World Wide Web

We need to connect our server to the outside world now that everything is ready behind-the-scenes.

- What are Computer Networks?
- How to Remotely Connect
- Virtual Private Network
    - Hosting a VPN Server
    - Connecting to Your VPN
    - Accessing Your Services
- Web Domain Name
    - Domains & URLs
    - Getting a Domain Name
    - Domain Name System
    - Reverse Proxy
    - Connecting Services to the Reverse Proxy
- Digital Stewardship
- Security & Privacy
    - Critical Thinking
    - Basic Authentication
    - Authelia
    - LAN-Only Access
    - fail2ban
    - CrowdSec
    - Restricting Access by Geographic Region
    - Hiding from Search Engines
- Router Configuration

What are Computer Networks?

We use the internet everyday, but it isn't often – if ever – that we need to consider how it works.

That's because the protocols powering the internet were intentionally designed to operate as invisibly as possible. Built around a common language, networks enable devices to communicate with each other and share resources. By standardizing how computers talk with each other, we have expanded the scale of networks over time seeking to achieve a global cloud infrastructure.

> *How we are at the small scale is how we are at the large scale. The patterns of the universe repeat at scale. There is a structural echo that suggests two things: one, that there are shapes and patterns fundamental to our universe, and two, that what we practice at a small scale can reverberate to the largest scale.*
>
> — adrienne maree brown

Connected Communities

Computer networks consist of nodes – which are devices that are seeking to communicate – as well as the links between them. Under some circumstances, nodes will connect to other nearby nodes and create a mesh that data can traverse while seeking it's destination.

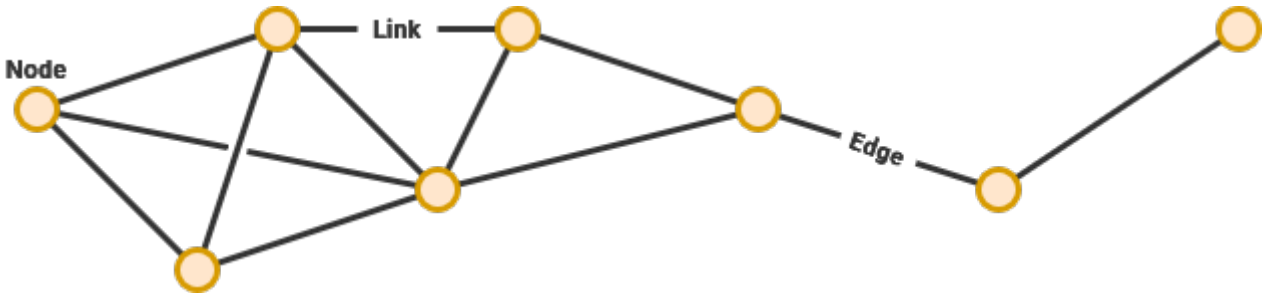

Edge networks form around bottlenecks that may arise in rural and disconnected areas. Functionally, this means that every person trying to connect to the World Wide Web must share a single Internet route – like an overly congested highway. This can increase the cost of internet service, as well as effect the overall speed, availability and reliability.

The ways that communities connect to the Internet relies entirely on what is available within their physical geography. While in a major city, there may be many options for access to fast and reliable fiber Internet. Trying to connect to the Internet from rural Alaska, communities may find

themselves restricted to a satellite connection.

These "digital deserts" can arise along geological boundaries – such as mountains or islands. More importantly, marginalized areas – such as Black, Indigenous and Hispanic communities – are not always offered equal or adequate Internet access. During 2021, it was estimated that over 42 million Americans do not have access to terrestrial broadband Internet – with 4 million in Texas alone.

> What reasons – geological, political and social – do you think contribute to "digital deserts" without internet access?

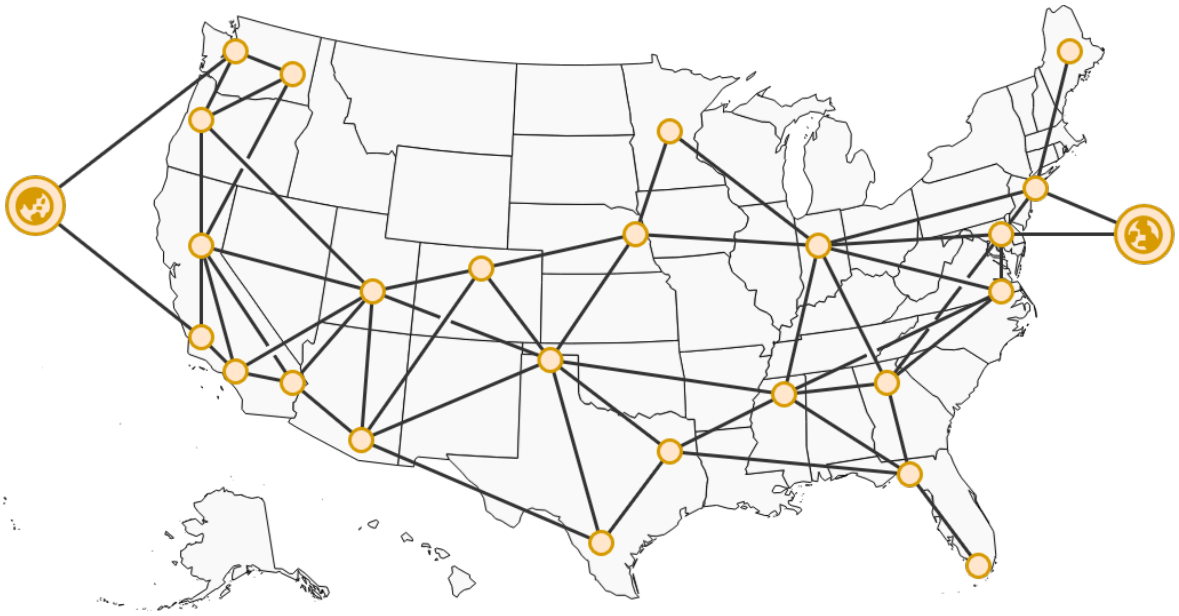

Networks, similar to the rest of computing, leans heavily on abstractions that enable people (and network engineers) to comprehend the infrastructure required to power telecommunications at this scale. In order to build a global infrastructure, digital technologies have created a stratified system that simplifies data shared in between these layers.

This foresight during the creative process has worked to simplify our relationship with technology. You don't need to understand electrical engineering to build a computer system from parts you bought off-the-shelf. Similarly, data isn't concerned about routing its own path across the internet and only follows the one assigned to it.

Scales of Connection

Scale is a foundational way these systems are abstracted when trying to classify them. This is important to consider because different networks may have unique requirements. Your home

network only needs to juggle a handful of people's data, while a college campus will be handling much more traffic from people located around the globe.

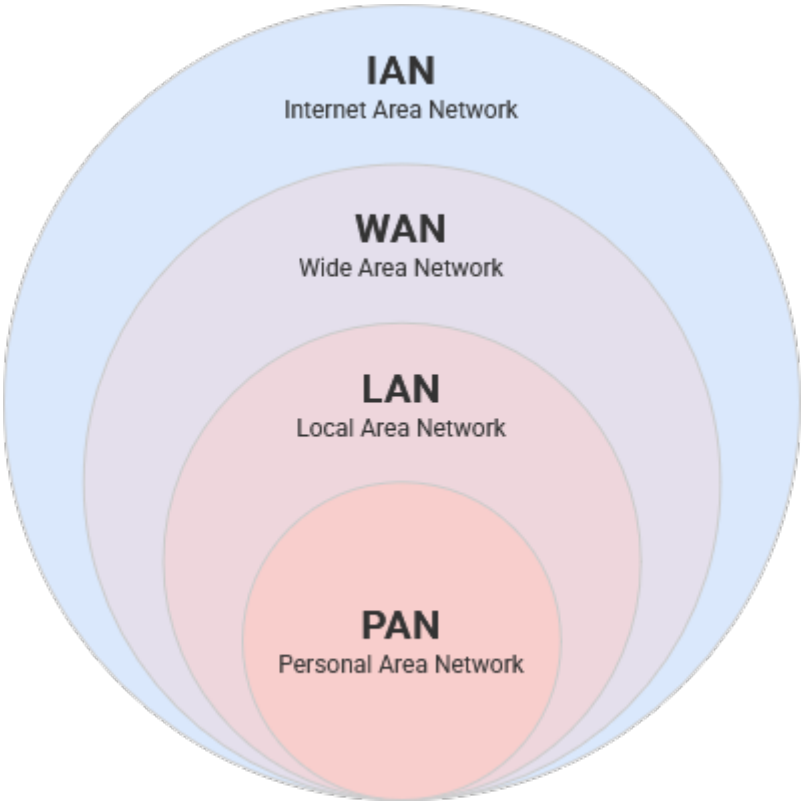

Personal Area Network

While your phone is connected to your headphones through BlueTooth, you are creating a Personal Area Network. These, as the name implies, operate on a smaller and more intimate scale. PANs utilize wired and wireless technologies to connect to each other, either through a hub device – like a cellphone or laptop – or directly to each other.

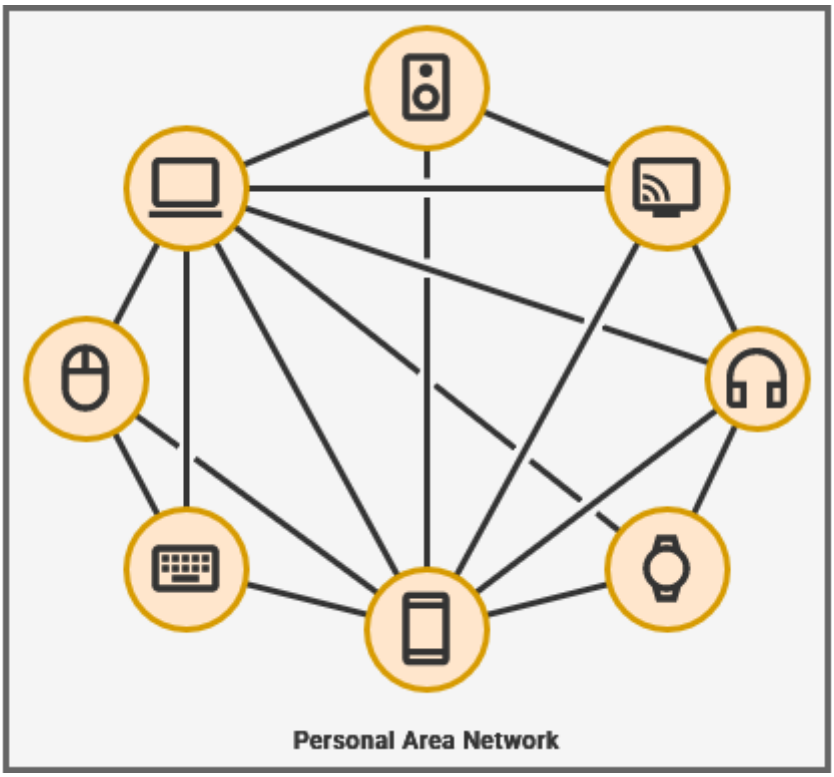

Local Area Network

A Local Area Network contains all of the nodes and links within a limited (often architectural or regional) area. This includes desktop, television and console devices attached by cable, as well as other devices connected wirelessly.

This could be as small as your home or some larger contained area – like a college campus or corporate headquarters. These institutions must subscribe to Internet service – just on a larger scale. They may hundreds of interconnected wireless routers blanketing a mesh network over a large physical area.

Universities can have several campuses and corporations may have branch offices at different scales. A Virtual Private Network creates a private tunnel connecting two geographically separated LANs into one that is accessible by both locations. This enables devices over vast distances to communicate as if they were nearby. This can be accomplished invisibly through hardwired infrastructure, as well as on a device-by-device basis by connecting to a VPN server using the appropriate credentials.

When connecting computers to a wired network, there are a few devices that can incorporate physical cables to facilitate links between nodes. Judging solely by appearance, it can hard to tell them apart.

Modem

Router

Repeater Hub

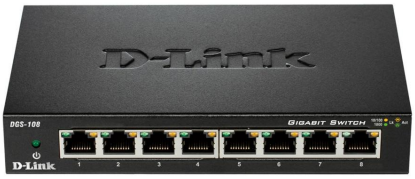

Switching Hub

Each Local Area Network has a modem responsible for transmitting data to and from an Internet Service Provider. This hardware is used to modulate – or translate – data into a signal that can be sent along a physical cable, radio wave or other connection.

The router connects to the modem and orchestrates communication between all the devices connected to it. While connected, each device is assigned a Private IP Address – a unique identification number on that network. This allows devices to quickly and intentionally exchange information over your network, even if there is no access to the outside of World Wide Web.

> Three Private IP address ranges have been reserved for LAN networks: *192.168.68.100, 172.16.0.0,* and *10.0.0.100*.

Ethernet is a standard for enabling network device communication over a wire. Ethernet cables are given a category designation – with higher categories meeting the performance requirements of data centers. Modern routers often incorporate wireless connectivity through the Wi-Fi standard – which turns data in radio waves that can be transmitted to devices through their wireless radio.

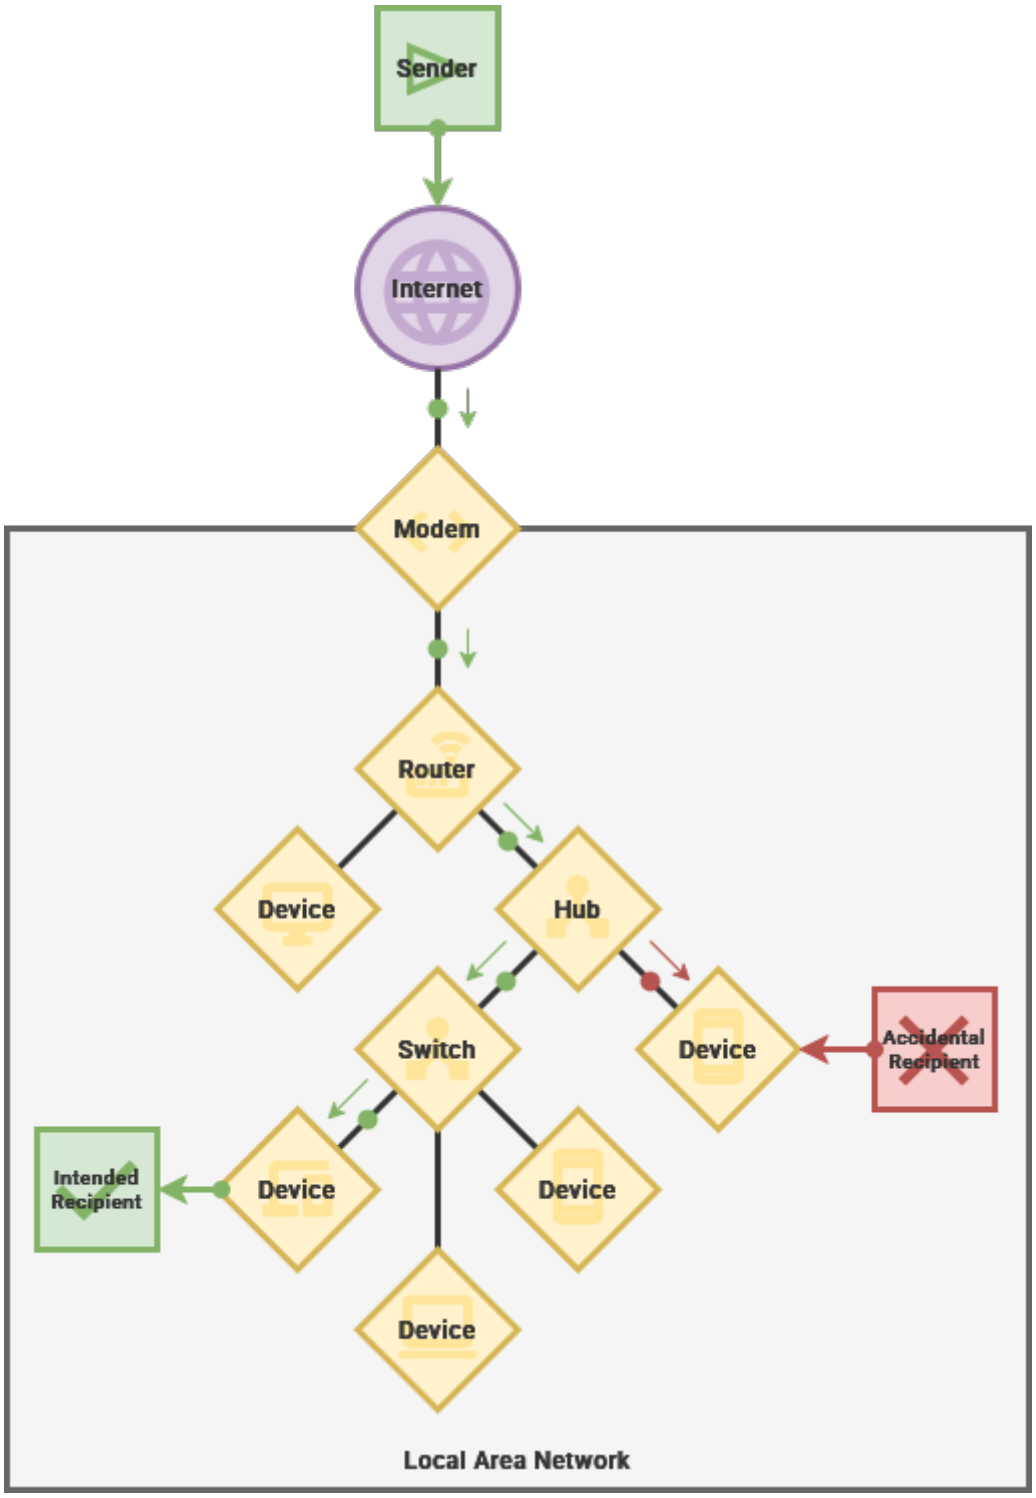

A repeater hub can connect many devices to the network on once, but will openly broadcasts all data it receives to every device connected to it. They can be cost-effective because of their simple design, but they greatly increase the potential for data sniffing – or the data being intercepted by someone other than the intended recipient.

> What could result from the wrong device accidentally receiving it's data?

On the other hand, a switching hub behaves more intelligently by only sending data to it's intended recipient. This requires electronics to process the information being transmitted through it, but

ultimately results in better reliability and security.

Wide Area Network

Your modem, acting as the gateway to the internet, is also assigned a Public IP Address. Similar to a phone number or street address, this is how networks find each other over the vast worldwide internet infrastructure. Whenever a device on your network contacts the World Wide Web, the router on the uses Network Address Translation to automatically convert data packets between your device's Private IP Address and the Public IP Address of your modem. This technique allows your LAN devices to access the Internet without the internet being able to access your LAN devices.

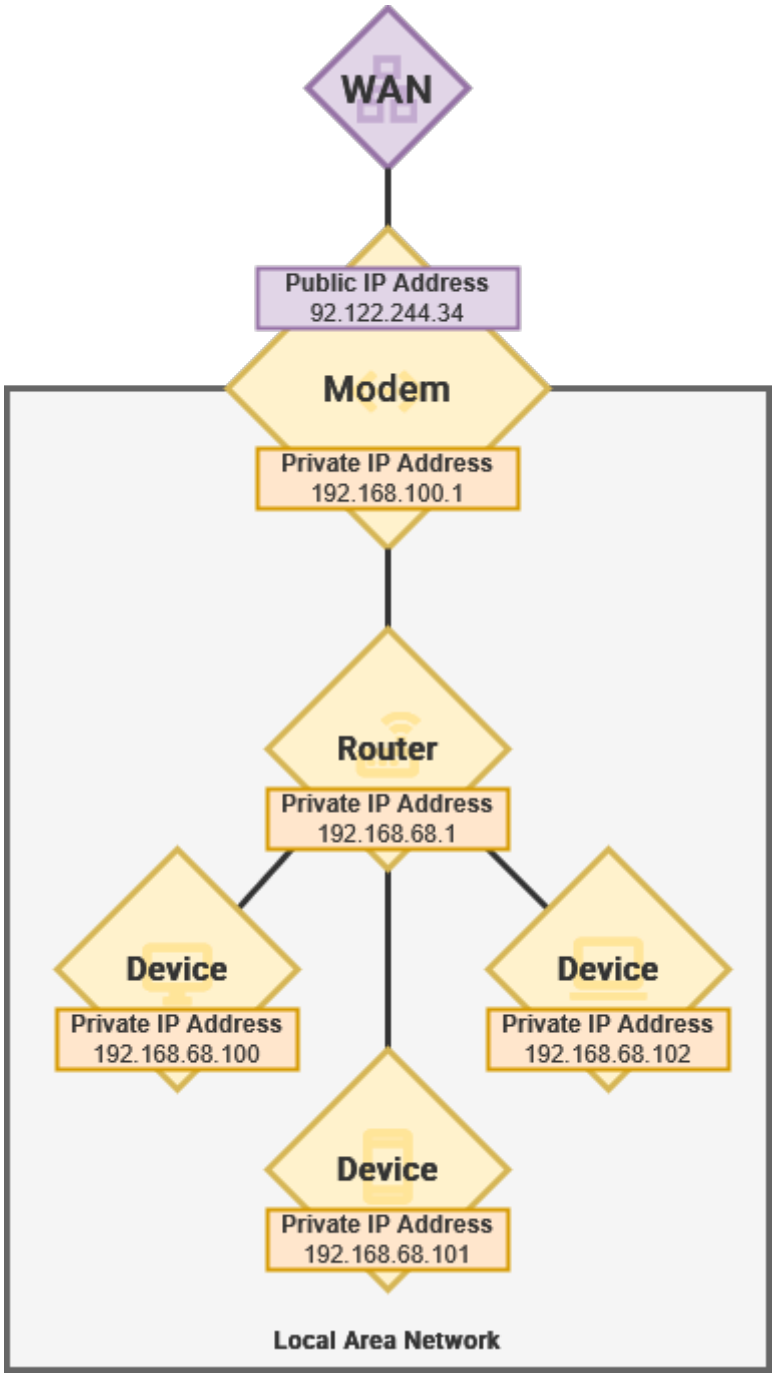

These disparate Local Area Networks – such as your home, your neighbors, city, county and state – are conglomerated together into a Wide Area Network or WAN.

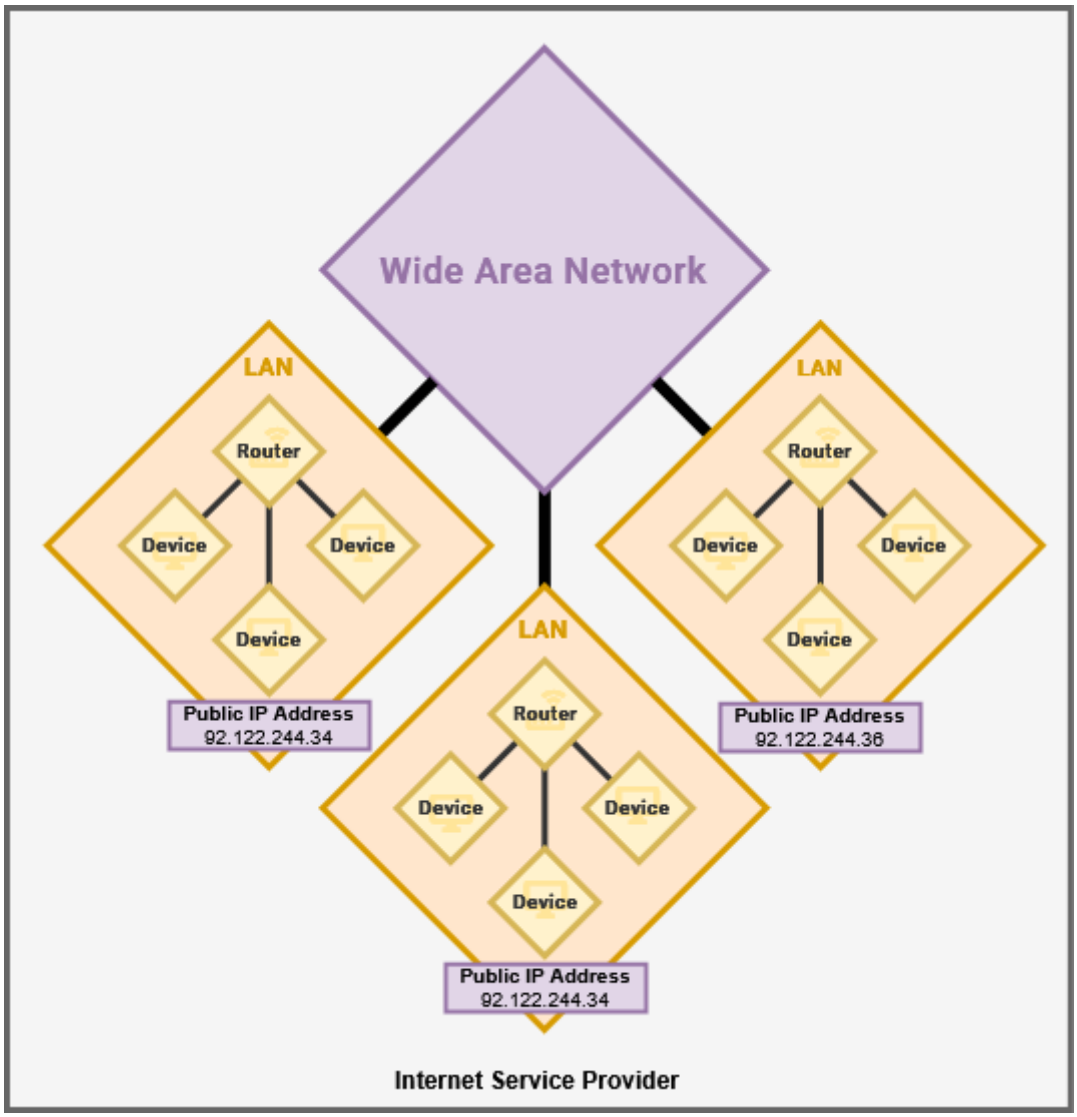

Distributed Applications

The modern internet, as we know it, predominately operates within the client-server model. This means that one computer – a server – is used to respond to the requests of other computers – known as clients. Perhaps the most well known example of the client–server model is the modern World Wide Web.

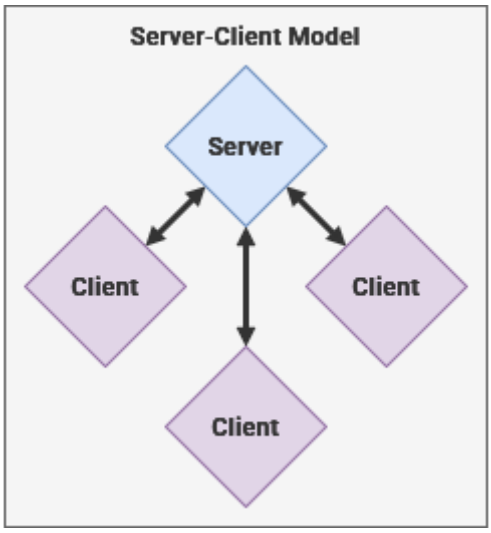

Through a Web browser, we can navigate to a server using a graphical interface and enter a URL – such as example.com. This is more specifically known as a domain name and points towards the address of a Web server on the open internet. By leveraging the HTTP protocol, we can request data from a server and receive the response back in the form of an interactive website.

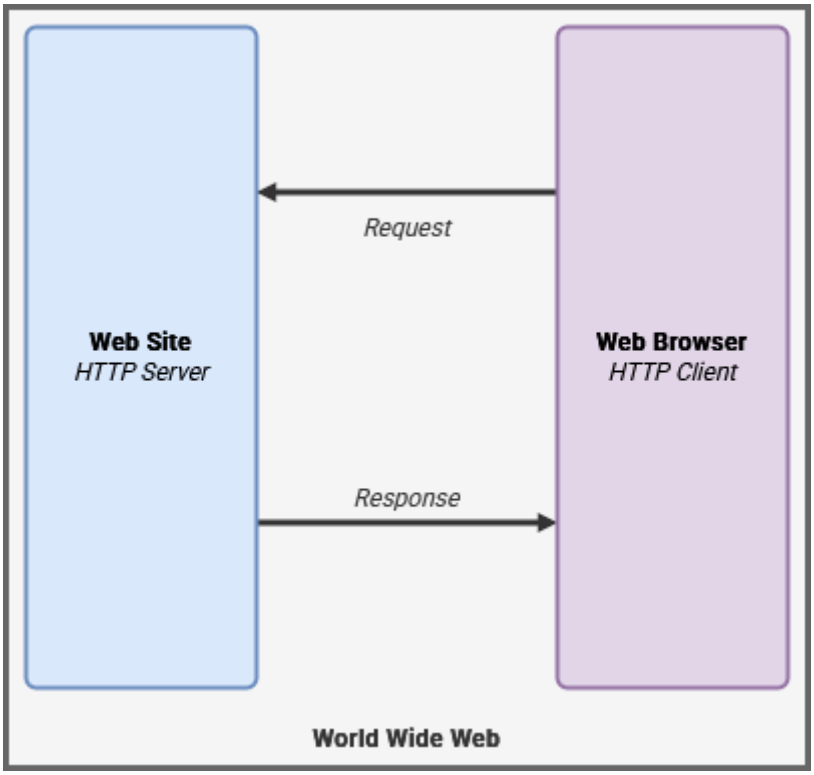

When you enter example.com into the browser's address bar, it needs to be translated into an IP address for our computer to connect to. The Domain Name System enables anyone in the world to know where to locate the web server over the World Wide Web. The predecessor to this infrastructure, which acts as the "phone book" of the entire Internet, was pioneered by Elizabeth Feinler.

> For example.com, the IP is "92.122.244.34".

While still less common for consumers, the peer-to-peer model is becoming more popular. These behave similarly to the mesh networks that allow ISPs to transmit data around the globe through interconnected networks. Within this network structure, each peer has the ability to act as both a server and a client to share data in a more efficient way. Each peer has the same privileges and power, creating decentralized networks – such as BitTorrent, OwnCloud and social media like the Fediverse or Bluesky.

The Path of Data

Data may need to travel vast distances to get from it's origin to final destination. This can include multiple internet service providers and connection types – ranging from physical cables to wireless connections. Carrying data around the globe can include anything from vast underwater cable networks to satellite relays in geostationary orbit.

Whenever data is transmitted over a network, it is first broken into small "packets" by your computer. These are transmitted to the network router and across the infrastructure laid by your Internet Service Provider to a local hub and, possibly, a regional or central hub.

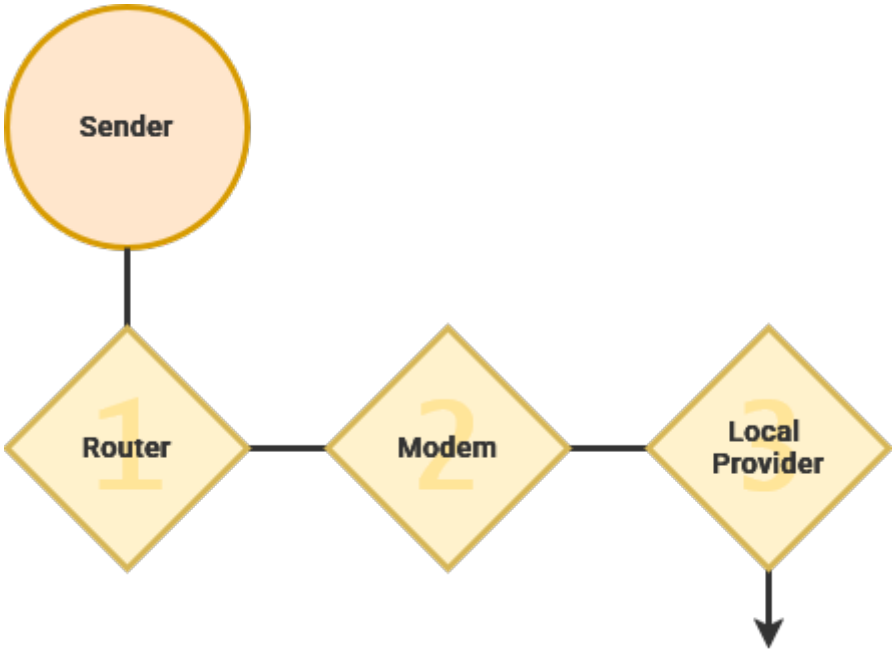

Your Provider contracts access to middle-mile and long-haul providers specializing in quickly transmitting data across a geographical region. These carry data from outlying areas into major metropolitan areas where, if necessary, it can be sent around the globe. The Eastern US shares many undersea cable connections with Western Europe, just like the Western US connects to China and Japan.

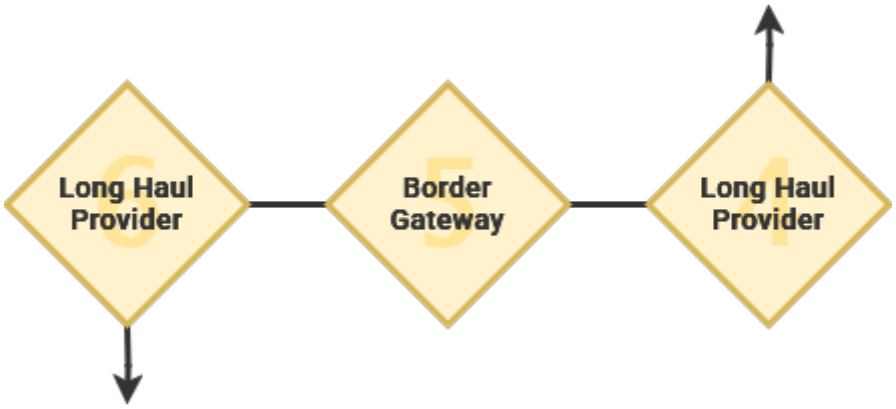

Through the Border Gateway Protocol, data can find a route across this patchwork of autonomous and independently-owned network systems. This process relies on the mutual agreement between ISPs that every network system will act as a neutral peer to all other networks ensuring that messages will always be passed along towards its destination. If these data packets contain any erroneous or fabricated metadata, they will likely get lost during this exchange process.

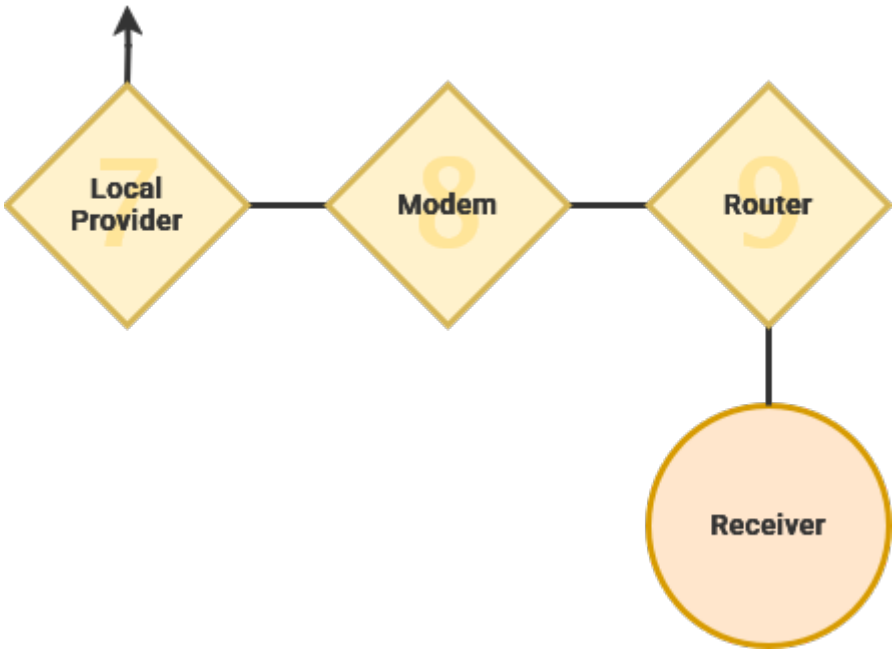

From here, packets will take the most direct route to it's destination. Performing the same process in reverse, data transits through middle-mile and long- haul providers, before filtering through regional and local internet infrastructures. Finally, the data enters the intended router before being delivered to its destination.

When the receiver wants to send a response back to the original sender, it must repeat this entire process again. Modern software systems often implement mechanisms that will remember the quickest connection between two points. Web browsers, for example, are built on top of open protocols that enable two computers to create a persistent connection that can be reused for transporting data.

Open Standards

By the 1960s, computers had begun gaining traction and the US Department of Defense created the first Wide Area Network by connecting research campuses across the US with those in the United Kingdom and Norway. Before the concept of the World Wide Web, this network was used to perform [illegible] digitally isolated computer systems.

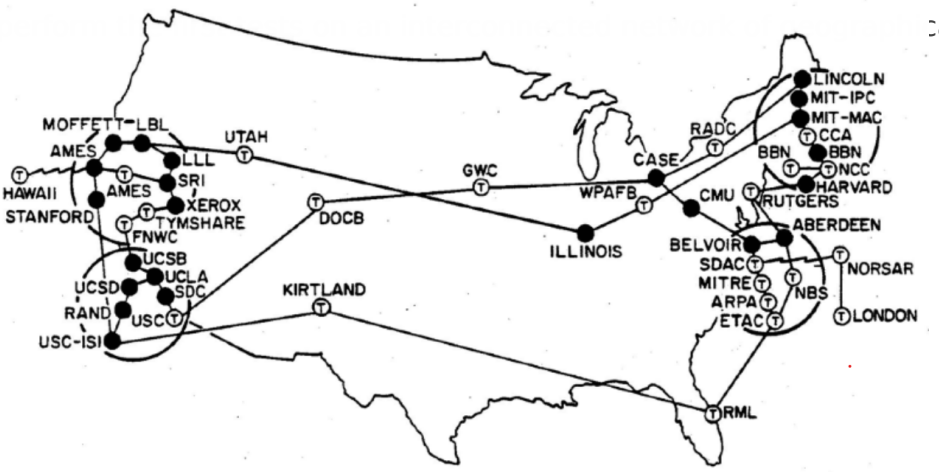

We quickly realized that we would need to create tools, protocol and standards to ensure that these networking systems didn't become too fragmented. for us to communicate amidst rapid technological advancement. By the late 80s, globally expanding internet networks were increasingly relying on the TCP/IP protocol standards – more commonly known as the Internet Protocol Suite.

Simplified, this protocol standardized access to the Internet by creating functional layers that depend on each other while operating in parallel and communicating through a common language. The five layers of the Internet Protocol Suite are:

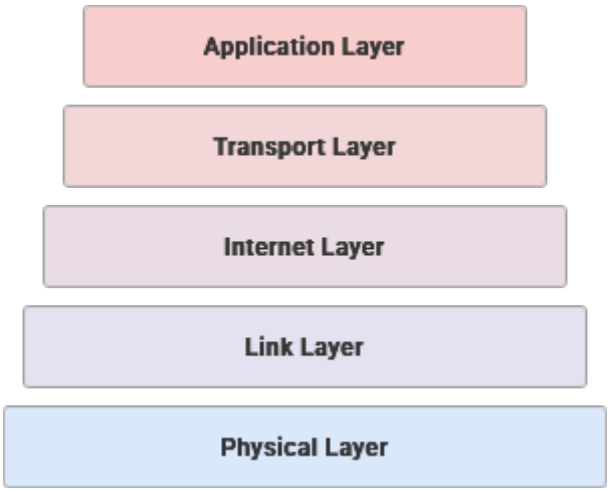

Each layer defines the requirements and expectations ensuring that disparate devices could mutually agree on a language to communicate with. At each level, we are ensuring that all nodes on the network are speaking a common language. By being able to effectively communicate, these layers can operate independently and only communicate tasks to the next layer as needed.

memory

Physical Layer

This level handles how a data signal is encoded and transmitted between hardware components using a range of connection types – such as electrical, optical or wireless links. Within this layer, each individual network-connected device is identified by a unique MAC address – with over 281 trillion possibilities.

Cable

Link Layer

This level defines the functional and procedural methods that are used to transmit data between nodes across their link. Within this layer, we are restricted to connections between nodes that are connected physically – such as through Ethernet and Wi-Fi. When necessary, a Virtual LAN can be used to segment a monolithic physical network into smaller virtual ones that operate in isolation.

lan

Internet Layer

This level covers the methods and specifications for transmitting data between intermediate routers along its path from origin to destination. When transmitting data across the world wide web, public IP addresses are used to route a path.

swap_horiz

Transport Layer

This level is responsible for entering that nodes can reliably communicate across this vast interconnected network. This is accomplished through an array of protocols that handle the security, reliability and flow logistics of data differently. Within this layer, any data to be sent over the network must be broken into smaller segments to facilitate smooth transmission.

api

Application Layer

This is the top-most layer of the Internet where familiar services exist – such as the World Wide Web, e-mail, remote desktop connections and file transfers. This is the layer that people interact with most often and data transfer happens almost invisibly.

While passing between layers, "metadata" encapsulates the data – or "payload" – being sent over the Internet. This process assists navigation across the globe and ensures that it arrives in one piece without errors.

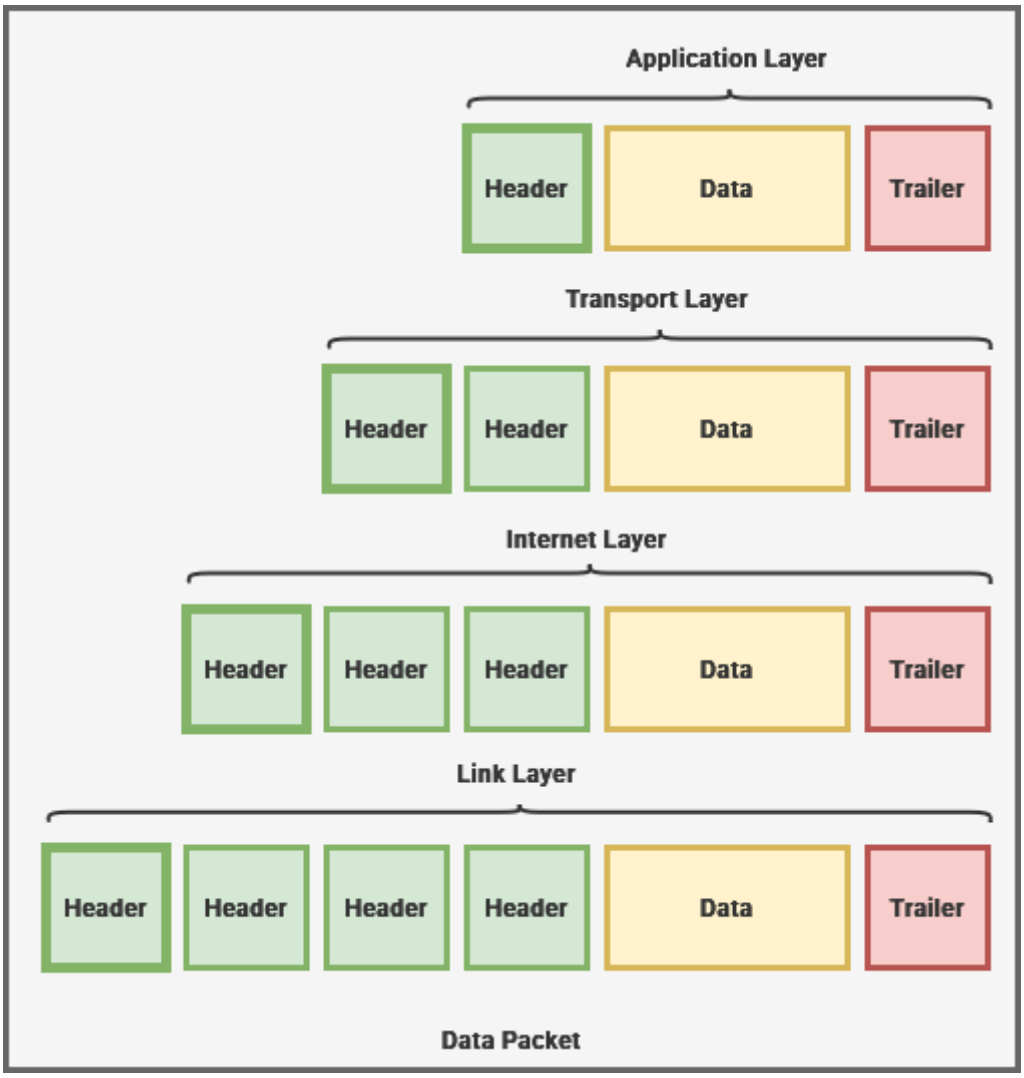

These "headers" detail how this data relates with other data, as well as the path taken and route still in progress. The "footer" will often contain a checksum – or abstracted data based on a mathematical equation – that can be compared to the received data and verify there aren't any errors.

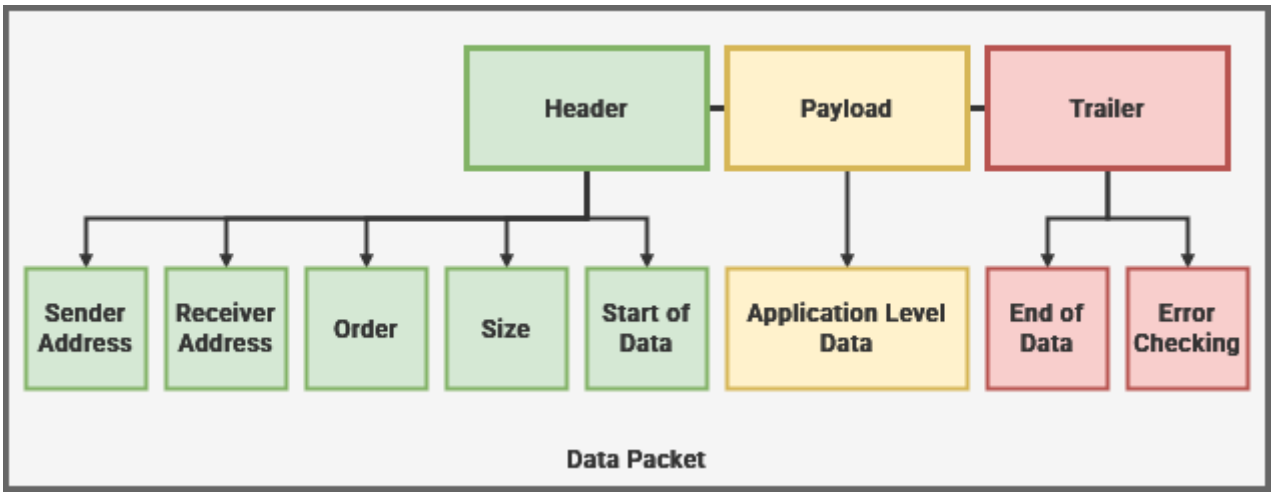

Before data can be sent anywhere, the applications being used – such as between a web browser and web server – must agree on a method of communication. While passing down from the top-most Application Layer to the Transmission layer, software developers can choose different protocols that define how the data will be transmitted over the open internet. During this process, data gets broken down into thousands of tiny "packets" that will need to be put back together on the other end.

Created alongside the internet in the early 60s, Transmission Control Protocol – more commonly known as TCP – is still one of the most widely used protocols. This protocol was built to ensure that all data will arrive without errors and in the exact order it was sent – all important properties for smooth web surfing, resilient email services and reliable data transfers. When using TCP, both devices must agree to "handshake" – or mutually agree on how they will be creating a link between them.

After the connection is established, each minuscule packet of information will be sent one after the other, ensuring each packet was received correctly before moving on to the next. When data becomes lost or corrupted, the sender is alerted so that it can be sent again. This process is extremely reliable, but can be slow because it requires more upfront coordination and processing power.

> On the World Wide Web, we depend on TLS – or Transport Layer Security – to serve us private and secure websites using HTTPS through a TCP connection.

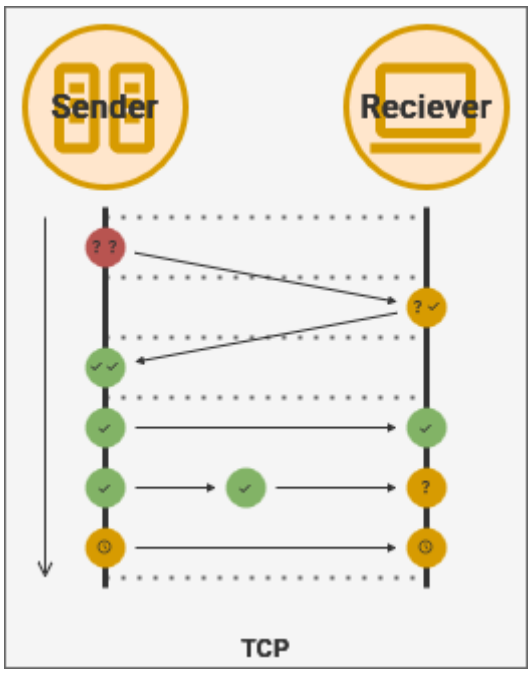

By contrast, User Datagram Protocol – or simply UDP – does not require a "handshake" to create an active connection. Instead, data packets can be sent directly to the server from a client device without needing explicit approval first. This requires that the server is always available to accept incoming connections without prior warning.

The server will respond to this request by also sending data without negotiation. Some applications – such as media services and video games – may begin to stream data by broadcasting packets in quick succession. Using this protocol, there is no guarantee that the packet will be received in any particular order or timeframe. Even when a server sends packets in order, they may be received out of order because of the underlying network.

Neither the server nor the client know when to expect data and, by extension, neither know when data intended for them never arrives. While the server and client will be alerted if the data was corrupted along the way, the application needs to send it again. This protocol is preferable for time-sensitive applications, online video games and media streaming services.

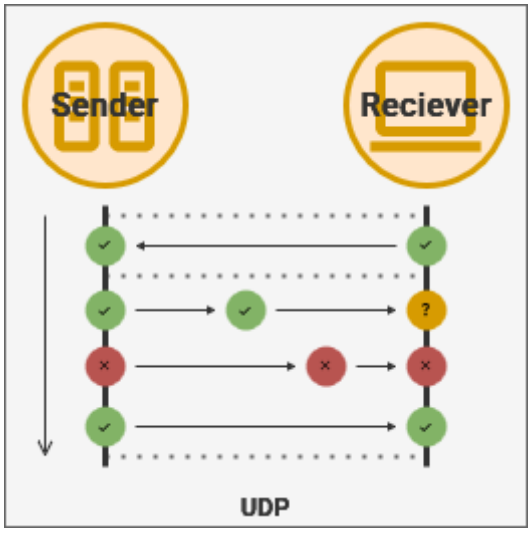

When hosting digital services for a network, ports enable a single physical server to create a dedicated sub-addresses for multiple running multiple applications. This allows each port to use different transmission protocols – such as TCP and UDP. Unlike a physical hardware port, network ports reside entirely in the virtual space and only exist within your operating system while it's loaded.

> Ports use a numeric identifier ranging from 0 to 65535.

If an IP address were to a building address, ports would be apartment numbers within that building. After receiving internet traffic, a computer will route data one final time to the applicable port. Services we use everyday are tied to a specific port on our computer. Often, these are defined through open standards that purposely reserve ports for specific services to reduce confusion.

Port Number	Protocol	Use
53	DNS	Web Browsing
80	HTTP	Web Browsing
443	HTTPS	Web Browsing
20	FTP	File Transfers
22	SSH	Remote Access

Port Number	Protocol	Use
25	SMTP	E-Mail
110	POP3	E-Mail
220	IMAP	E-Mail

How to Remotely Connect

When it comes to connecting to your services while away from home, there are two common methods to approach this: through connecting to a self-hosted Virtual Private Network or broadcasting your services to the World Wide Web. These techniques can be used individually or combined to create a tailored experience.

Virtual Private Network

Just like a corporation or university, you can self-host a Virtual Private Network server from home. This enables anyone with the proper credentials to securely connect individual devices to your Local Area Network. This way, your services can be available to you, friends and family – even while away from your home – without making them accessible to the public internet.

This is clearly not the ideal for hosting a WordPress blog or Flarum forum intended for an online audience. VPN access can be the perfect balance of security and convenience for a small or exclusive audience – such as Bookstack for a tabletop roleplaying campaign. Some services may require using a web domain to properly function, but access can still be restricted to access from your LAN.

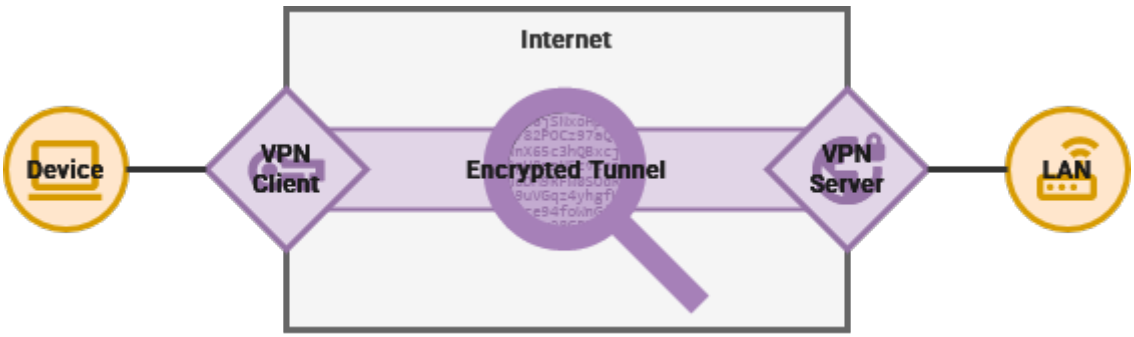

For accessing services that handle private personal information such as Actual Budget or Paperless-ngx, this is the most secure option. By requiring authorization to remotely access your Network, you can greatly decrease your attack surface – or the amount of publically-accessible software that may contain software vulnerabilities that can be leveraged by malicious actors.

> These type of software vulnerabilities are commonly called zero-day exploits because they are either unknown or unfixed.

While open-source software [can improve security by putting more eyes on potential vulnerabilities](), it does not mean there will not be breaches. Software projects are written by developers with varying priorities, including security and privacy. You do not need to be as concerned about the security of individual software programs when everything is protected behind a singular VPN program.

Comparison	
Security	**Security** starsstarsstars By requiring authentication before even connecting to any services, you can greatly decrease your overall attack surface.
Shield_with_heart	**Convenience** **stars**circlecircle This will need to be configured on a device-by-device basis. Once the service has been setup, you just need to make sure you stay connected.

Virtual Private Networ

Web Domain Name

Self-hosting a web domain involves connecting your server to the World Wide Web. We accomplish this by linking the Public IP address assigned by your ISP to a domain name you control. This adds your public IP address to the [Domain Name System]() registry that helps web servers locate each other.

> Web domains – such as example.com – are hierarchical with deeper levels appended to the front. The URL above contains a top-level domain ("*com*") and a second-level domain ("*example*"), joined by a period. When you own a domain name, you can create additional sub-domains – like app.example.com.

Broadcasting your server on the World Wide Web makes it extremely simple to access your services from anywhere in the world using only a web browser. This is equally true for every person in the world who has access to the World Wide Web. At the end of the day, we are opening our server to the whims of the open internet – and any potential malicious actors.

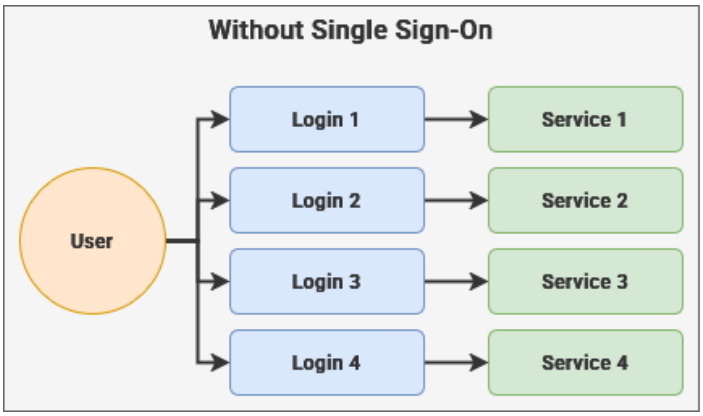

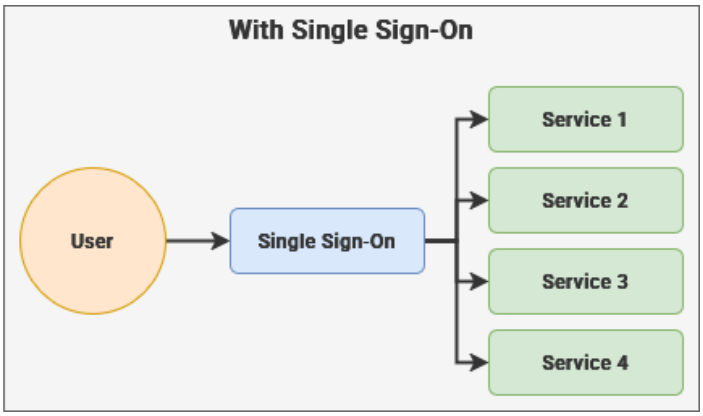

We will take proactive steps to harden security, preempt vulnerabilities and limit fallout. SWAG makes it simple to setup secure encryption for our web domain. Authelia is a single sign-on service that can decrease your attack surface by protecting your individual services with the same trusted authentication system. Fail2Ban and CrowdSec are open-source solutions for automatically identifying and intercepting malicious actors.

Comparison	
Security	**Security** **stars**circlecircle While you can take proactive steps to protect your data, it is still connected to the open internet.
Shield_with_heart	**Convenience** starsstarsstars Once setup, you can access your services anywhere in the world with just a web browser.

Combination

You can tailor your Web server as needed to find your preferred balance between security and convenience. We can leverage the convenience and memorability of web domain names while still retaining the security of a Virtual Private Network. This enables websites to be easily accessible while still denying access to anyone outside of our Local Area Network.

We can provide access to Cockpit at cockpit.example.com, but deny access to anyone attempting to access it from outside your Wi-Fi or Ethernet network. At the same time, we can provide public access to a personal WordPress blog. When combined with a VPN, you can still provide secure remote access to private data and services.

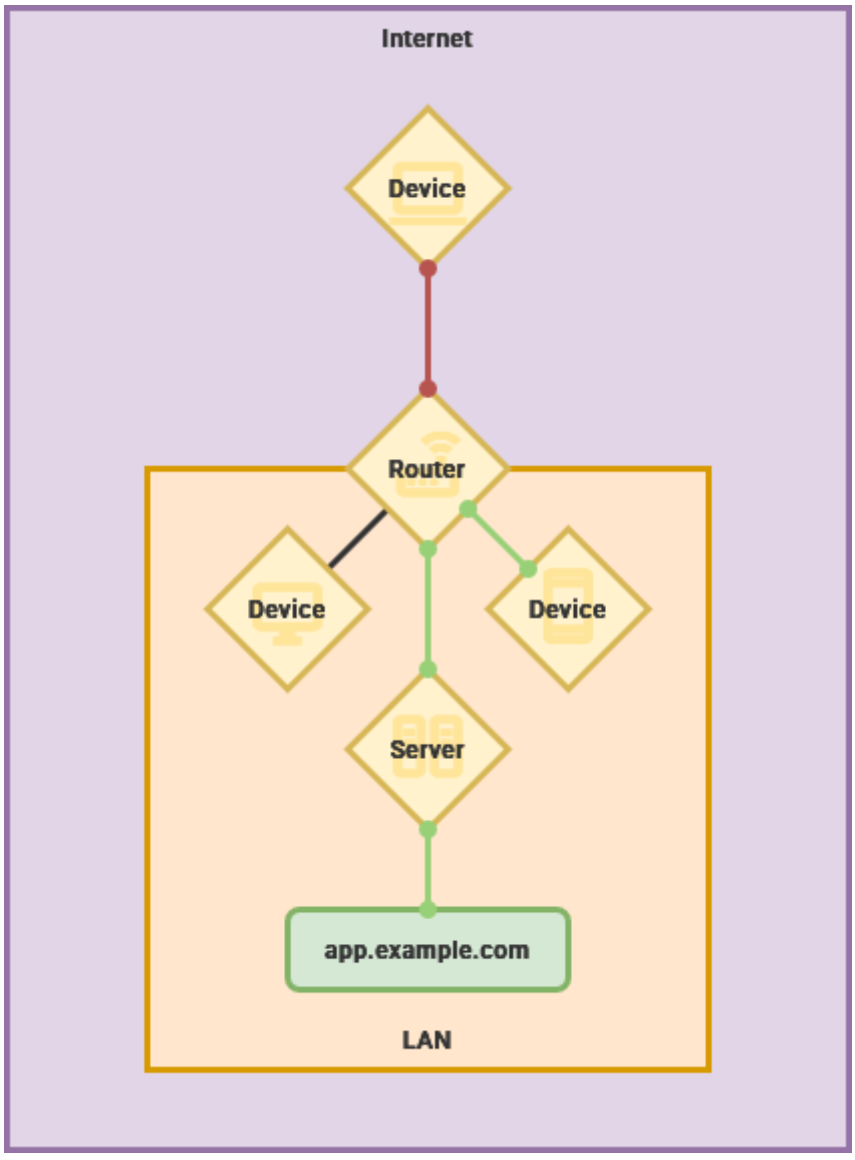

diagram showing inside and outside access to a local restricted address.

Comparison	
Security	**Security** starsstarscircle By requiring local access for critical services and leveraging single sign-on, you can have the best of both worlds.
Shield_with_heart	**Convenience** starsstarscircle Your Web server will always be accessible through a browser, with device-by-device setup required for accessing critical services.

LAN-Only Access **k**

Virtual Private Network

By hosting your own VPN server, your devices can remotely connect to your Local Area Network as if they were connected to your Wi-Fi.

Hosting a VPN Server

Virtual Private Network

Connecting to Your VPN

Accessing Your Services

Web Domain Name

After purchasing a domain name, we can configure SWAG to generate a secure certificate that ensures our server's security and authenticity.

Most people will only need a single domain address with sub-domains used for each of our web services. If you are trying to run multiple different brands, you can also host multiple domains from the same web server.

Web Domain Name

Domains & URLs

When navigating the World Wide Web through a browser, we enter an address in the form of a URL - or Universal Resource Locator – that contain the information necessary for your computer to connect to another located anywhere in the globe. Each URL has a few core components that help specify a location and how to reach it.

> A full web domain URL might be https://www.example.com.

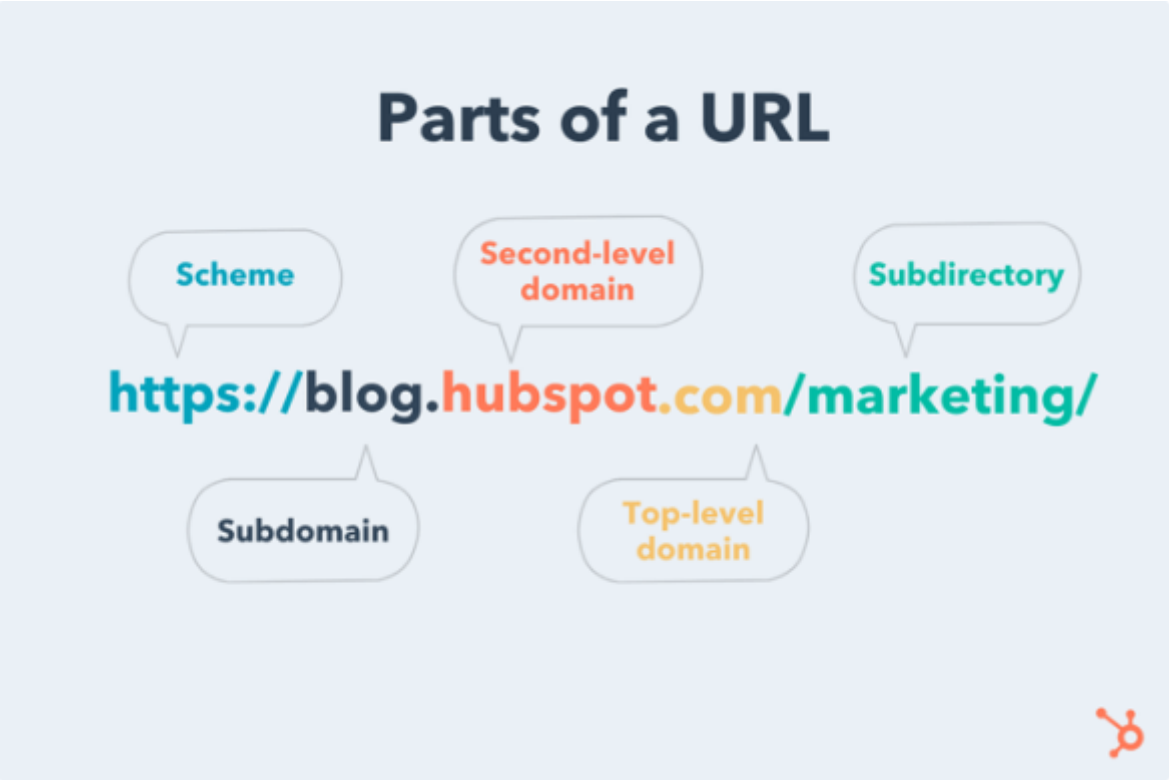

Domain names are used to identify autonomous areas of the World Wide Web under the control of a person, corporation or other entity. These are used for gaining access to websites, email and other services. Web domains are listed on the Domain Name System – a distributed registry to help locate servers around the globe.

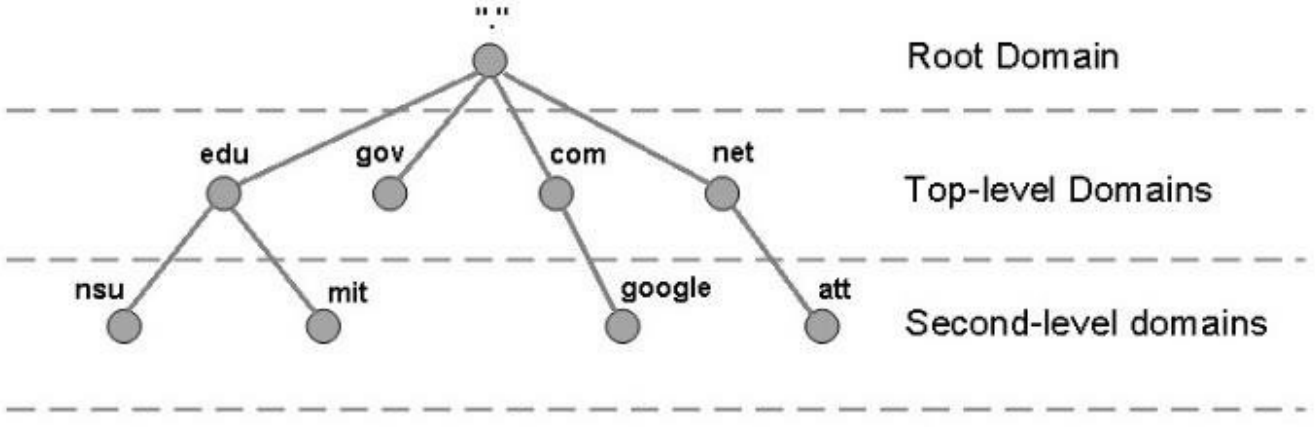

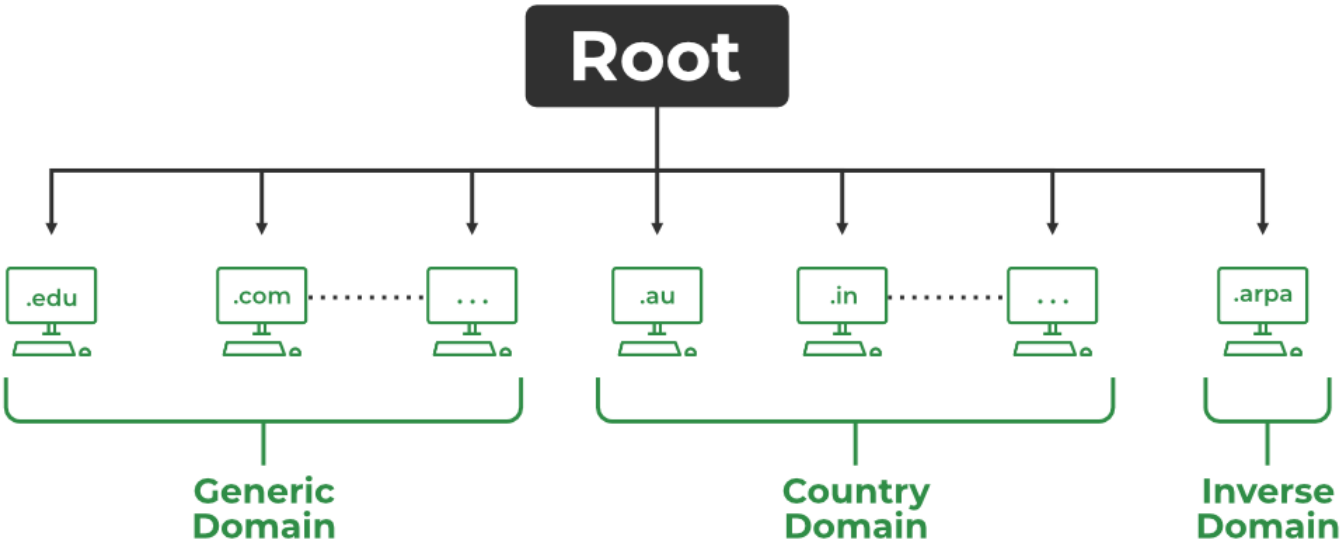

Much like the file system on the Linux operating system, Web addresses start from a root that contains all domains. Stemming from the root, there are top-level domains like the popular .com, .info, .net, .edu, and .org. There are also special country code web domains, as well as the recently expanded purpose-specific domains – such as .news, .tech, .name or .blue.

These "TLDs" are carefully maintained by the Internet Corporation for Assigned Names and Numbers (ICANN) – a global non-profit that handles the development of policies and procedures for the Domain Name System.

Second-level domains offer personality and customization to a URL. They are combined with a top-level domain using a period to create a full web domain name. When purchasing a domain name, you are generally renting access to this specific combination of top- and second-level domain for a pre-determined period of time.

> Domain registrars often provide discounts for the first year, but following years come with higher annual renewal fees.

Single Domain

For many people, you will probably only require one domain name. We will be using SWAG to connect our server to the World Wide Web and this project is configured to use a single domain with sub-domains as a default.

Sub-Domains

When offering multiple independent services from your server – such as OwnCloud or Radarr – we can use sub-domains to separate these web applications into different URLs.

> Using this example, radarr.example.com could direct you to the Radarr web application.

These are an excellent way to build a digital community and brand identity around the same domain name. With this technique, we could host a primary website at example.com, as well as a Flarum forum at forum.example.com and a WordPress blog available at blog.example.com.

Multiple Domains

SWAG can configured to act as the access point for multiple different web domain names – such as example.com and example.org. This makes it so you can run a personal and professional web domain from home using the same server.

> Each of these domains will be listed under the same SSL certificate and therefore linked.

Getting a Domain Name

Before you can connect your server to the World Wide Web and access it through a Web browser, you will need a domain name. This address will be used to generate the SSL certificate that web browsers rely on to create a secure network connection and verify we are who we say we are.

Domain Registrars

While ICANN – an American 501(c)3 non-profit located in California – is tasked with the development of the global Internet infrastructure and security, they do not have the capacity to orchestrate over 350 million global web domain names. Instead, over 2,800 domain name registrars have been accredited to operate in their capacity.

This process requires ongoing adherence to strict regulations including additional clearance required to offer special or country-specific domains. These rules, set by ICANN, require that the ownership of domains can be transferred between registrars. They set explicit pricing restrictions on some top-level domains – such as .com and .net – while also imposing a maximum domain registration period of ten years.

While the ICANN has requirements for domain registrars, they are not all created equal. Some employ annoying marketing strategies while others have lax security or even actively malicious practices. Aside from select pricing restrictions that are required for accreditation, domain registrars are free to operate their service however they desire. This includes pricing schemes that reflects expected market desirability and popularity – such as home.tech costing $650,000.

Domain registrars are required to provide information to WHOIS – a public-access database about domains, including contact information for the person who owns the right to it. Registrars commonly offer privacy services that withhold personal identifying information.

> You don't need to pay for an SSL certificate through the domain registrar because we will be generating them for free using SWAG.

Cloudflare

Based in America, this company provides cybersecurity services to nearly 20% of websites, including a free consumer tier. They are also an accredited registrar that provides at-cost domain name services for most top-level domains.

Micro.Domains

This service leverages Namecheap's infrastructure to sell domain names that are 5-characters or less. While these domain names are often random, they offer marginal 'security by obscurity' for accessing personal web services at a reasonable and transparent price.

PorkBun

This service is operated by the American business Top Level Design, offering an intuitive experience and transparent rates for domains down to $2.

NameSilo

This service is operated out of America and focuses on customer privacy. They offer transparent pricing and accept many different payment options.

Dynamic DNS

Most consumer Internet Service Providers assign Public IP addresses to residential networks that can change at any time. Proprietary Dynamic DNS services offer tools for consumers to quickly update their Public IP address and make sure their server is always accessible. There are several free services available, but they come with drawbacks.

> SWAG can generate a working SSL certificate to ensure data privacy, but browsers may not be able to verify this certificate and can lead to browser-based warnings.

DuckDNS

This free service offers a free sub-domain under duckdns.org – such as example.duckdns.org. They use Amazon servers, but can be easily integrated into your Web server setup.

DDclient

This open-source utility makes it simple to keep one or more Dynamic DNS services up-to-date with your current Public IP address. This will require more initial setup, but it can create redundancy through multiple services. DDclient currently supports over 30 different dynamic DNS services. During the creation of an SSL certificate, you will need to manually validate your server using the HTTP mechanism provided by SWAG.

Manufacturer DDNS

Consumer routers from mainstream technology manufacturers – such as ASUS, TP-Link and Netgear – are providing built-in Dynamic DNS features. This will provide a sub-domain under their dynamic domain service. During the creation of an SSL certificate, you will need to manually validate your server using the HTTP mechanism provided by SWAG.

Domain Name System

We need to configure a DNS service to handle the translation of our domain name into the Public IP address provided by your Internet Service Provider. There are three ways to accomplish this: one requiring a paid domain name, while the two others are free subscription services.

Using Cloudflare

> This is the recommended method.

Cloudflare is a content delivery and cybersecurity services company that offers free basic-tier solutions for anyone running a web service. We will be using their DNS service and nameservers to direct traffic to our server. They offer protection from Distributed Denial of Service – more commonly known as DDoS – attacks as well as defensive tools in the event that you are targeted. They will provide another layer of security to our web services.

Nameservers

When purchasing a domain name through domain registrar, they will generally to use that registrar's web services. This includes an authoritative nameserver that serves as a directory of domain names they are providing service for. The Domain Name System powering the internet is decentralized and no one entity owns it explicitly. A recursive resolver is used to systematically search these disparate nameservers and find the desired domain.

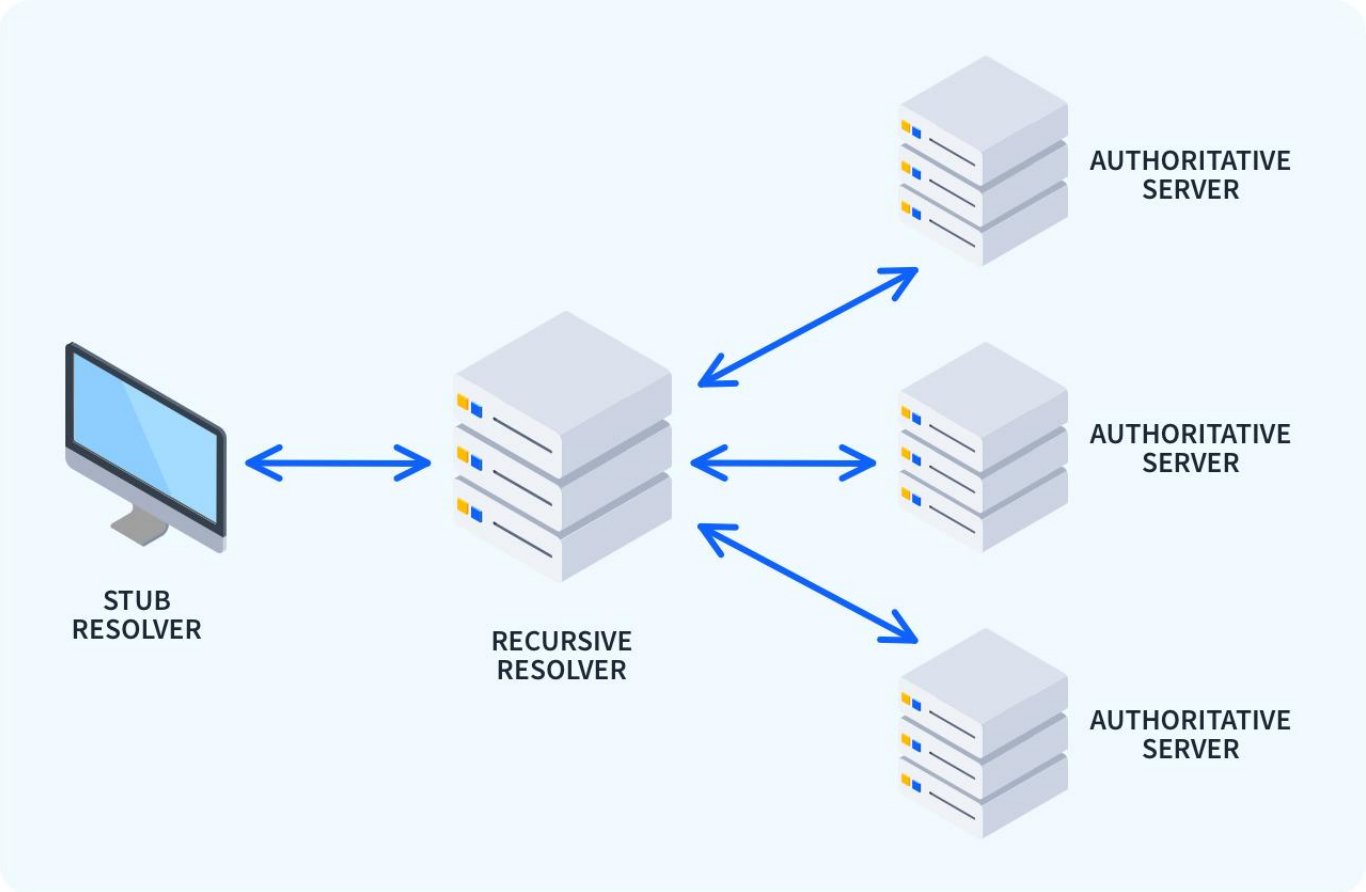

We will need to configure our domain name provider to use the CloudFlare nameservers. This will enable us to leverage their free services. The process to configure your domain's nameserver will be different based on the registrar you used. We use redundant servers to ensure that at least one is always available even if their are outages.

> When purchasing a domain through Cloudflare, they are pre-configured to utilize their nameserver and security services.

These can provide insights for a select few domain name registrars:

- NameSilo
- Porkbun
- NameCheap

We will need to use the Cloudflare nameservers to leverage their services. They host an assortment of decentralized nameservers to split up the workload.

> You will need to create a Cloudflare account. If you want privacy and anonymity, ProtonMail allows you to create separate email aliases.

When creating an account through Cloudflare, we will first need to add our site to their dashboard and then they will assign you two nameservers.

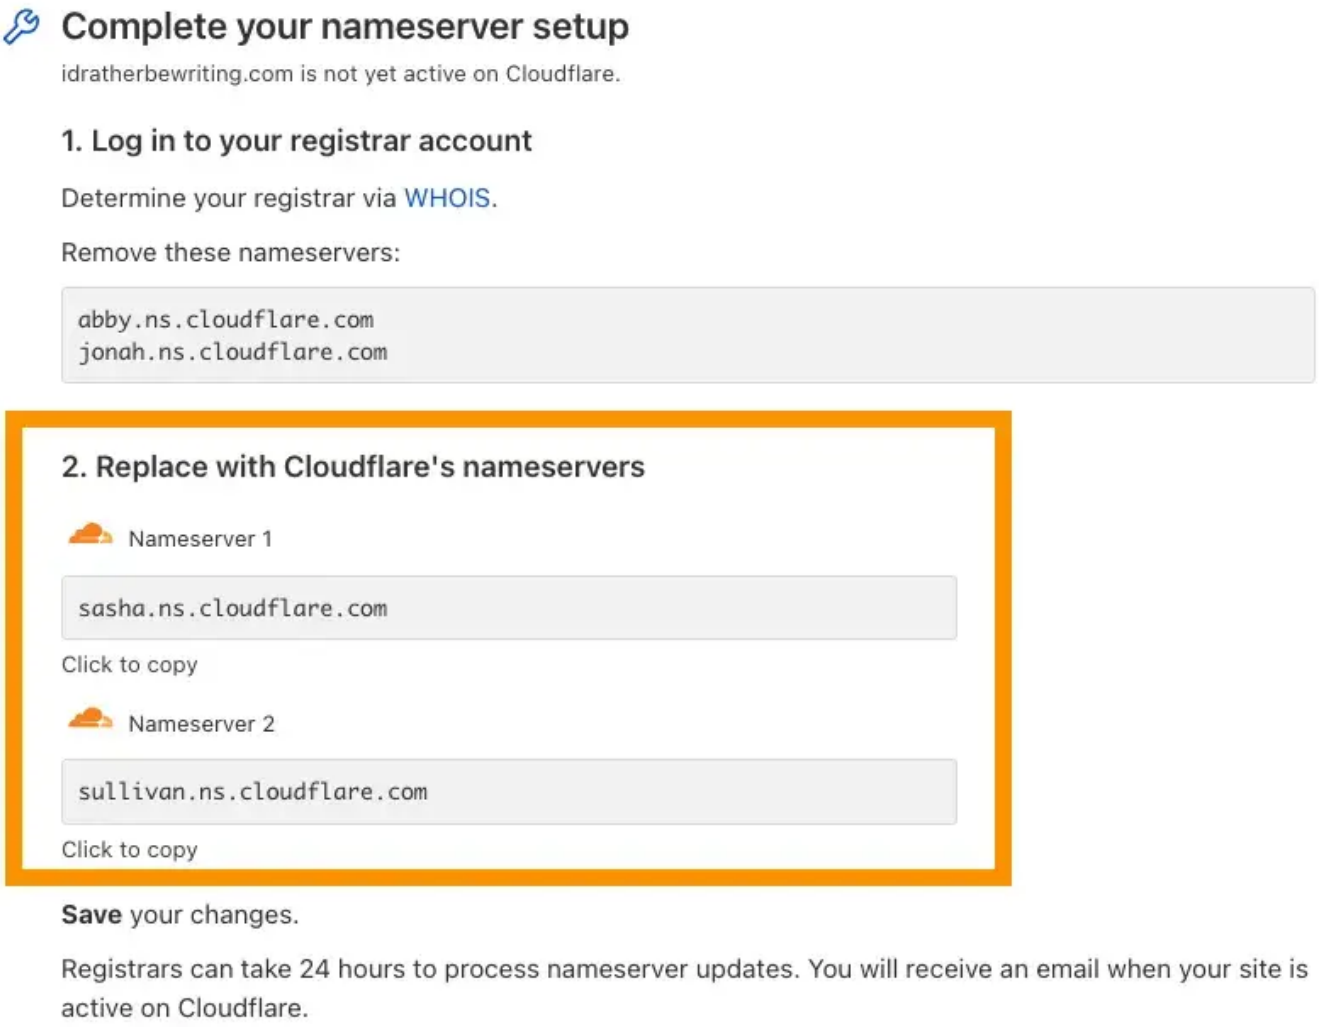

> Nameservers generally update quickly (every ~15 minutes) but it may take up to 24 hours.

You will receive an email notification once this process is complete. After you have added their nameservers through your domain name registrar, you can complete the Cloudflare domain verification.

> For enhanced security, you should follow the Cloudflare guide for enabling DNSSEC.

DNS Records

Once our domain is configured to use Cloudflare's nameserver, we will need to configure traffic received at our domain to be directed to the server located at our Public IP address. DNS records – much like a label on a filing cabinet folder – explains what can be found within. These can also be used to configure email addresses, social media handles or even store public notes.

If you pay your Internet Service Provider for a static Public IP address, you can direct the domain to your server and you'd be done. We need to create an A Record to direct traffic to our IP address using the Cloudflare dashboard.

This record should have the name '@' to signify we are setting the IP address for the root of our server – such as example.com – as opposed to a sub-domain. For the IP address, we need to add the Public IP address provided by your Internet Service Provider. If you are unsure, you can view your public IPv4 address by visiting a web service like What Is My IP?.

Dynamic Addresses

Most residential Internet service plans do not come with a Static Public IP address by default. This is generally restricted to commercial business internet plans for an additional fee. Home Internet connections generally use dynamic IP addresses that may change at any time.

There are open-source software options to automatically update our IP address within DNS records. LinuxServer.io maintains a Docker image for ddclient which can connect to CloudFlare through their API to ensure the IP address is always accurate. This will require creating an API key for Cloudflare, installing ddclient and configuring it with a text editor.

Install ddclient keyboard_arrow_right

Using DuckDNS

DuckDNS is a free and proprietary service where you can reserve a sub-domain – such as example.duckdns.org. This domain can be directed to your server and configured to work with your individual services. You will need to make an account with the service by logging in using to Google or GitHub through their homepage.

Once you have logged in, you can register your sub-domain through the service and assign your IP address to it. On your account page, there is a private token that can be used to automatically update your IP using a DuckDNS Docker image.

Install DuckDNS keyboard_arrow_right

Manufacturer Dynamic DNS

Consumer routers from mainstream technology manufacturers – such as ASUS, TP-Link and Netgear – are providing built-in Dynamic DNS features. By creating an account with them, you can access a sub-domain under their dynamic domain service. this will happen automatically without intervention.

Web Domain Name

Reverse Proxy

We need to install a reverse proxy to safely access our web server over the World Wide Web. This specialized server software sits in front of other servers and retrieves websites on behalf of the people trying to access them. They act as the public-facing front for a hidden network of computer servers operating behind-the-scenes on your local area network.

When passing along web site data, a reverse proxy will actively overwrite any information about the server it came from. During this process, they can also modify the HTTP headers used to silently communicate information between a web server and browser. This also can be used to inject information into the stream, changing the way a website looks and operates from the top-down.

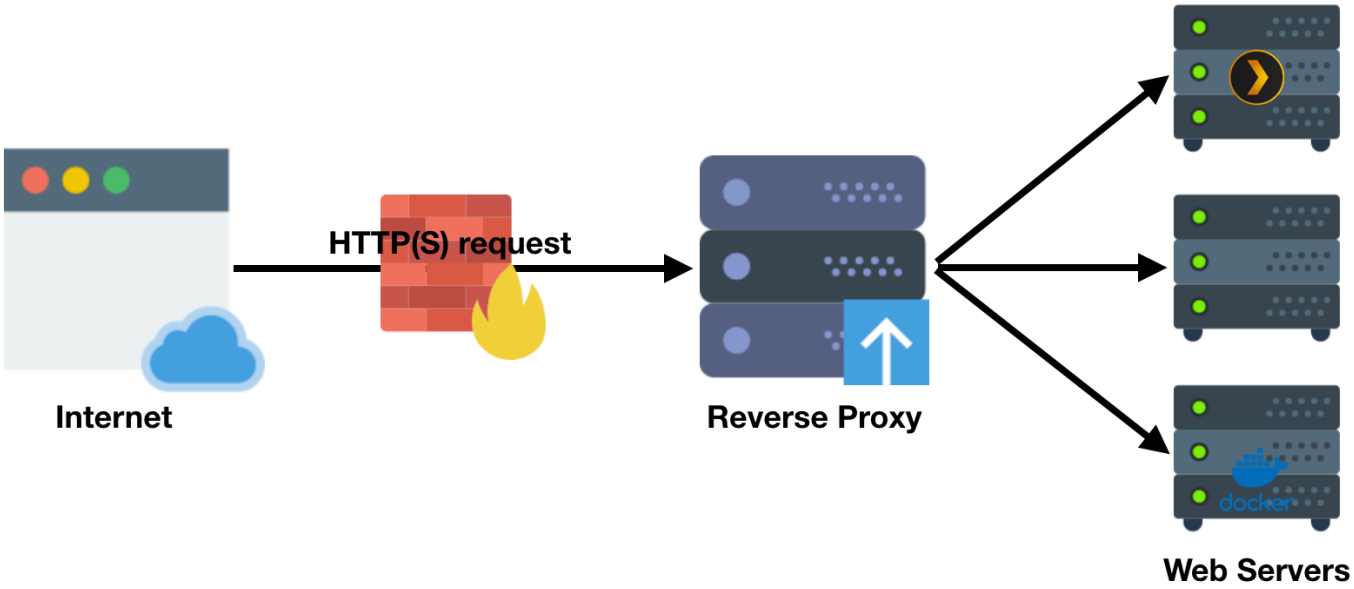

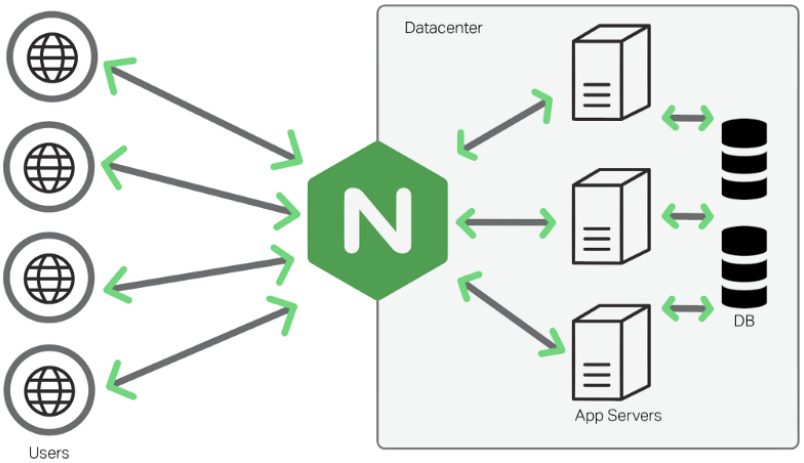

This lets us keep access to our services behind the privacy of our Local Area Network and ensures no one can ever communicate directly with an origin server. This decreases our attack surface – and improves overall security – by only giving the Internet access to the reverse proxy instead of by individual service. In the event that our server is hacked, they only gain access to the reverse proxy and not any of our underlying services. However, a compromised reverse proxy has the potential to cause a great deal of damage and should be safe-guarded.

{{Show difference between reverse proxy and connecting individual services to the internet.}}

For cloud computing environments, reverse proxies provide powerful speed benefits on multiple fronts. Acting as a cache – or a copy of frequently requested data – a reverse proxy can step in and take the strain off of individual services by handling simple requests. They can also work as a load balancer to spread out users across multiple independent servers. This is how large cloud websites provide access to millions of users.

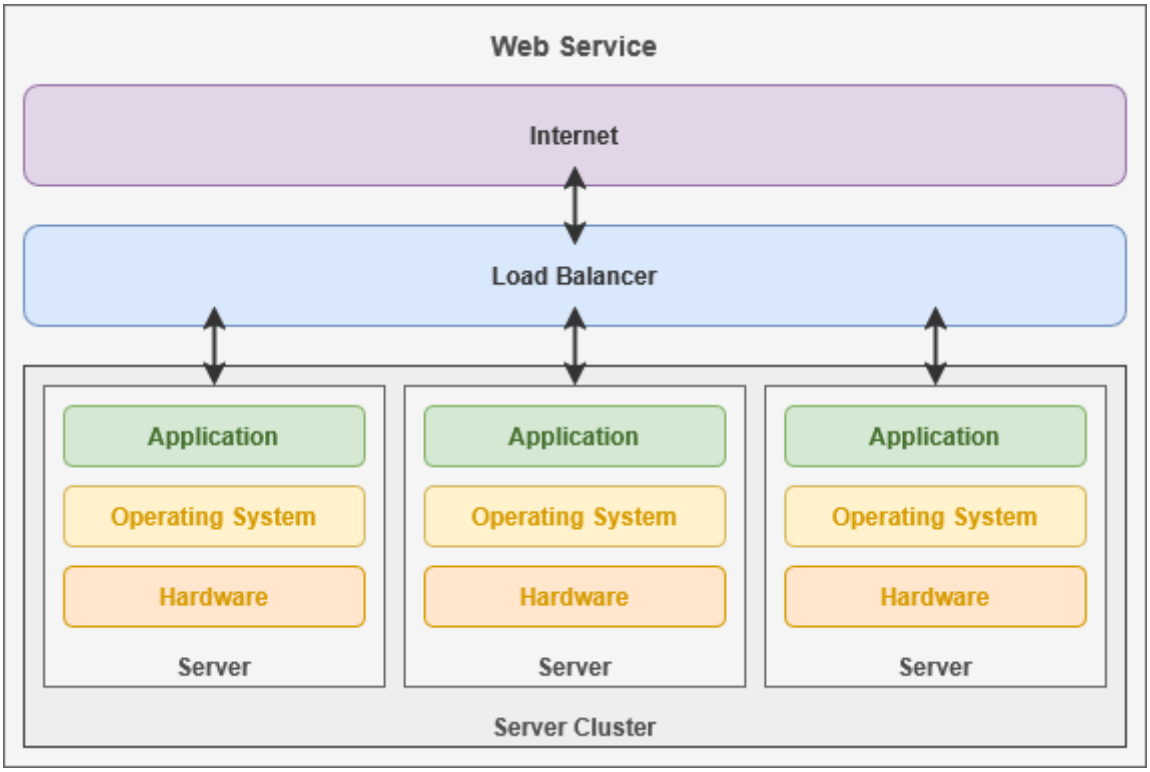

Encrypting data for secure transmission over the open internet can require a great deal of hardware resources. A reverse proxy can be indispensable because they can handle the encryption (and decryption) of secure data with clients outside of the Local Area Network. This stream can also be compressed so that less data needs to be transmitted over the internet.

By acting as the singular access point for a plethora of behind-the-scenes services, they can also represent a single point of failure where one malfunction can leave you without any access. This is a contingency that we prepared for by setting up local network access protocols, but it might be prudent to operate your web server in a place where you can access a mouse and keyboard.

Secure Web Application Gateway

Secure Web Application Gateway – more commonly referred to as SWAG – is a community-driven project by LinuxServer.io to host a secure and easy-to-use reverse proxy and web server.

Nginx ("engine x") is an open-source HTTP server, reverse proxy, web cache and load balancer that is used to power the majority of corporate domains. This is the core of the SWAG self-hosting project – including pre-configured add-ons and templates for accessing popular self-hosted services behind a reverse proxy. We will be configuring this to access our individual services from a centralized location.

Nginx can be used to host an HTTP website with full PHP functionality and add-on modules for accomplishing specific tasks. By default, SWAG will host a basic website until we configure our reverse proxy to access our self-hosted services.

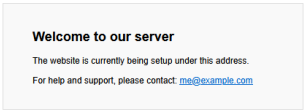

When we use the Internet to connect to our SWAG container running in Docker, we are connecting through an HTTP (port 80) or secure HTTPS connection (port 443). When the reverse proxy receives traffic, it will inspect the data to see where and how it should be delivered. When accessing the website through the primary domain – such as example.com – the traffic will be transmitted to the self-hosted website.

When accessing a sub-domain – such as jellyfin.example.com – the data will be transferred to the Jellyfin server in order to handle the client's request. All data transmitted between the reverse proxy server and Jellyfin is hidden from the World Wide Web and protected behind our Local Area Network.

> A reverse proxy forwards a website as it is – meaning that you need to ensure security through encryption and password-protection before connecting it to the Internet.

We are hosting services using our server and attacing their web interface to local ports that are accessible to other computers on our local network. For example, this is how we access Cockpit at its default port 9090 through a Web browser. By using a reverse proxy, we can route this same access through a web sub-domain – such as cockpit.example.com.

> It is recommended that services like Cockpit are restricted to Local Area Network access.

SWAG makes it easy to automatically generate an SSL certificate using a variety of mechanisms. These form the foundation of the HTTPS protocol by verifying the identity of the server and encrypting the data sent through a secure TCP connection.

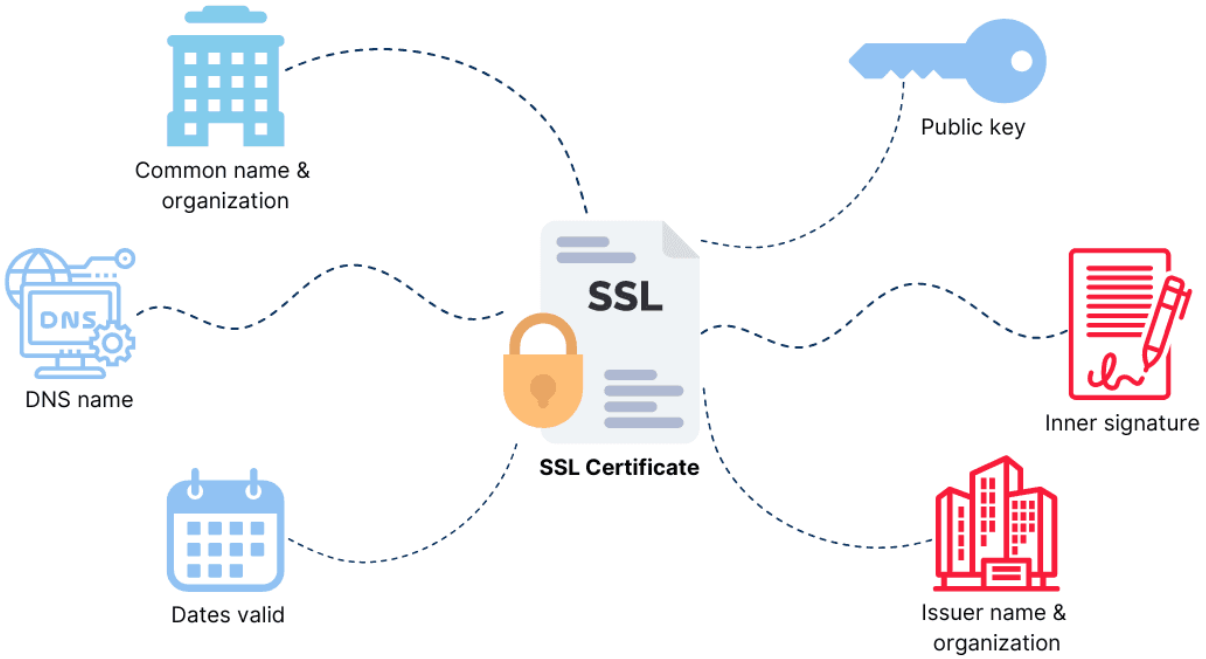

SWAG also provides pre-configured settings for integration with other security-focused add-ons:

Language

Geo-Fencing Services

This service offers the ability to either block – or selectively allow – clients seeking access from specific geographical locations.

SWAG comes with accessible defaults that balance security and convenience. The software it's built on is fully open-source and that means you can configure it however you need – even if that means modifying the underlying software.

The repository includes community-tested templates for each of the services available here, as well as many more. This will require first configuring and installing the software through Docker Compose using Portainer.

Installing SWAG keyboard_arrow_right

Web Domain Name

Connecting Services to the Reverse Proxy

How to edit the SWAG configuration files to connect a service to the internet using the pre-made templates.

Digital Stewardship

Digital self-determination

https://en.m.wikipedia.org/wiki/Digital_self-determination

Digital self-determination is a multidisciplinary concept derived from the legal concept of self-determination and applied to the digital sphere, to address the unique challenges to individual and collective agency and autonomy arising with increasing digitalization of many aspects of society and daily life.

There is no philosophically or legally agreed-upon concept of digital self-determination yet. Broadly speaking, the term describes the attempt to comprehensively project the pattern of human self-determination (as first explored in disciplines like philosophy and psychology, and in the law) into the digital age.

The concept has been included in an official document for the first time by ARCEP, the French Telecoms Regulator, in a section of its 2021 Report on the State of the Internet,[1] exploring the work on "Network Self-determination"[2] conducted by Professor Luca Belli.

The concept of self-determination relates to concepts of subjectivity, dignity, and autonomy in classic central-European philosophy and derived from Immanuel Kant's conception of freedom. Self-determination presupposes that human beings are entities capable of reason and responsibility for their own rationally chosen and justified actions (autonomy), and ought to be treated accordingly. In formulating his categorical imperative (kategorischer Imperativ), Kant suggested that humans, as a condition of their autonomy, must never be treated as a means to an end but as an end in itself. The pattern of self-determination similarly aims at enabling autonomous human beings to create, choose and pursue their own identity, action, and life choices without undue interference.

The increasing digitization of most aspects of society poses new challenges for the concept and realization of self-determination.[7] While the digital sphere has ushered in innovation and opened up new opportunities for self-expression and communication for individuals across the globe, its reach and benefits have not been evenly distributed, oftentimes deepening existing inequalities and power structures, commonly referred to as a digital divide. Moreover, the digital transformation has enabled, oftentimes unbeknownst to individuals, the mass collection, analysis, and harvesting of personal data by private companies and governments to infer individuals' information and preferences (e.g., by tracking browsing and shopping history), influence opinions and behavior (e.g., through filter bubbles and targeted advertisements), and/or to make decisions about them (e.g., approving or not a loan or employment application), thus posing new threats to

individuals' privacy and autonomy.[8]

Although the definition of digital self-determination is still evolving, the term has been used to address humans' capacity (or lack thereof) to exercise self-determination in their existence in and usage of digital media, spaces, networks, and technologies, with the protection of the potential for human flourishing in the digital world as one of the chief concerns.[9]

Data Sovereignty

Data sovereignty is the concept that data is subject to the laws and regulations of the country or region where it was generated or is stored. It essentially means that a country has the authority to govern the data within its borders, including how it's stored, processed, and protected.

https://en.m.wikipedia.org/wiki/Data_sovereignty

Data sovereignty is the idea that data are subject to the laws and governing structures of the nation where they are collected. In other words, a country is able to control and access the data that is generated in its territories.[1] An example of a nation's data sovereignty policy would be Australia's Privacy Policy guidelines, also known as APP. [2] The APP contains 13 principles for how all personal or organizational data in Australia is meant to be kept.[2] For many countries, the issue of data sovereignty is presented as an issue of national security with concerns over being able to protect citizens' personal data.[1] Data can be used to help improve medical care, reinforce national security as well as have a positive impact on many economic and social infrastructures but may also be used for identity theft and other data related attacks.[1]

The concept of data sovereignty is closely linked with data security, cloud computing, network sovereignty, and technological sovereignty. Unlike technological sovereignty, which is vaguely defined and can be used as an umbrella term in policymaking,[3] data sovereignty is specifically concerned with questions surrounding the data itself.[4] The issue of managing data sovereignty can be considered more complex when introducing the idea of cloud computing, where data can be accessed globally; meaning organizations and companies must comply with multiple nations data laws.[5] Data sovereignty is also associated with data localization, the requirement that data be stored within a specified region, and data residency, the actual location in which the data is stored, such as cloud servers.[1]

Data sovereignty as the idea that data is subject to the laws and governance structures within one nation, is usually discussed in one of two ways: in relation to Indigenous groups and Indigenous autonomy from post-colonial states, or in relation to transnational data flow.[6] With the rise of cloud computing, many countries have passed various laws around the control and storage of data, which all reflect measures of data sovereignty.[4] More than 100 countries have some form of data sovereignty laws in place.[7] With

self-sovereign identity (SSI), the individual identity holders can fully create and control their credentials, although a nation can still issue a digital identity in that paradigm.[8]

https://en.m.wikipedia.org/wiki/Data_sovereignty_(data_management)

Data sovereignty is the ability of a legal person or an organisation to control the conditions that data is shared under, and how that shared data is used, as if it were an economic asset.[1][2] It can apply to both primary data and secondary data derived from data, or metadata.[3] In order to use restricted data, data consumers must accept the conditions that it is provided under.[4] In turn, the legal persons sharing data must trust other entities with it. Trust can be supported through the use of a suitable secure information system (such as a data space) which identifies, authenticates, and certifies users.[5]

The data sovereignty of individual legal persons can conflict with national data sovereignty.[6] Currently, a natural person does not have a statutory right to exclusively control how their data is shared and used. However, they can make it part of a contract, and offer it as payment.[7] The most common method for a legal person to impose its data sovereignty is through contract law.[8] Such a contract includes the terms of use, access and control policies, commercial conditions and jurisdiction.[3]

The use of cloud services can make it difficult to determine where data is physically located and which law applies

A common criticism of data sovereignty brought forward by corporate actors is that it impedes and has the potential to destroy processes in cloud computing.[30] Since cloud storage might be dispersed and disseminated in a variety of locations at any given time, it is argued that governance of cloud computing is difficult under data sovereignty laws.[30] For example, data held in the cloud may be illegal in some jurisdictions but legal in others.[4]The concept of a sovereign cloud is proposed as a solution to address this challenge.[31][32]

Data governance

Scholar Shoshana Zuboff popularized the term surveillance capitalism to refer to the private sector's commodification of users' personal data for profit (e.g. via targeted advertising), leading to increased vulnerability to surveillance and exploitation. Surveillance capitalism relies on centralized data management models wherein private companies retain ownership and control over the users' data. To guard against the challenges to individuals' privacy and self-determination, various alternative data governance models have been recently proposed around the world, including trusts,[32] commons,[33] cooperative,[34] collaboratives,[35] fiduciaries,[36] and "pods".[37] These models have some overlap and share a common mission to give more control to individuals over their data and thus address the current power imbalances between data holders and data subjects.[38]

Freedom & Privacy

Intellectual freedom & censorship

Censorship online can be carried out (to varying degrees) by actors including totalitarian governments, network administrators, and service providers. These efforts to control communication and restrict access to information will always be incompatible with the human right to Freedom of Expression.5

Censorship on corporate platforms is increasingly common, as platforms like Twitter and Facebook give in to public demand, market pressures, and pressures from government agencies. Government pressures can be covert requests to businesses, such as the White House requesting the takedown of a provocative YouTube video, or overt, such as the Chinese government requiring companies to adhere to a strict regime of censorship.

People concerned with the threat of censorship can use technologies like Tor to circumvent it, and support censorship-resistant communication platforms like Matrix, which doesn't have a centralized account authority that can close accounts arbitrarily.

You must always consider the risks of trying to bypass censorship, the potential consequences, and how sophisticated your adversary may be. You should be cautious with your software selection, and have a backup plan in case you are caught.

Representation of diverse realities and viewpoints

Internet activist Eli Pariser coined the term filter bubble to refer to the reduced availability of divergent opinions and realities that we encounter online as a consequence of personalization algorithms like personalized search and recommendation systems.[30] Filter bubbles have been suggested to facilitate a warped understanding of others' points of view and the world. Ensuring a wide representation of diverse realities on digital platforms could be a way of increasing exposure to conflicting viewpoints and avoiding intellectual isolation into informational bubbles

Personal privacy

A common counter-argument to pro-privacy movements is the notion that one doesn't need privacy if they have **"nothing to hide."** This is a dangerous misconception, because it creates a sense that people who demand privacy must be deviant, criminal, or wrong.

You shouldn't confuse privacy with secrecy. We know what happens in the bathroom, but you still close the door. That's because you want privacy, not secrecy. There are always certain facts about us—say, personal health information, or sexual behavior—that we wouldn't want the whole world to know, and that's okay. The need for privacy is legitimate, and that's what makes us human. Privacy

is about empowering your rights over your own information, not about hiding secrets.

Take cookie consent forms, for example. You may encounter these dozens of times per day on the various websites you visit, with a nice array of checkboxes and sliders which allow you to "curate" your preferences to exactly fit your needs. In the end, we just hit the "I Agree" button, because we just want to read the article or make a purchase. Nobody wants to complete a personal privacy audit on every single website they visit. This is an exercise in choice architecture, designed to make you take the easy route out instead of delving into a maze of configuration options that don't need to exist in the first place.

Control over your privacy inside most apps is an illusion. It's a shiny dashboard with all sorts of choices you can make about your data, but rarely the choices you're looking for, like "only use my data to help me." This type of control is meant to make you feel guilty about your choices, that you "had the choice" to make the apps you use more private, and you chose not to.

Privacy is something we need to have baked into the software and services we use by default, you can't bend most apps into being private on your own.

Corporate surveillance

For many people, tracking and surveillance by private corporations is a growing concern. Pervasive ad networks, such as those operated by Google and Facebook, span the internet far beyond just the sites they control, tracking your actions along the way.

Additionally, even companies outside the *AdTech* or tracking industry can share your information with data brokers (such as Cambridge Analytica, Experian, or Datalogix) or other parties. You can't automatically assume your data is safe just because the service you're using doesn't fall within the typical AdTech or tracking business model.

Technological Stewardship

and Responsible Innovation

The call for responsible innovation is a call to address and account for technology’s short- and long-term impacts within social, political, environmental, and cultural domains.

Technological stewardship stands as the corollary of this mindset: a commitment to anticipate and mitigate technology’s potential for disruption and especially harm and to guide innovation toward beneficial ends.

Adopted from the domain of engineering ethics, where it “involves taking a value sensitive approach to embedding ethics, sustainability, and EDI (equity, diversity, and inclusivity) principles

into the practice and culture of engineering" [3, p. 34], technological stewardship belongs to any who would take up the mantle.

To be a technological steward means to be committed to an ethos of *care*: to understanding the systematic impacts of innovation and its philosophical dimensions, and to cultivating the capacity to apply this knowledge to drive development into new realms of humane engineering and design. To be a steward is to see care as the foundation of one's professional identity and not as an add-on or bonus. Technological stewardship is therefore about an optimistic turn toward the future of more responsible innovation—about recognizing the possibility for technology to transform culture and civilization and to steer away from the maelstrom of unfettered innovation.

Dialogue and collaboration across diverse perspectives is essential for developing actionable technological solutions that attend in responsible ways to the evolving needs of society.

Human-centered design of user interfaces and experiences

Scholars have coined the term attention economy to refer to the treatment of human attention as a scarce commodity in the context of ever-increasing amounts of information and products. In this view, the increasing competition for users' limited attention, especially when relying on advertising revenue models, creates a pressing goal for digital platforms to get as many people as possible to spend as much time and attention as possible using their product or service. In their quest for users' scarce attention, these platforms would be incentivized to exploit users' cognitive and emotional weaknesses, for example via constant notifications, dark patterns, forced multitasking, social comparison, and incendiary content.[8] Advocates of human-centered design in technology (or humane technology) propose that technology should refrain from such 'brain-hacking' practices, and instead should support users' agency over their time and attention as well as their overall wellbeing.[31]

Digital literacy

Digital literacy and media literacy have been proposed as necessary for individuals to acquire the knowledge and skills to use digital tools as well as to critically assess the content they encounter online, create their own content, and understand the features and implications of the digital technology used on them as well as the technology they consciously and willingly engage with.[28] In addition to basic digital navigation skills and critical consumption of information, definitions of digital literacy have been extended to include an awareness of existing alternatives to the digital platforms and services used, understanding how personal data is handled, awareness of rights and existing legal protections, and of measures to independently protect one's security and privacy online (e.g., the adoption of obfuscation techniques as a way of evading and protesting digital surveillance[29]).

Access to digital infrastructure and tools

Bridging the various forms of existing digital divides and providing equitable and fair access to digital technologies and the internet has been proposed as crucial to ensure that all individuals are able to benefit from the digital age, including access to information, services, and advancement opportunities.[24][25]

In this sense, the concept of Digital Self-determination overlaps with the concept of "Network Self-determination"[26] as it emphases that groups of unconnected and scarcely connected individuals can regain control over digital infrastructures, by building them and shaping the governance framework that will organise them as a common good.[27] As such, Belli stresses that network self-determination leads to several positive externalities for the affected communities, preserving the Internet as an open, distributed, interoperable and generative network of networks.[2]

Security & Privacy

When connecting your server to the open internet – whether by VPN, reverse proxy or a combination of both – it is also important to focus on sustainable security and privacy solutions.

Security & Privacy

Critical Thinking

Take about criticality and it's importance. Just because we've always done in this way doesnt mean we need to.

Attack surface

Close unused containers

Close unused ports

Uninstall unused software

Disable unused hardware

What is the difference between the two?

Security vs convenience

Privacy

The ability to control who can access our personal information and what they can see.

Security

The proactive measures we take to keep ourselves free from digital threats.

Obscurity

When we self-host our own services, we remove ourselves from the more typical security threats. While OwnCloud may have software vulnerability, they are quickly patched. Actors more commonly attack large companies where more data can be taken. Unless specifically targeted, you are less likely to suffer a security breach

Two-Factor Authentication

Passwords & Passphrases

Basic Authentication

SWAG makes it easy to use basic HTTP authentication through your web browser to password-protect a basic website.

Authelia

SWAG can also be easily integrated with Authelia, a service that can provide a unified login experience through single sign on. Anytime someone tries to access a page protected by authelia, they will be forced to login. Once logged in, you can access all Authelia protected applications without signing it again. For added security, you can configure 2FA with your BitWarden service.

Security & Privacy

LAN-Only Access

How to configure swag to restrict access to your current IP address.

Security & Privacy

fail2ban

This service is provided by SWAG by default it is used to automatically ban someone trying to access your server and using invalid authentication. It does this by banning the IP address of the attacker for an increasing amount of time for each offense.

Fail2ban works out of the box with intrusion detection for basic http authentication as well as Authelia. For other applications that do not use these, such as Plex, Jellyfin, and other apps that have their own login screen, you will need to manually configure them. Fail2ban works with these services by monitoring their log files for evidence that someone failed to login.

Security & Privacy

CrowdSec

CrowdSec is an open-source security solution for responding to malicious actors on your services. They take a unique approach by levering the power of the open-source community to actively share information about previous cyber attacks and protect your community.

Security & Privacy

Restricting Access by Geographic Region

https://www.linuxserver.io/blog/securing-swag

https://github.com/linuxserver/docker-mods/tree/swag-dbip

Security & Privacy

Hiding from Search Engines

https://www.linuxserver.io/blog/securing-swag#search-results

Blocking robots.

Router Configuration

We need to set up our router to make sure it's ready to send and receive traffic through the World Wide Web.

Accessing our Router Dashboard

Connecting our web server to the internet will require gaining access to our router's administrative dashboard. We can find our routers address by using the command:

```
ip -o route show | grep default | perl -nle 'if ( /via\s+(\S+)/ ) {print $1}'
```

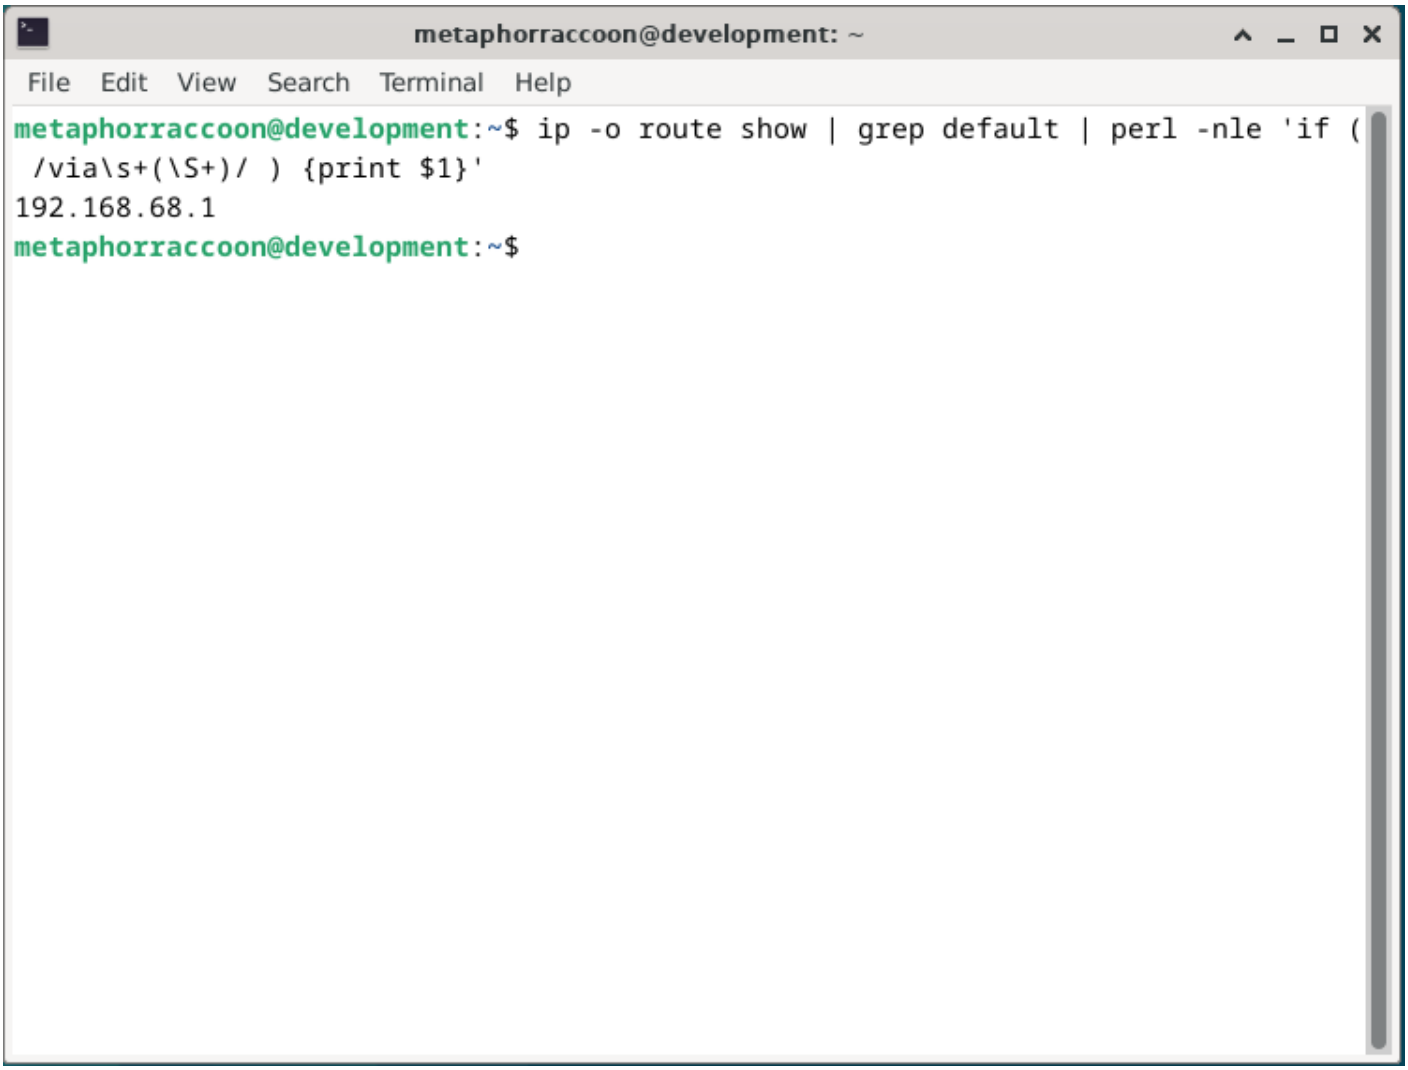

Going to our web browser, we can enter the IP address returned by our command:

```
192.168.68.1
```

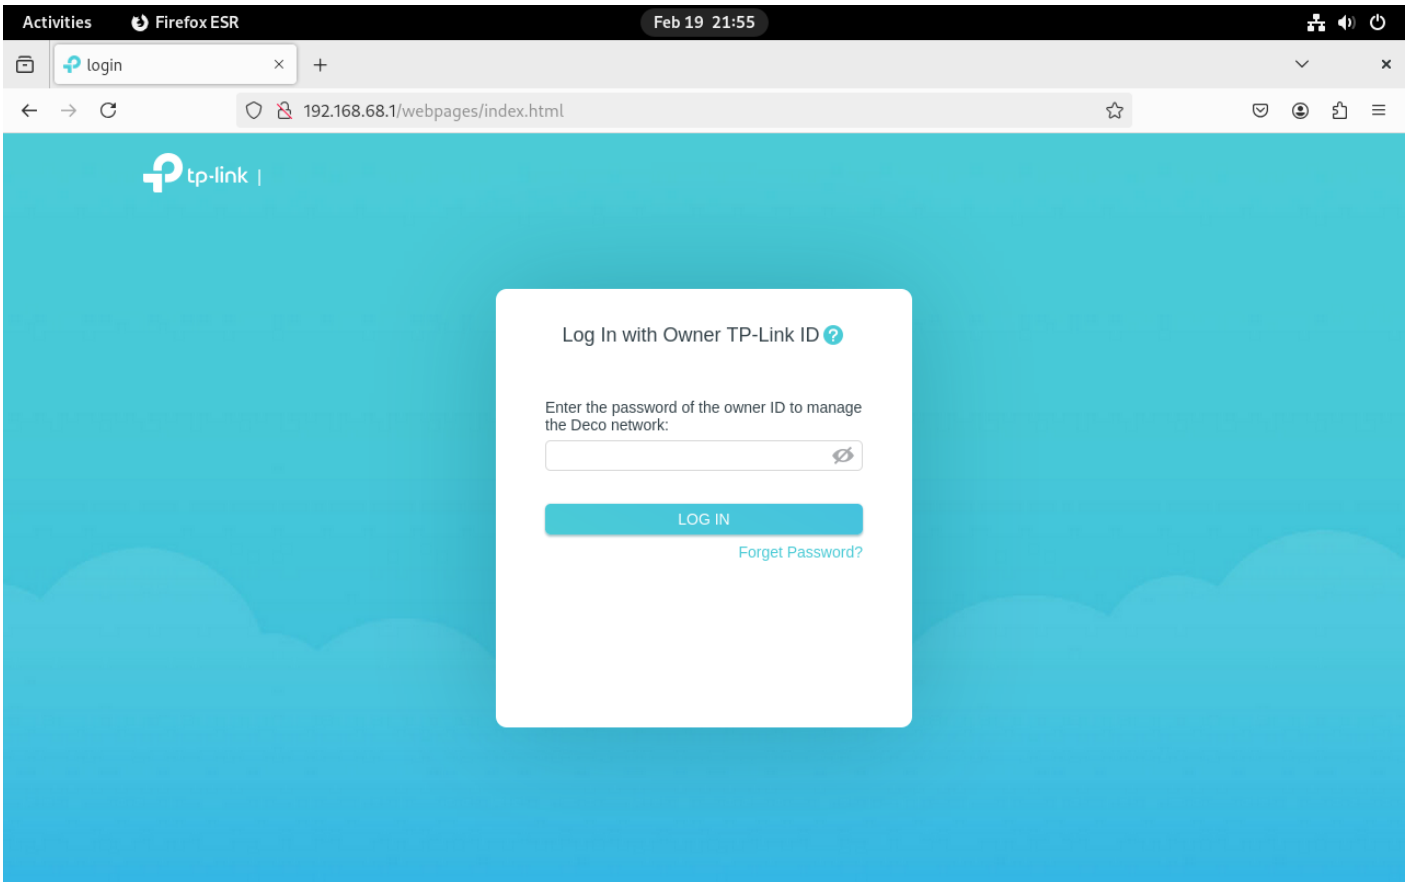

> Some routers may require an app to access all configuration settings.

Sign in. Default username and password.

Router Configuration

Securing the Administrator Account

We need to make sure that our router uses a secure password.

Reserving an IP Address

We've been requesting the same IP for our computer from the router, but this doesn't stop another computer on the network from requesting the same IP address and causing a conduct. It's always best for security and stability to reserve the IP address for our server at the router level.

DHCP address reservation for common routers

Router Configuration

Connecting Your Personal Server to the Internet

We will need to set our router to forward any requests it gets through port 80 and port 443 to our server. This is how we'll connect our server to the Internet to accept request for web traffic. Web browser traffic uses port 80 for insecure HTTP requests as well as port 443 for secure HTTPS requests.

Even though port 80 is used for insecure traffic, our server is configured to forward all traffic from this port to 443 to create a secure connection. Our server will not allow insecure connections. If we close port 80, the browser will not know how to respond if we go to a website without the protocol. Many web browsers allow insecure web site connections by default and will default to checking for the website using http. If we have the http port closed, the traffic can't be channeled into a secure connection.

Find port forwarding or nat forwarding.

Forward port to our servers IP Address

openvpn server port

https://portforward.com/router.htm

What Next?

Imstall services, create another server, awesome self hosted,

how you can help. all of these software packages are open source. they seek and thrive off of feedback from their community of users. contribute to them. there are problems within the open source community, but we can help work to improve them. talk about how demographics about coding and software development. there's so many ways to help a community and create space within an open community than just coding.