



Collabora  
Productivity

# Collabora Online Technical Documentation

Version 2024-10-16



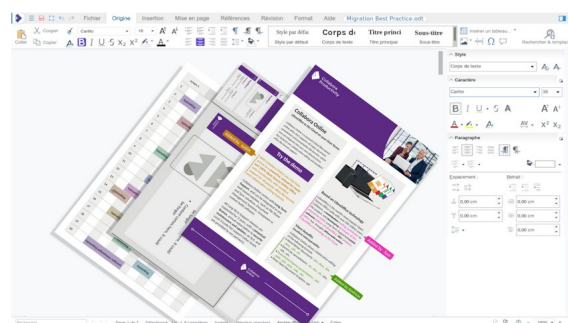
Collabora Online is a software component which can be integrated into your web application and enables viewing and editing of office documents in dozens of file formats. Collabora Online provides the un-paralleled file format support and rendering capability of LibreOffice.

## Key Features

- Collaboratively edit text documents, spreadsheets, presentations & more
- Preservation of layout and true WYSIWYG formatting of documents
- Excellent interoperability with Microsoft file formats
- VBA-macro's partially supported
- Basic editing features
  - Download as PDF, ODF or MS Office
  - EPUB export, full CSV support
  - Undo/redo & document repair
  - Rich character & paragraph formatting
  - Insert charts, tables, images
- Advanced editing features
  - Change tracking & manage changes
  - Live notifications of users entering or exiting
  - Insert, edit and reply to comments
  - Rich form fields, with protected sections
  - Full page styling, columns etc.
  - Powerful and inter-operable formulae calculation engine
  - Data entry validation
  - Pivot table editing
  - And much more
- Powerful scripting interface
  - Over 1500 interfaces with 4000 attributes and methods
  - Rich manipulation of document content and structure
- Admin console and API for monitoring system utilization.
- Extended subset of WOPI (Web Application Open Platform Interface) protocol for communication with your web application (authentication and file transfer)
  - Consultancy available to integrate with any solution.

## Deployment options

- RPM or DEB packages for all modern Linux distributions (minimal requirement: glibc 2.17)
- Docker image
- Other (source code is available)



## Enterprise-ready

- QA'd and tested, LTS, Security updates, SLA
- Roadmap – Product management interaction/input
- Apps for mobile devices, and for desktop.

## Scalability

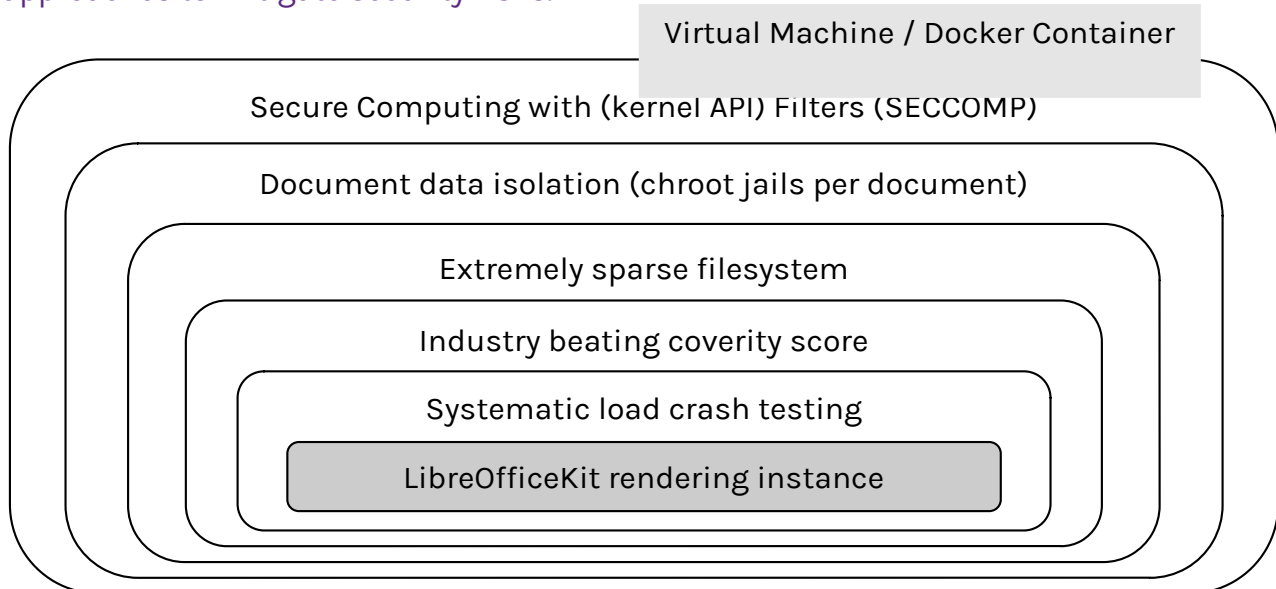
- Round numbers to provide an upper bound for provisioning – combined with diversity factors for server sizing.
- 15 active users / CPU thread
- 10 active users / Mbit/s
- 50 MB RAM / active user

## Browser support

Collabora Online uses a WebSocket protocol. Recent browsers versions, supported by their vendors, of Chrome, Safari, Firefox, Edge, Brave, iOS and Android are supported.

# Security

A significant focus of our design is to ensure good isolation of document data between users. In order to achieve this we adopt a layered approach to security with multiple approaches to mitigate security risks.



The main source of risk is the eight million lines of LibreOffice code we re-use from our existing Collabora Office product. Naturally these share many of the mitigation strategies we have:

## Virtual Machine / Docker Container

Separation and encapsulation of applications on the same host can provide an outer layer of containment.

## SECCOMP (BPF)

Linux's Secure Computing with Filters are utilized with a Berkley Packet Filter (BPF) disabling many Kernel APIs which have no legitimate use by LibreOfficeKit, and have previously caused local privilege escalation issues, in addition debugging APIs like `ptrace`, and server calls like `listen` are disabled.

## Document Data Isolation

Each document is isolated in its own `chroot` jail running its own instance of a LibreOfficeKit process, and runs as a non-privileged 'cool' user.

These documents are located inside a path including a 16 character hard-random name.

## Extremely sparse file system

These `chroot` jails contain only the bare minimum of files (libraries, fonts, etc.) needed for running Collabora Office (LibreOfficeKit).

## Systematic Load Crash Testing

The primary entry point for any attacker is maliciously constructed files designed to attack our file filters. To avoid issues with our filters, in addition to reactive security issue fixing, combined with timely updates. LibreOffice conducts extensive automated crash testing, loading all previous problem documents from relevant public bug trackers - loading over 800,000 files regularly to ensure that there is no unexpected behavior. This is often combined with ASAN address checked builds.

## Industry Beating Coverity Score

LibreOffice uses routine, weekly static checking via Coverity, during development and has an exceptionally low defect score.

# Configuration

Collabora Online has to be configured before use. Many options have sensible defaults. There are command line options of `coolwsd`. `coolwsd --help` lists them all. The configuration file is `/etc/coolwsd/coolwsd.xml`. Besides these, there are environment variables which have effect on behavior. In this document we mention only the most important options.

## SSL settings

Collabora Online uses WOPI protocol, which mandates SSL. However, it is possible to run Collabora Online server without SSL, it is configurable. Basically there are 3 modes:

1. SSL
2. SSL termination
3. No SSL

When SSL is enabled, in `/etc/coolwsd/coolwsd.xml` the path to SSL key, SSL certificate and SSL CA certificate has to be given in the `ssl` block.

The SSL termination option in the config file enables integration of Collabora Online with SSL termination proxies, which handle incoming SSL connections, decrypt the SSL and pass on the unencrypted request to the server. In this setup only the proxy server has to have proper SSL settings, Collabora Online server is hidden behind it, and Collabora Online communicates unencrypted with the proxy.

If you set both enable and termination settings to false in `/etc/coolwsd/coolwsd.xml`, then Collabora Online can be used in a HTTP-only

environment, without encryption between browser and server. It is not recommended to use Collabora Online in this mode, but for testing only it is OK.

## Trusted WOPI hosts

Collabora Online accepts connection requests only from trusted WOPI hosts. The administrator has to list the host names and/or IP addresses of these trusted WOPI hosts in the `storage.wopi` block. Please note that connection requests from the same machine are always accepted.

## Logging

The level of logging can be set in `/etc/coolwsd/coolwsd.xml`. Set the log level and verbosity to one of: none (turns off logging), fatal, critical, error, warning, notice, information, debug, trace. When `coolwsd` is started from the console, it logs to the standard error (screen). When `coolwsd` is started by `systemd`, it logs into `systemd` journal. It is possible to redirect log to a file by enabling this feature in `/etc/coolwsd/coolwsd.xml`.

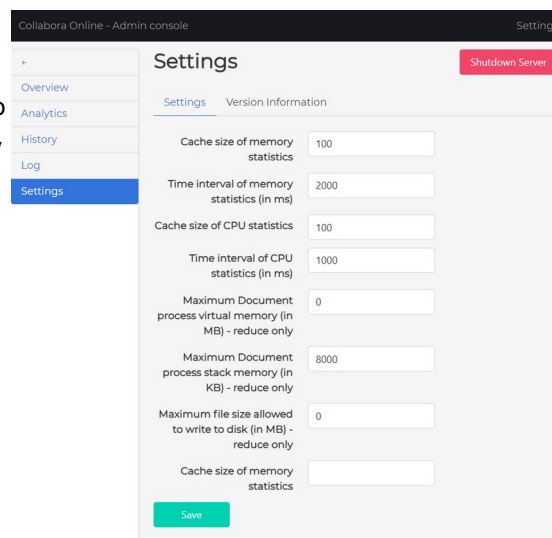
NB. high levels of logging have a significant performance impact and are extraordinarily verbose, and should be avoided in production.

# Admin Console

Do live monitoring of all the user sessions running on Collabora Online instance by accessing The Admin Console: <https://hostname:port/browser/dist/admin/admin.html>

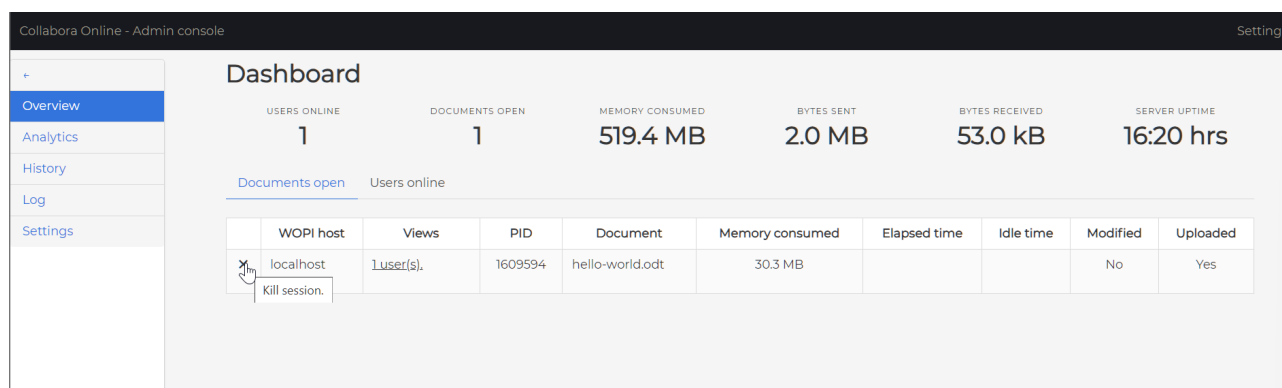
Port is 9980 by default. It will ask for username and password which can be set with the coolconfig utility. After entering the correct password you should be able to monitor the live documents opened, total users, memory consumption, document URLs with number of users viewing that document etc. You can also kill the documents directly from the panel which results in closing the socket connection to the respective document.

The admin-console front-end presents and fetches its data via a defined websocket protocol, which can be used to collect information programmatically to allow integration with other monitoring and control solutions. To subscribe and receive client notifications, query the open documents and change server settings.



For the websocket protocol details of Admin Console, see the Admin Console section in the protocol documentation:

→ <https://github.com/CollaboraOnline/online/blob/master/wsd/protocol.txt>



WOPI host	Views	PID	Document	Memory consumed	Elapsed time	Idle time	Modified	Uploaded
localhost	1 user(s)	1609594	hello-world.odt	30.3 MB			No	Yes

Admin Console Dashboard. It is possible to terminate a session by clicking on the × icon of the document

# Deployment options

## RPM and DEB packages for 64-bit Linux

Customers are provided with customer specific secret URLs to package repositories. Binary packages of Collabora Online can be installed on wide variety of Linux distributions, including openSUSE/SLE15, RHEL 7, 8, and 9 and compatible distributions, Debian 10, 11, 12, Ubuntu 18.04 LTS and higher.

To get full installation repositories, and commands to install binaries – please see the Partner Portal at <https://support.collaboraoffice.com/> which will provide the set of commands needed to configure repositories and to install Collabora Online on your selected Operating System.

## Docker image

Collabora Productivity provide scripts and Dockerfiles to create a Collabora Online Docker image. Docker images can be created on demand from the latest version of Collabora Online and the underlying system components. The repository is on GitHub.

<https://github.com/CollaboraOnline/online/tree/master/docker>

Customers can get access to Collabora's private docker hub upon request, and pull Collabora Online docker image.

# Hardware Provisioning

Calculating the hardware requirements for a deployment of Collabora Online requires several factors:

## Usage factor

This is what proportion of users are expected to be using Collabora Online at any given moment vs. other tasks such as E-mail and/or web browsing. A conservative value here for enterprise users would be a 25% usage ratio. For a large scale hosted solution focused on E-mail, or file-synchronization the usage ratio can be as low as 1 to 2%.

## Viewers vs. Editors

This gives an indication of the work-load likelihood. Again a conservative enterprise user value would be 50% viewing, 50% editing, whereas for a hosted, E-mail focused solution the viewing proportion could easily be over 90% viewing, 10% editing.

## Concurrent Users

To simplify this calculation we convert each Viewer into one tenth of an Editor, and add the Editors.

We recommend provisioning for Collabora Office

- **15 Concurrent Users per CPU thread.**
- **10 Concurrent Users per 1Mbit/s**
- **1 GB of RAM per server plus**
- **50MB of RAM per concurrent user**

These ratios are rather conservative, round numbers for easy mental arithmetic and a rough rule of thumb.

## Disk Storage

In a normal logging mode, Collabora Online produces reasonably light logging output, although is capable of producing very significant amounts when being debugged – however high volumes of logging also have a significant performance impact.

For scratch temporary storage in `/tmp` around four times the required RAM size is recommended for each machine.

Disk storage for jails in `/opt/cool/child-roots` can vary significantly. With a properly configured system with the privilege required to bind-mount, or that can hard-link document jails, under 100GB of fast local storage per system is required. Otherwise if the system setup or systemd policy forbids mounts or hard links from `/opt/collaboraoffice` then system images will need to be copied at considerable CPU and disk space – around 1GB per concurrent document. This is to be avoided. A normal, well configured system should function with just 500GB of storage. Most use-cases are not effectively limited by disk storage.

## Worked examples

### A. On-premise Enterprise

For 10,000 users, with a usage factor of 25%, and a View/Edit ratio of 50% we have:

- **Concurrent Users =**  
 $10,000 * 25\% * (50\% + 50\%/10)$   
**= 1,375 Concurrent Users**
- **CPU threads:** 138
- **Bandwidth:** 138 MBit
- **GB of RAM** 69 + server count

### B. E-mail / File Sync. focused Hoster

For 1,000,000 users, with a usage factor of 2%, and a View/Edit ratio of 90% we have:

- **Concurrent Users =**  
 $1,000,000 * 2\% * (10\% + 90\%/10)$   
**= 3,800 Concurrent Users**
- **CPU threads:** 380
- **Bandwidth:** 380 Mbit
- **GB of RAM:** 190 + server count

# Load Balancing & High Availability

There are many possible ways to load balance Collabora Online. However there are a set of basic principles that help correctly configure and use load balancing / HA solutions.

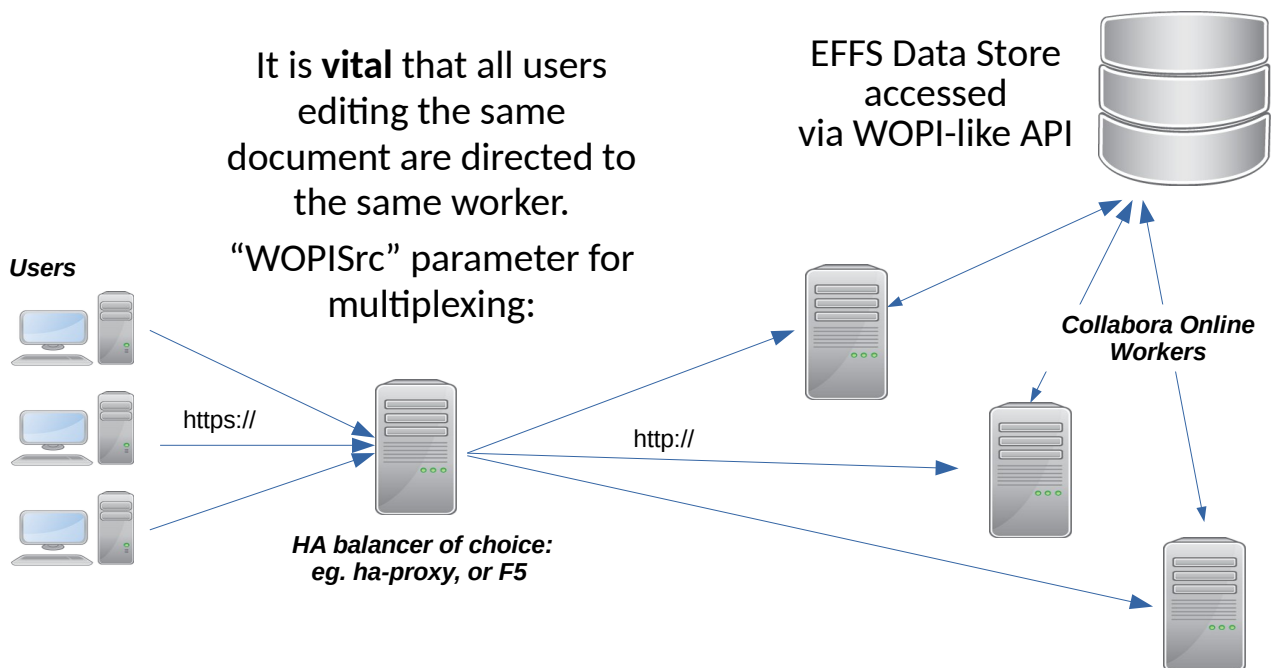
The key principle here is that Collabora Online that no long term persistent data is stored inside the Application – all long term persistent document data is saved to the data layer accessed via a protocol such as WOPI or WebDAV – but it does contain short term state. A Collabora Online instance has transient cached tile data but also, critically, the current document model's editing state.

In order for Collaborative Editing to function correctly, it is vital to ensure that all users editing the same document end up being served by the same Collabora Office instance. All load balancing approaches must keep this in mind.

Using the WOPI protocol, the https URL includes a WOPIsrc parameter that uniquely identifies this document. Thus load balancing can be done based on this URL parameter – ensuring that all identical WOPIsrc values are sent to the same Collabora Office instance. There are few new connections made for a given client, since this socket is upgraded to be a WebSocket which is subsequently used for low-latency bi-directional communication.

Further specifics on configuring HA proxy are provided in our Installation Guide, and more details on more advanced load balancing via eg. F5 are available on request.

For advanced load balancing in Kubernetes environment Collabora recommends the [Collabora Online Controller](#) product.





## General Disclaimer

Collabora makes no representations or warranties with respect to the contents of this document, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. The development, release, and timing of features or functionality described for Collabora products remains at the sole discretion of Collabora. Further, Collabora reserves the right to revise this document and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.