Free / Open Source
Software Solutions
For Educational Institutions

Technical Details, Common Problems and Benefits

Self Hosted
Education

https://selfhosted.education

Contents

Overview

Presented below is an overview of solutions that are essential and indispensable for any college campus or departmental network.

The goal of this note is to explain what these solutions are, how they work, why they are useful and how they provide more value to everyone while enabling us to become more effective in whatever we do.

There are three core ideologies which form the basis for our discussion here:

- Free / Open Source Software,
- Self-Hosting, and
- Build-it-Yourself

Free / Open Source Software

Free / Open Source Software is software that provides the following four freedoms to us:

- Freedom to use the software for any purpose
- Freedom to study the source code of the software and learn from it
- Freedom to modify the software and make changes to it
- Freedom to share the software (in its original and modified form)

These freedoms are essential as they provide us with the essential framework for learning and sharing knowledge.

The freedom to use software for *any purpose* ensures that *anyone* can use the software for any endeavour *(educational, personal, commercial, non-profit, governmental, defence and more)* and in any region. Free Software[1] does not discriminate.

The availability of source code provides us an opportunity to study and understand how something works. Without this freedom we would never be able to learn from what others have done or apply that learning to our own lives.

The freedom to modify the software enables to learn and gain confidence by doing things. It allows us to extend the software, fix bugs in it, translate or localise it and adapt it for other purposes.

Finally, the freedom to share a program (in its original or modified form) allows us to help others while also providing subsequent users with the same sort of *freedom* that we enjoy. It ensures that the software remains free and provides every user, teacher and developer *equitable freedom*.

Proprietary software *(ie. software that is distributed only as a binary and where we are restricted from accessing its source code or modifying it or sharing it freely)*, on the other hand, does not provide us with any opportunity to learn or understand how things work or make changes to them to test and validate our understanding.

[1] For the purpose of this document, we use the terms "*Free Software*" and "*Free / Open Source Software*" interchangeably to refer to software that provides freedom to us. While "Free Software" is the original term that was used to describe such a body of software, the term "Open Source" has also gained a lot of popularity . We consider "Free Software" to be a more useful term since it captures the essence of this philosophy and does not dilute it for whatever pragmatic or marketing reasons.

Self Hosting

Today, it is also possible to use and consume software that is hosted on “somebody else’s computer” (a.k.a “the Cloud”). Such software (or services) are beyond our control in every possible way because we do NOT host it on our own computers and hence, don’t enjoy the necessary freedoms discussed above.

Self Hosting is essential for educational institutions because it is the only way to test and validate if a solution is really “free”. It is the only way we can have control over what software we use, how we deploy it, what it takes to deploy and manage it, who uses it and how.

Unless we can self-host a solution, we lose the ability to choose or consider Free / Open Source Software and also protect our rights and freedom.

Build it Yourself

Choosing to “*build something*” instead of “*buying it*” is another way to choose free software and self-hosting. When we build something ourselves, it helps us cultivate understanding about the system. This understanding liberates us, makes us self-sufficient and gives us confidence in our knowledge and how we choose to apply it. Building something makes us “true owners” of a solution instead of being plain consumers and users of it – true ownership stems from our ability to build, maintain and extend a system independently.

Why is this relevant for us?

If we choose to have no control over or understanding of the technology that drives our educational institution, we eliminate the opportunity to acquire and share knowledge. If a core part of learning is to “learn by doing things”, then choosing technology that is proprietary and not self-hosted makes it impossible to be effective at learning or teaching; and instead, just makes us “passive users”.

While it is true that not everyone is equally curious about how things work and not everyone desires this freedom or opportunity to control technology that affects their life and work, that is not an argument against having such freedom and opportunity. Unless we mindfully build our “technology stack” by choosing Free Software and Self-Hosting, we are making dangerous choices for future generations – choices that they did not agree to make or were not even a party to.

Hence, in an educational setting, it is our responsibility to protect our own freedom, promote such choice in technology and create greater opportunities for learning and building things.

What can we build?

Cloud & Hardware infrastructure for IoT projects and implementation	ESSENTIAL SOLUTIONS
Virtualisation, Private Cloud & Infrastructure / Platform as a Service	
Centralised Log Management and Monitoring	
Source Code Management, Versioning, Collaboration, Review, Release, Packaging, Documentation & Deployment	
File Sharing, Collaboration & Media Archival	
Help Desks & Request Tracking	
Email, Communication & Collaboration	
Lab Automation & Local repository mirrors	
Identity Management & Centralised Authentication	
Network Access Control, Captive Portal, BYOD Management, Centralised Wired and Wireless Management	NETWORK SOLUTIONS
Campus wireless networking and Access Points	
Network security and Intrusion / Malware Detection	
IP Address management, name services, remote access and software-defined networks	
Firewalls, Bandwidth Management and Link Aggregation and Failover	

Network Solutions

Firewalls, Bandwidth Management and Link Aggregation and Failover

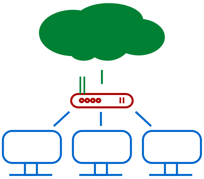

Overview

A firewall is a hardware + software solution that connects and shares Internet bandwidth with your local network. It isolates and protects the local network from the Internet and also allows you to securely expose specific applications, services and servers on the local network.

A bandwidth management solution helps you monitor your bandwidth usage and then allows you to define bandwidth and Internet usage priorities and policies.

A link aggregation and failover solution lets you connect multiple Internet links to your firewall so that you can have more bandwidth available to users on the local network. This also means that you don’t rely on a single Internet link or ISP for your Internet connectivity needs.

Common Problems

- You have a firewall but you need to repeatedly call a service provider or vendor to manage the rules for you.

- You have a lot of Internet bandwidth but your Internet usage and experience is slow.
- When large or bulk files are repeatedly downloaded, it causes a slowdown for other users.
- It is impossible to account for your bandwidth utilisation.
- When your main Internet link goes down, or is slow, it cuts off the entire network from the Internet
- You are having to depend on third-party ISPs to provide high-speed and unrestricted access to the Internet
- Even when you have multiple links to the Internet, you might not be able to use them all at the same time. Or specify policies on who can use which one at what time.

Benefits

- No dependency on 3rd party for firewall management.
- No risk of usage data being sent to a firewall vendor's cloud setup for analysis - thereby, ensuring better privacy for users
- Optimal utilisation of all the available Internet links simultaneously and automatic failover
- Fair sharing of bandwidth amongst users so that prioritised traffic and users always get the required bandwidth
- Actionable Internet utilisation reports
- Caching of frequently accessed content from the Internet saves bandwidth and improves user experience
- Extended caching of bulk content like OS updates, Anti-Virus updates etc
- Granular Web Access controls depending on user, time of day, day of week, content-type, request-type etc

- Logs can be retained for as long as required for data and trend analysis
- User identity and authentication is mapped from a singular, campus-wide service that is used for all other applications as well

Network security and Intrusion / Malware Detection

Overview

Keeping the network secure is as important as keeping it fully-functional. The traditional approach of just using an equipment at the gateway is not sufficient. The threats can emanate from inside the networks as much as they come from the Internet.

A functional and up-to-date Intrusion Detection and Prevention System (IDS/IPS) is an integral part of a secure network in addition to anti-malware software. An IDS/IPS inspects all the traffic passing through the network in real-time and matches them with pre-defined signatures to generate alerts and/or block the malicious activities from outside or inside the network.

A SIEM (Security Incident and Events Monitoring) system can collect logs generated from various components like firewalls, servers, IDS/IPS, network access controllers, DHCP Servers etc to provide a holistic view of the security situation of the network.

Common Problems

- There are infected machines in the network which are generating huge amount of traffic affecting the whole network. These things can happen silently without any outward symptoms till it begins to affect users in a negative way.

- There are Trojans installed in the machines which spread through the network or can be part of a bigger attacks and are just waiting for a trigger from “Command and Control” (CnC) server on the Internet
- You are not even aware that there are a lot of Intrusion or Denial-of-Service (DoS) attacks are happening on your network which are eating system resources of your gateway firewall
- Students could be trying to learn about security by attacking the servers inside the college network (or on the Internet)
- Unauthorised access attempts to the servers or computers inside the network are not noticed

Benefits

- Visibility into malicious activity present in the network - originating from either the Internet or within the local network
- Automatic blocking of the hosts generating such activities from the Internet
- Automatic isolation of the hosts generating malicious traffic from within the network with the users and administrators get notified of the same
- Various vulnerabilities get highlighted on the SIEM which when fixed on time save the network from catastrophic downtime
- The possibility of machines inside the network getting involved in large scale attacks on the Internet gets minimised to a great extent

Campus Wireless Networking and Access Points

Overview

To provide network access to the users in the campus everywhere, wireless networks are required. A large campus requires hundreds of access points to be installed and managed in order to provide equitable, highly available, stable and high-performance wireless service to users. Additionally, the wireless network is required to be compliant with all wireless standards to support legacy as well as latest user devices.

The free software friendly approach to wireless deployment is to first standardise on a set of hardware platforms, technology stacks and operating systems to build such infrastructure. By using free software and minimal operating systems such as OpenWRT, we can build a unifying operating environment that remains the same across hardware from multiple (and compliant) manufacturers.

Additionally, it is possible to build an integrated and centralised management solution for all these wireless access points that monitors, manages and controls them – irrespective of their hardware features, lineage or age.

Common Problems

- Vendor does a site survey and deploys access points as per their observations. Conditions and wireless coverage and performance requirements could change over time, requiring repositioning or addition of access points. These additional access points would, then, need to be from the same vendor / make / control system.
- Users complain about poor connectivity and speeds on the wireless network
- It is difficult to identify problems caused due to wireless interference or overloaded access point
- It is difficult to provide the same sort of network access to users on wired and wireless networks alike
- It is cumbersome providing network access to the guests
- Any hardware upgrade or vendor change for access points puts you through steep learning curves repeatedly.
- Replacing or upgrading wireless infrastructure every few years (by changing vendors or deployment methodology) might still not address the real problem with providing reliable wireless acess to users

Benefits

- You learn to perform site surveys on a regular basis without having to depend on external vendors and service providers
- You can centrally manage all the access points from a single management interface
- Better visibility about SNR, connected stations and interference

- Firmware upgrades across different access point hardware doesn't change the management interface and hence, there is no repeated learning curves
- Best available off-the-shelf hardware irrespective of the vendors can be used with the latest released firmware providing more choices and freedom

Network Access Control, Captive Portal, BYOD Management, Centralised Wired and Wireless Management

Overview

A campus network is bigger and more complicated than a large corporate network considering the number of type of users and devices it caters to. Devices in the departments, labs, libraries, offices and hostels need network connectivity and have varying access requirements. On an average, a typical user can have upto 5 devices like laptops, phones, Raspberry Pi and other single board computers, smartwatches, IoT boards / sensors etc.

To manage such a huge network, it has to be broken down into smaller subnets - one each for different departments, labs, administrative offices, libraries and hostels etc.

A Network Access Control solution verifies every device and user when they try to access the network and only allows authorised devices and users. This role based access ensures that a user on specific devices gets access only to certain resources on the network e.g. an MCA student's laptop connected to the wireless from anywhere in the college will still be put into the MCA department's network.

A Captive Portal can be used for authentication of the campus users or the guests and grant them specific access only.

A campus is a very large Bring Your Own Device (BYOD) environment where differentiating between user's own devices and institute's machines is important for granting specific access to the users. A user from a mobile phone might not be allowed to access lab servers but is only able to access Internet and library's catalogue server.

The same user when connected through a laptop is granted access to the labs, library and the Internet.

It is imperative to manage official devices and provide requisite access to user's own devices in a secure manner irrespective of the network media (wired or wireless) the devices are connected over.

Common Problems

- You are unable to control mobile device usage in the network as you are not able to identify the device type reliably
- It can be frustrating and unsustainable to build access control rules based on IP addresses
- It is difficult to manage access control for multiple IP addresses of a user's device when using wired and wireless networks
- You are only able to control the Internet access of a user, but all other resources inside the campus network are exposed to the user and there is no mechanism to grant or deny access to specific resources on the internal network.

Benefits

- Simply connecting a computer to the network doesn't grant access to whole network. User and device based roles decide what access is granted.
- It becomes possible to restrict access to specific internal network resources to a user or a device
- Access management shifts focus from IP Addresses to the user's identity and the type of device

- BYOD access control becomes manageable and device types (laptop, mobile, tablet) etc. are identified automatically
- Providing network access to the guests becomes easier with self-service portals

Essential Services & Solutions

Identity Management & Centralised Authentication

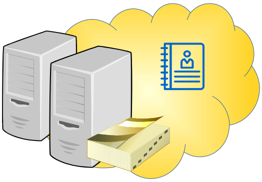

Overview

Each network deserves to have a solution and vendor independent identity management system. Such a system builds a centralised authentication service which enable users to use the same authentication credentials across all network and server applications.

Identity management solutions also offer users a self-service method of resetting their own passwords as well as managing their own account information.

Common Problems

- Your user database was created as a part of the network firewall implementation and those user accounts can not be used for other purposes.
- You have other special purpose user databases which are not comprehensive and are limited to users who need to access a particular applications
- Users in the campus typically have to remember a bunch of username + password combinations for accessing various facilities in the local network.

- Your email service is provided by a cloud service provider and that user database has no connection with your local network user databases
- There are actual financial costs associated with adding more users to some of your user databases and hence, you only add enough users as you have licenses for.
- Users create and share their login credentials with others and that makes it difficult to really identify users.
- You would like to add a second-factor authentication system (based on OTPs or hardware USB tokens) but it would not be possible to integrate that system with all applications.

Benefits

A identity management solution implemented using Free / Open Source Software attempts to solve the problems listed above.

- You build the user database in a stand-alone identity management solution.
- You can securely access the user database over the standards-based LDAP protocol. This makes it trivial to integrate almost any sort of application or service and have it authenticate users from the same identity management solution
- You can organise users into groups and map various attributes to them. These groups and attributes can be used in any application that can map access control based on groups or arbitrary attributes.
- You can integrate GNU/Linux-based desktops into this authentication system so that all authentication on the

desktops is done centrally. This eliminates the need to create the same set of users on each GNU/Linux desktop or manage the credentials for these users.

- An LDAP-based user directory is extensible and flexible. You can add custom attributes and interact with it using variety of command-line, GUI and web-based tools.
- You can also write code in all programming languages to talk to the user directory over LDAP and query or authenticate or manage users.
- You can use RBAC (Role-based Access Control) and HBAC (Host-based Access Control) features to build all sorts of flexible access control policies for applications
- Centralised sudo policies can be implemented so that they apply all over the network
- You can build a distributed and hybrid user directory as well where-in each department or user group manages their own sets of users – instead of someone doing it at the institute level
- An identity management solution can also serve as a central store for encryption public keys and SSL certificates. These keys can then be used on any server, desktop or device and can be used to implement a variety of security techniques and protocols.
- For Students: When students build projects that need an authentication source, they can either replicate this setup for their own project. Or they can simply implement LDAP-client functionality and plug in to the already existing user directory.

- For Faculty Members: Faculty members can deploy any application which can use LDAP. They don’t need to worry about adding users again. Almost all free software web-applications can talk to LDAP and it is easy to use existing user accounts with them.

Lab Automation & Local Repository Mirrors

Overview

Lab Automation enables administrators of computer labs to provision and deploy operating systems and then manage them centrally and automatically over the network. Such a solution eliminates repetitive work for installing operating systems and software and adding configuration to each lab computer.

Common Problems

- Lab Administrators typically re-install the operating systems on lab machines 2 to 4 times every semester.
- This re-installation is needed to have a clean operating systems on each lab computer devoid of all user data.
- A re-install could also be required in case of hardware problems or failure
- Apart from the OS, common software would also need to be installed on each computer in the lab
- Doing these things manually can can not only be cumbersome but also error-prone
- When users install common GNU/Linux packages from the package repository of a distribution, the installation is more often

done over the Internet. Installing the same software on all lab computers results in the same software being downloaded multiple times from the Internet. Not only is this time-consuming but also a waste of bandwidth.

Benefits

A lab automation solution enables rapid and unattended deployment of operating systems as well as application software over the network. One can effectively boot machines off the network (using PXE or EFI boot images) and choose to install a variety of operating systems automatically on a given machine. Once the process is automated, it can be done more frequently and on-demand.

Automation also allows us to do extremely complex software installation and configuration on each lab computer. Such tasks would otherwise be very difficult to do manually and hence, very few people would attempt them owing to the complexity and effort involved.

Once an OS is provisioned, the ongoing maintenance, upgrades and customisation can be done using automation tools such as Ansible. These playbooks allow one to perform a set of tasks on a group of computers in an unattended manner.

This lab automation framework is readily coupled with a local mirror of various GNU/Linux distributions. This greatly speeds up package installation, saves Internet bandwidth and opens up so many possibilities for rapid and repeated software installation and OS deployments.

Faculty members and lab in-charge do not have to worry about students making mistakes and messing up the OS on a desktop – when the cost of re-installation is so minimal (and completely non-interactive and automated), there is no incentive to be “careful”. Instead, students can be motivated to experiment and learn from experience, instead of trying to be “safe and correct” all the time out of fear of spoiling something.

When a workshop or training is being organised, the process of installing software on laptops and even single-board computers can be speeded up because of the local repository and package/software mirror.

Email, Communication & Collaboration

Overview

Email, communication and collaboration solutions offer simple, yet effective ways in which everyone on the campus can interact, discuss, communicate and keep in touch.

These solutions cover the following services:

- Email – SMTP / IMAP / Webmail services for simple and reliable email communication
- Group Chat – Public and private group (rich text) chat services that allow real-time collaboration amongst users in a group
- Instant Messaging – Private one-to-one real-time rich text chat services
- Newsletters – The ability to send out email newsletters to various interest groups on and off the campus
- Conferencing – Real-time audio + video conferencing solutions for discussions, training programs and meetings
- Discussion Forums – Public (and private) discussions forums to allow users to collaborate to solve problems publicly and use the same resource to also search for solutions over a period of time.

Common Problems

- Users on the campus try to solve their collaboration requirements in a variety of ad-hoc ways. This causes three problems:
    - Everyone chooses a different platform with varying levels of usability, constraints and features
    - There is a proliferation of user accounts and hence, there can not be any standardisation of user names across multiple third-part communication platforms
    - There is no indelible or permanent archive of important discussions; sometimes, the data is locked into a proprietary cloud-hosted service
- Data related to various types of collaboration and communication is distributed across multiple, sometimes-proprietary and cloud-hosted platforms
- When people are not aware about solutions that promote privacy and security, they can sometimes opt for solutions that might be “dangerous” to them in the long run
- Many proprietary and cloud-hosted platforms have tracking features (in lieu of their “cost-free” nature)

Benefits

- Local network email services will add tremendous value to students, faculty members and others to use email securely within the local network and beyond
- Students can also develop software that uses email and do setups for systems that can use email for notification and other purposes. A local network email does not penalise

students to have more complexity in their programs. Addtionally multiple types of mail servers can be deployed for production, research, testing and other needs.

- All email alerts generated by the local infrastrucutre setup, by Gitlab and other such local services will remain within the college infrastructure itself and can be easily accessed by everyone
- Group chat is the new real-time collaboration paradigm. Instead of requiring Internet access and access to cloud-hosted proprietary group chat systems, a locally hosted group chat system powered by Free Software enables students to collaborate and communicate with faculty members and peers with security, privacy and ease

Helpdesks & Request Tracking

Overview

Helpdesks and Request tracking solutions serve a very niche and important need on a college campus. They enable us to setup and track requests of all types and build general purpose as well as special purpose helpdesks such as:

- IT Helpdesk – user's can report IT infrastructure related issues so that they can be looked in to and handled
- Students Helpdesk – for students to raise queries with college authorities
- Hostel Helpdesk – for hostel level tracking of problems or incidents
- Department Helpdesk – for department-level issue tracking
- "HR" Helpdesk – for the college faculty and staff members to raise requests with the college
- Account Helpdesk – for accounting and finance related queries
- RTI Helpdesk – to track and handle "RTI" queries

Common Problems

- Users on the campus report and track issues in an ad-hoc manner – worst case by trusting their memories
- There is no central record of how many problems were reported, who handled them and at what stage of resolution these requests are at

- In the absence of a tracking system, it is also impossible to “report” on the performance and effectiveness of a problem or issue resolution system
- When users can’t access a self-service system to know the status of their issues and requests, they need to spend time and effort following-up, making calls or sending messages.

Benefits

Organising requests into a helpdesks provides the following benefits:

- The one place where everyone can go to to raise requests or report issues
- Everyone is a requester at some level, an “agent” (users who actually handle / address the reported issue) at some level and a helpdesk “administrator” at another level
- People don’t have to rely on their memory to fix problems
- Reporting and analytics on raised issues can give in critical insights into the type of problems that happen more often and in large numbers – this provides an opportunity to make pre-emptive steps to avoid such problems in future and improve processes and systems to make them more effective, efficient and useful
- Reminders can be sent out, escalation workflows can be built and “SLAs” (request / issue resolution timelines) can be implemented
- It can also useful to extend helpdesks to record, track and respond to queries generated via various websites (college, department, event, club etc)

File Sharing, Collaboration & Media Archival

Overview

Users often resort to using a variety of cloud-hosted services to share files with each other (on-campus) as well as others beyond the campus.

File Sharing services aim to provide an integrated, local-network and default solution for this requirement of sharing and syncing files. Furthermore, they can collaborate on working with files and content.

Media archival solutions aim to address the issue of being able to locate and access all types of photos or videos of various events on the college campus from single and reliable location.

Common Problems

- Everyone, at all levels creates documents – text documents, presentations, documentation, spreadsheets, papers, thesis, project reports. When they need to share these documents or collaborate on them with others, they typically use a communication system (email, instant messaging etc) to share these documents
- People typically use public services such as Google Drive or Dropbox to share files with others – sometimes even within the same lab or local network
- There are other types of files that users create and need to access and share. When there is one single source from where

files can be accessed or shared, it is confusing to organise the files

- Files proliferate all over the place and we end up having multiple copies across all our devices (desktops, mobile phones, laptops, servers)

Benefits

- A file sharing system such as NextCloud enables users to effortlessly access, sync and share files across all their devices.
- Since it integrates with the local identity management and centralised authentication system, it is trivial to share files with everyone
- Files and documents shared by users on the campus stay within the institution network. This promotes privacy, can contain and prevent data leaks and provides high speed access to these files
- Users have one default and recommended system to access and share files – they don't need to make decisions on what system to use every time they need to share such files
- Users stop using communication platforms as file sharing platforms – instead of sharing files via email or chat or discussion forums, users can share links to the files
- Users of this system get personal as well as group-based collaboration workspaces. Groups can be project-bases, course-based, department-based and so on.
- A document management and collaboration system enables people to work on documents together without having to share then back-and forth with others

- A media archival solution enables the college to organise media (photos, videos, audio recordings) of various events (seminars, sports events, conferences, lectures) in one central place. These media can then be shared securely. Additionally access control can be built so that only those allowed to access a specific portion of the content get access to it.
- Have a central and reliable media archive eliminates the need for users to keep such files on a variety of disparate media at their personal or department levels.
- Proper data integrity, access control, encryption and backup procedures can be instituted to protect users’ data in a central and local-network server

Source Code management, Versioning, Collaboration, Review, Release, Packaging, Documentation & Deployment

Overview

Gitlab is a comprehensive Free / Open Source Software solution that aids software development collaboration using a variety of development tools and frameworks. It allows us to build a very capable self-hosted solution around the concept the git version control system.

Common Problems

- Most users are not aware of version control (or git) based software development and collaboration workflows. However, almost 100% of software development (also 100% of Free / Open Source Software development) collaboration + contribution happens over version control systems
- When (if) people realise that version control can add value, they either pay to use a cloud-hosted private repository or use a public repository free-of-charge
- In absence of version control systems and habits closely integrated with academics, ad-hoc strategies for project management, tracking and reporting are not as effective or productive

Benefits

For Students

- Students can build private, shared, collaborative and college-wide projects and code repositories without any restrictions, constraints or costs
- Students' work is private and secure at the college level and so they can feel free to experiment and write / share their source code and work without the fear of others reviewing or judging their work publicly on the Internet
- Access to the source code repository on Gitlab is very high speed
- Students can learn how to use the git version control system – the most popular version control system among software developers world-wide
- Students can use Gitlab to implement the entire "software development life-cycle". This includes project and task planning, scheduling, development, project release, bug-fixing, packaging and so much more.
- Gitlab can also enable students to understand concepts of continuous integration (CI) and continuous deployment (CD) – so that software they write or collaborate on can be automatically tested, packaged and deployed whenever new code is pushed to the repository
- Gitlab will also enable students to write, collaborate on and publish documentation along with their projects

For Faculty Members

- Git version control is an indispensable skill in software development and systems administration today.

- Git is also a very power tool (and workflow + habit) to make source code, assignment and project submissions.
- Gitlab offers very useful project management and planning features as well as source code review and feedback features. These can be used by faculty members to build and institute practices for project planning and execution by students.
- Students can be asked to use git push to upload assignments (source code, plans and documentation etc) to the Gitlab server so that faculty members can review and critique the submission.
- Git also enables the creation of an indelible record of when and what was submitted. Faculty members can use this feature to track timely progress on assignments and projects so that students put in consistent work instead of doing all the work in the last few days.
- Since Gitlab has a CI/CD system integrated with it, faculty members can also write their own test cases to automatically check if the submitted source code compiles and executes, if it follows coding guidelines and if it gives output as expected.
- Faculty members can give feedback to students via the ticket system so that their feedback and the student's response is recorded as a part of the project or assignment process and is available for learning by other students and faculty members as well.
- Faculty members can also teach the basics of how to collaborate to write, release and deploy projects using the "pull requests" model
- Gitlab ships with a built-in container registry. So software hosted on Gitlab can be easily packaged and deployed from this registry.

- This automatic deployment of software can happen to a private cloud infrastructure consisting of virtual machine and container hosting platforms.
- Compiling and publishing documentation also becomes a part the same workflow as writing source code

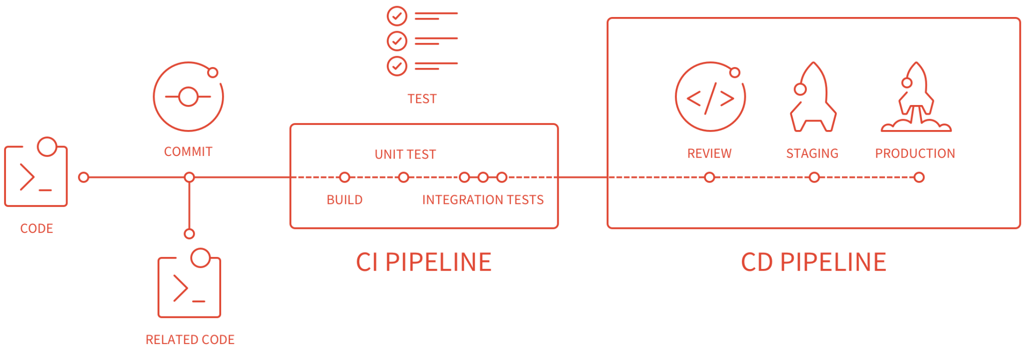

The Gitlab CI / CD Workflow

(Source: https://docs.gitlab.com/ce/ci/README.html)

Virtualisation, Private Cloud & Infrastructure / Container / Platform as a Service

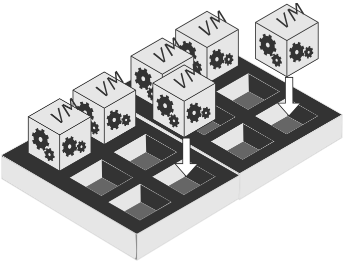

Virtualisation & Instrastructure as a Service (IaaS)

Overview

A Virtualisation Solution available on the local network will provide students and faculty members with the freedom, performance and flexibility to deploy operating systems and work in Virtual Machines. They can do all this without any added or running cost for using these VMs (as on the pubic cloud) and also without Internet access.

Students can learn to deploy virtual machines, install and configure operating systems within them and also learn automation and scripting techniques that are so relevant on the cloud today.

This virtualisation server also provides exposure and a learning opportunity to how people deploy and run software in production.

Common Problems

- It is difficult to create development or testing environments

- Some development or teaching platforms are complex and time-consuming to deploy
- There is no way for students to create long-running and dedicated “computers” to demonstrate or develop or collaborate on their projects
- Code or setups which might work on one computer don’t work on another because it is complex to build or re-build the same environment multiple times
- When projects require more than one machine, it is difficult to allocate such computing resources to projects and keep them available for a considerable period of time
- Faculty research projects which require a complex mix of computing infrastructure are difficult to execute
- It can take considerable time to provision or reserve computers (or VMs) – especially when their lifetime is short (ie. someone needs 1 or 2 computers just for an hour or a day)
- Many times educational institutions use the public cloud to do their work – this can become very costly at scale if done repeatedly

Benefits

- The large scale benefits of using the public cloud are obvious. One can provision a new server in minutes and get productive with it immediately. Building such a local and private cloud would require hardware investments (proportional to the desired scale of implementation). However, these investments provide extremely high long-term value and extensive learning and teaching opportunities.

- Virtual machines can be provisioned quickly once the infrastructure is setup and made reliable
- Users can trust these VMs to offer good performance and reliability
- VM deployment can be automated using cloud-init and Ansible
- VMs can be easily integrated into the Gitlab CI/CD pipelines
- Faculty members can create and share development environments using tools such as Vagrant. These development environments can serve as a reference for production environments as well.
- Virtual Machines can be allotted to students for project work and demos. These VMs will remain alive for the duration of their project and can be accessed (or access controlled) by others as well (even from the public network)

Container, Function and Platform as a Service

- While VMs provide a rich, secure and isolated environment to run any operating system, such systems also require setup and administration.
- Using Linux containers, Docker and Kubernetes, today we can provide multi-level deployment environments to users.
- These can be used to deploy simple docker containers and clustered Docker containers for testing and using applications.
- Platform as a Service (PaaS) can provide the highest level development and deployment environment to users. Using simple git-push deployment techniques, code pushed to the PaaS can be automatically deployed without having to setup web services, key value stores, databases and such.

Other Benefits for Students

- Students will gain complete hands-on experience on the tools that are regularly used in industry and the FOSS community. Pro-actively building skills in such tools will make it possible for them to contribute to FOSS and also pitch for systems design and administration, research and software development opportunities professionally.
- Awareness of methods and practices used in industry or FOSS community: It is not just tools that are essential – awareness and experience with established and ad-hoc methods and practices will enable students to fit in easily and quickly into global communities and participate effectively.
- Opportunity to experiment and learn safely without affecting the production deployment of services in the lab. Students should not just be users and consumers of technology. By demonstrating the fact that such infrastructure can be built effectively and affordably by choosing Self Hosting and Free Software, we can enable them to also attempt to do the same and learn and build so much more.

Self Hosted Education

https://selfhosted.education

The Self Hosted Education Initiative is a public service initiative led by like-minded Free Software Businesses like DeepRoot Linux and Unmukti Technology.

The Visvesvaraya Technology University (VTU) whole-heartedly endorses this initiative and encourages colleges to use Free Software for building infrastructure to provide better and more effective learning, research, development and teaching opportunities in educational institutions.