Tealeaf cxReveal API Guide</i> • <i>IBM Tealeaf cxReveal User Manual</i>
IBM Tealeaf cxVerify	<ul style="list-style-type: none"> • <i>IBM Tealeaf cxVerify Installation Guide</i> • <i>IBM Tealeaf cxVerify User's Guide</i>
IBM Tealeaf cxView	<i>IBM Tealeaf cxView User's Guide</i>
IBM Tealeaf CX Mobile	<ul style="list-style-type: none"> • <i>IBM Tealeaf CX Mobile Android Logging Framework Guide</i> • <i>IBM Tealeaf Android Logging Framework Release Notes</i> • <i>IBM Tealeaf CX Mobile Administration Manual</i> • <i>IBM Tealeaf CX Mobile User Manual</i> • <i>IBM Tealeaf CX Mobile iOS Logging Framework Guide</i> • <i>IBM Tealeaf iOS Logging Framework Release Notes</i>

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](http://www.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