

Next generation Linux Applications Gateway for CERN accelerator control systems

Mario Rey Regúlez^{1*}, and Thomas Oulevey²

¹CERN IT Department, CD Group, 1211 Geneva 23, Switzerland

²CERN BE Department, CSS Group, 1211 Geneva 23, Switzerland

Abstract. CERN has been providing central Windows remote desktops via the Windows Terminal Infrastructure service for several years and aims to provide a similar experience for Linux graphical environments. Different communities and experiments offer a series of tools to their users with this goal in mind, but the solutions are far from ideal and generate a support overhead for their respective providers. The Linux Applications Gateway project (LAG) was born to provide this functionality centrally from the IT department. After an extensive market research, the tool FastX was identified as an enabler, and to set up a closed, internal pilot for evaluation. These efforts led to the creation of the Remote Operations Gateway (ROG) service with a high approval rate. We aim to further extend the usage of FastX at CERN, reaching out to other communities and experiments, and to provide a better support coverage for them all.

1 Introduction

CERN's remote desktop solution "Windows Terminal Infrastructure" provides a central solution for 34,000 users, enabling them to connect world-wide to 10,000 CERN remote desktops and servers inside the Organization's enterprise network [1].

Microsoft's Remote Desktop Services has been widely used at CERN for several years, and was proved to be highly stable and reliable, providing both generic and dedicated functionality through more than 40 Terminal Services clusters across six generations of Windows operating systems.

1.1 Technical context

Part of the Windows Terminal Infrastructure service offering includes several clusters for CERN's Beams department, dedicated to specific technical and scientific communities and their daily operations, such as Cryogenics, Vacuum or CO₂ control. They are critical for the monitoring and maintenance of the Large Hadron Collider.

However, not all their operations are possible through Windows operating systems products, and they also require Linux environments. Several communities and experiments at CERN offer different solutions to their users so they can seamlessly connect to remote

* Corresponding author: mario.rey@cern.ch

Linux graphical user interfaces (GUI). These solutions have limitations and generate a significant support overhead for each team and difficulties for users who are less used to highly technical computing environments. Additionally, the different implementations and tools they use aren't as safe as they should be and could put the Organization at risk.

The Linux Applications Gateway project (LAG) was born under the MALT [2] umbrella at CERN within the Information Technology department, trying to find a replacement for Windows Terminal Infrastructure clusters, on top of improving the service catalogue and solving the Accelerators and Technology Sector's long-awaited needs.

1.2 Objectives

To enlarge the service catalogue the IT department can offer to its users, we undertook to accomplish the following objectives:

- Provide a public cluster to all CERN users with a curated set of common applications of general utility (web browsers, email client, office suite...), protected behind CERN Single Sign-On and Multi-Factor Authentication.
- Offer private, dedicated clusters with specialized applications to those communities and experiments that would need them, through delegated administrators.
- Define the procedure to install and configure user-managed installations and provide them with licenses if required.
- Keep up with new-software releases and security patches centrally, properly validating them before being released to the CERN public.
- Support end-users and delegated administrators with their requests and solve their incidents.
- Collaborate with open-source projects and software vendors to follow up bug reports and propose required functionality, or actively enhance the software with recommendation and patches.

2 Market research

There is no native Linux tool or way to offer a service such as the one we were aiming to provide, like Windows does. Some methods and technologies are available, such as implementations of the Remote Desktop Protocol (RDP) or Virtual Network Connection (VNC) tools, but those are sub-optimal in a centralised, multi-user environment.

2.1 Technical requirements

To provide a successful and secure service it was required that all or as many as possible of the following features were available.

- Session persistence: ensures that the state and context of a user's interaction with the applications are preserved within a specific period of time.
- Session brokering: intermediation between user and session servers, to manage and distribute remote desktop connections across the deployment, guaranteeing easy session reconnection.
- Load balancing: distribute incoming user connections across session servers to optimize resource utilization and prevent server overload.
- Encryption: understandable, plain data is converted into a secure and unreadable format to safeguard sensitive information.

- OpenID Connect (OIDC): authentication protocol built on top of OAuth 2.0 designed to enable secure user identity verification, facilitating the integration with CERN Single Sign-On and multi-factor authentication.
- Window resizing: dynamically change and adapt the dimensions of the working user interface to support different resolutions and preferences.
- Multi-platform: ability to run on multiple operating systems without changing much the user experience and system interaction.
- Support: assistance, troubleshooting and bug fixing directly from the vendor and/or an active user community behind the software.

2.2 Product comparison

With the technical requirements clarified, in early 2019 we started searching for both open-source and commercial products that could meet our needs, looking both into industry-leaders and new products. The selection of the most promising were individually tested and assessed in lab environments, in collaboration with key stakeholders in Beams department. The results are shown in Table 1.

Table 1. Linux Applications Gateway feature comparison.
 Legend: ✓ supported; — partially supported; ✗ not supported.

	TigerVNC [3]	X2Go [4]	XRDP [5]	NoMachine [6]	FastX [7]
Persistence	✗	✓	✓	✓	✓
Brokering	✗	—	—	✓	✓
Load balancing	✗	—	—	✓	✓
Encryption	—	✓	✓	✓	✓
OpenID Connect	✗	✗	✗	✗	✓
Window resizing	✓	✓	—	—	✓
Multi-platform	✓	✗	✓	—	✓
Support	✗	✗	✓	✓	—

2.3 Product selection

FastX, from vendor Starnet, was identified as the most comprehensive tool as it covered all the requirements or had the potential and willingness of its developers to act on specific technical requests from CERN. Additionally, CERN already had a very successful working

relationship with Starnet, as they are also the providers of X-Win32 tool and licenses were already included in our current contract with them.

3 Service architecture

A typical FastX cluster deployment consists in the following components:

- Client: application on the user’s device that is used as entry point to the remote desktop services. It can be a web browser or a Windows, Linux or MacOS application, and communicates via HTTP/S.
- FastX gateway (optional): the access point to the rest of the system, used to physically separate sensitive information from the outside network and achieve network isolation.
- FastX manager: core component of the deployment, used as the central hub for all interactions and operations. It’s an Xorg X server that implement the FastX protocol.
- Database: used to store configuration and other information. It’s a MongoDB nedb database that can operate in both standalone and cluster mode.
- License server: to distribute and keep track of issued user licenses. The only supported system is a Reprise License Server (RLM).
- VirtualGL (optional): used for direct access to video card hardware when available on the session hosts.
- FastX hosts: Linux servers where the graphical X11 applications run and are finally displayed to the clients.

All components can be made redundant to achieve high availability and business continuity. A typical, simplified FastX cluster deployment is shown in Figure 1.

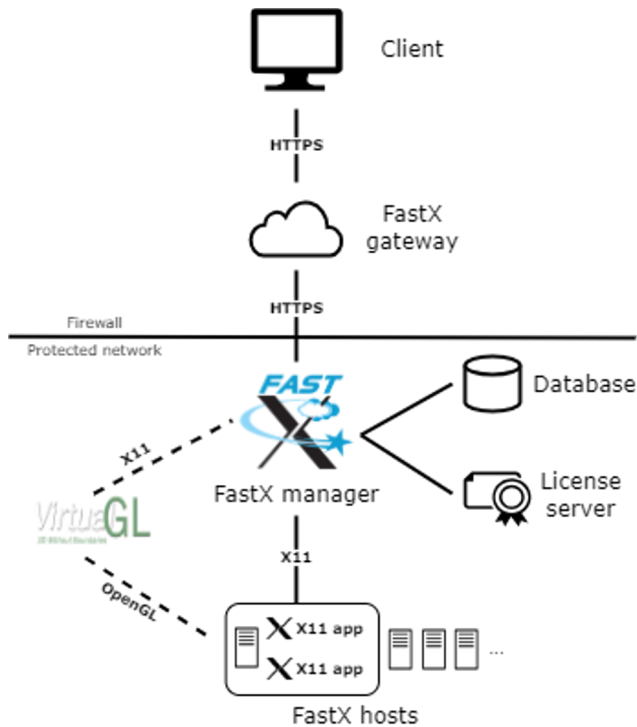


Fig.1. Diagram of a typical FastX cluster deployment.

For a smaller, simplified FastX deployment where no clusterization is needed, service administrators can install all components in a single server, where the gateway component is unnecessary as there is no isolation. This is what is called a standalone deployment.

4 Service implementation

4.1 Internal IT tests and pilot

Once we identified FastX as the tool to move forward, during the summer of 2019 we started the provisioning and installation of the machines. We decided to use CERN Cloud Infrastructure service, based on OpenStack, and Puppet as a configuration management tool, also centrally available and supported at CERN. The operating system would be CERN CentOS 7, as the time of the tests it was the recommended one and fully supported by the IT department.

The goal was to run an internal pilot service to test the performance, capabilities, potential and user satisfaction with the product. A baseline of common applications would be installed for general usage: Mozilla Firefox as web browser, Mozilla Thunderbird as e-mail client, LibreOffice as office productivity suite. Also, we could offer different desktop environments to satisfy a larger set of users: GNOME, Xfce and KDE were selected as the most versatile. From that starting point it would be possible to build a more stable service, aiming for the general CERN population or dedicated communities that could profit from it.

Coincidentally, FastX version 3 was in a beta phase and all support efforts were being focused on it, so we decided to skip version 2 and start profiting from new features, plus reporting back to Starnet bugs or feature requests from that version on.

During installation and configuration, several issues were found, mainly coming from the way Starnet packaged and installed its software: rather than providing a repository we could install and use, they heavily relied on REMI repository [8] which caused dependencies issues with our systems and internal repositories. We had to adapt their installation scripts to match our available packages.

They also packaged the application with their own internal builds of other components like NodeJS or Redis, sometimes relying on old releases whose support was abandoned. Updating to a supported version was sometimes a challenge for both teams.

Being a beta release, FastX 3 presented numerous bugs that we helped Starnet developers understand and fix: internal components not properly configured on installation, services crashing, unimplemented API calls, security concerns, etc. There were also fixes and requests to better adapt the tool to the CERN environment, like better handling of home directories on our AFS system, or a proper implementation of the OpenID Connect authentication protocol to match industry standards.

One of the major pain points was the performance: while on standalone installations FastX worked fine, a degradation was clearly noticeable on cluster deployments. This led to the release of FastX 3.2, where major internal components were redesigned: it changed from a distributed deployment where all nodes had the cluster manager and CouchDB component installed and continuously communicating, to a central deployment on which a main node had the manager role and MongoDB database. This drastically improved the performance overall and could finally be considered for a pilot deployment.

Under this context we launched a closed, internal pilot that we called CERNLAG, for CERN Linux Applications Gateway, that was presented and offered to several technical communities in CERN, especially in the Beams department, so they could test its potential. In Figure 2 we can see a screenshot of a working LAG session, running a Xfce desktop from

a web browser, with a LibreOffice document and a regular terminal showing the user home directory on a network EOS path.

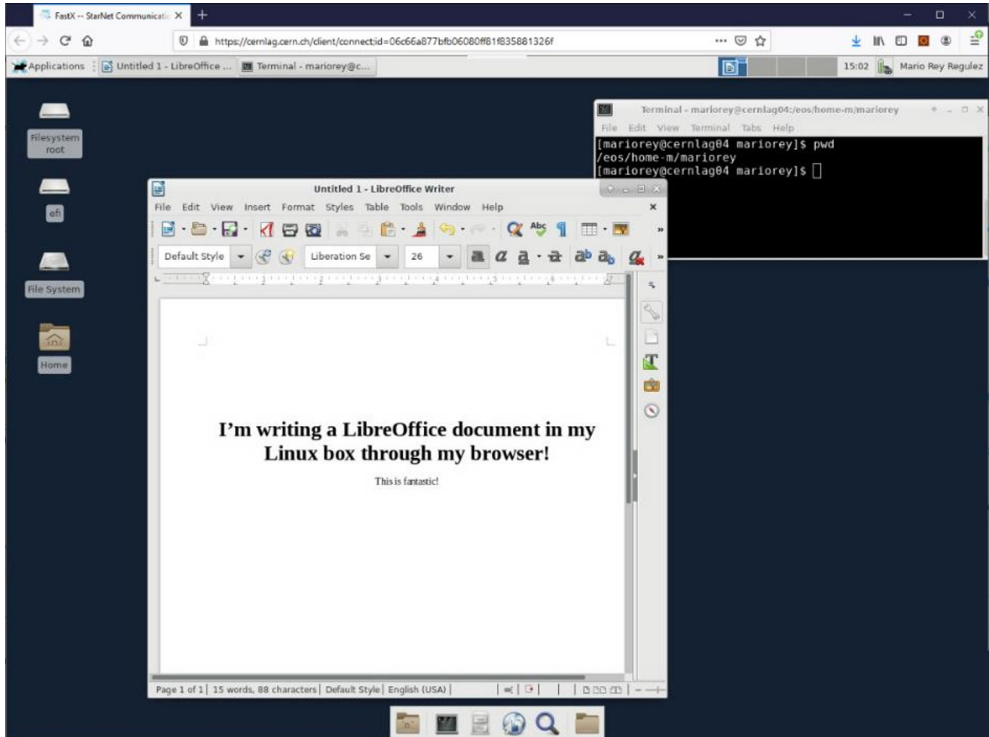


Fig. 2. Screenshot of a working FastX session from a web browser.

It was a fruitful collaboration appreciated by both Starnet and CERN: they were satisfied with us testing and improving their product, and we were happy to see the product evolve to better match the needs of the HEP user communities.

4.2 LAG as a Service

After the initial pilot was released in October 2019, we offered the service to user communities that would like to test their setups for their own use-cases, without the IT department intervention except for general installation support and license management.

With this purpose in mind, we prepared a thorough documentation for administrators, giving them instructions on how to install and configure both standalone deployments and cluster installations. Complementing that, we released a Discourse forum for community engagement, important announcements, and new version releases, additionally to a mailing list.

That initiative led to the creation of a growing community and the proliferation of small FastX installations, both for individual use and for limited use-cases where this technical solution was a great fit. The interest was more prominent in the Accelerators and Technology Sector (ATS), where they saw FastX and LAG as an enabler to improve the catalogue of services offered to their users, and the possibility to reducing their dependency on Windows Server, also alleviating their support overhead.

4.3 Remote Operations Gateway

In collaboration with the IT department, Beams department took on the lead in 2020 of creating a central cluster for the ATS sector, entirely managed by Beams department, which they called Remote Operations Gateway (ROG) and became operational in July 2021.

ROG provides access to CERN Technical Network (TN) and software from a web browser. It supports maintenance and monitoring operations of the accelerator complex. It is used by external companies but also by Controls operators and experts when they are not on CERN site. It is the first CERN secure gateway which gives a direct access to the TN from external locations.

Compared to other remote access facilities at CERN, ROG has a distinct set of features:

- Enforce 2FA for all users.
- No access to Internet.
- Access to CERN shared filesystem (CERNBox & EOS, AFS), but no persistent home directory.
- Same environment as operational technical consoles (XFCE with custom launcher).
- Central security monitoring aligned with similar services expectations.
- Centrally managed updates aligning with the LHC and injectors schedule.

This is a balanced set of features designed to offer a simple method for creating sessions while also reducing the need for support calls outside of working hours. An Ansible [8] (configuration management tool used in ATS sector) role has been created to configure the different parts of the cluster as mentioned in Section 3 of this document.

To improve users' accessibility, additional widgets and tools have been installed or developed, for example: a Kerberos widget reminding users to renew their ticket when it expires, or a custom launcher to have the same environment as in the Controls rooms.

Several FastX bookmarks have been created to assist users in transitioning from Windows to Linux. Users who are accustomed to a particular remote access technology might find it challenging to adapt to a new one. This requires training and support to ensure a smooth transition. A video training was provided and can be updated as needed.

FastX bookmark system also allows Controls sysadmins to test new operating system or updates in an easy way without compromising access at any time. In 2023, both CERN CentOS 7 and Red Hat Enterprise Linux 9 (RHEL9) sessions are available, which make ROG one of the first CERN service partially migrated to RHEL9.

The main challenges are the migration path between minor releases; version 3.1 to 3.2 changes were quite disruptive (main database technology changed) and has proven to be a challenge to do an in-place upgrade. Version 3.2 to 3.3 migration path is far smoother but changes to the main web interface layout, which presents another challenge for users.

Starnet provides security updates and new features only to the latest supported minor release (3.X), which forces system administrator to upgrade in a timely manner, constituting a challenge in our setup if the architecture changes.

Starnet support is reactive, and all our proposals were considered. Even HTML/CSS customizations have been discussed: they proposed to give us access to part of the sources to apply our patches.

5 Conclusions and future work

In a work environment where remote connections to distant desktops or protected networks is necessary, organisations must provide the best and simplest tools for its users, regardless of their required operative system or preferred client.

The aim of ROG is to guarantee a similar experience to being in a CERN Control room for the LHC operators, while ensuring that both security and ease of access are not compromised. For this reason, CERN aims to expand the project to reach more communities, providing them with a wider and more adaptable range of services to meet their needs.

The ATS sector's move to FastX is focused on two primary aspects: transitioning from Windows to Linux for some users and switching to a new remote access method (through the web browser). After 24 months ROG is widely used across ATS sector users, with 160 concurrent daily sessions and 600 registered users on 10 virtual machines.

As Kubernetes takes the lead as the primary orchestration platform, efforts are being made to explore the possibility of migrating future versions of FastX to be compatible with cloud-native environments. Additionally, there are ongoing projects aimed at evaluating container-based solutions for xrdp to enhance isolation and security measures.

CERN will continue to investigate Linux remote desktop solutions and profit from modern orchestration tools to maximise resource usage without sacrificing the user experience.

References

1. S. Bukowiec, R. Gaspar, T. Smith. *Windows Terminal Servers Orchestration*, 22nd International Conference on Computing in High Energy and Nuclear Physics (CHEP 2016)
2. *The MALT Project. Re-assessing the IT provisioning Strategy for Core Services at CERN*, [Online]. Available: <https://malt.web.cern.ch/malt>
3. *TigerVNC*, [Online]. Available: <https://tigervnc.org>
4. *X2Go – everywhere@home*, [Online]. Available: <https://wiki.x2go.org>
5. *xrdp, an open-source Remote Desktop Protocol server*, [Online]. Available: <https://www.xrdp.org>
6. *NoMachine – Free Remote Desktop for Everybody*, [Online]. Available: <https://www.nomachine.com>
7. *FastX – Best Linux Remote Desktop Software – Starnet*, [Online]. Available: <https://www.starnet.com/fastx>
8. *Remi's RPM repository*, [Online]. Available: <https://rpms.remirepo.net>