Mikko Karjalainen

From On-Premises to Serverless

Metropolia University of Applied Sciences

Bachelor of Engineering

Information Technology

Bachelor’s Thesis

19 October 2021

Abstract

Author:	Mikko Karjalainen
Title:	From On-Premises to Serverless
Number of Pages:	65 pages
Date:	19 October 2021
Degree:	Bachelor of Engineering
Degree Programme:	Information Technology
Professional Major:	IoT and Cloud Computing
Supervisor:	Janne Salonen, Head of School, School of ICT

This thesis focuses on the customer's on-premises server infrastructure and their recent serverless goal. The aim of this thesis is to migrate all the local servers found in the Nordic and Baltic countries into a hybrid cloud model and to decommission the local servers.

In the assessment phase, all on-premises services and setups were visited and a migration plan was set for each of them. While the sites' migrations were done in parallel, each site was done as its own individual project. They were progressed server-by-server due to the large size of the project that required a lot of coordination and communication between different departments across several countries.

The project was concluded when all the targeted servers were successfully migrated to public or private cloud. Users were able to work regularly but reap the benefits of the new cloud services.

Throughout the project, a runbook was constructed to aid other IT departments with their on-premises server decommissioning. Based on this thesis it should also be easier to judge if cloud migration is the right solution for certain infrastructures, what the prerequisites could include, and which steps can be executed to undertake a migration.

Keywords: server, cloud services, on-premises, serverless

Tiivistelmä

Tekijä:	Mikko Karjalainen
Otsikko:	Konesalista Palvelittomaan Ratkaisuun
Sivumäärä:	65 sivua
Aika:	19.10.2021
Tutkinto:	Insinööri (AMK)
Tutkinto-ohjelma:	Tieto- ja viestintätekniikka
Ammatillinen pääaine:	Verkot ja pilvipalvelut
Ohjaaja:	Janne Salonen, Osaamisaluepäällikkö, ICT ja tuotantotalous

Insinöörityössä tarkoituksena oli toteuttaa paikallisten palvelimien migraatio pilvipalveluihin. Projektin pyrkimyksenä oli saada asiakas pois perinteisestä konesaliratkaisusta kohti hybridipilviratkaisua, jossa kaikki paikalliset palvelimet ajetaan alas ja palvelut siirretään yksityiseen tai julkiseen pilveen. Työn vastuualueeseen kuuluivat Pohjoismaat ja Baltian maat, joissa viidessätoista toimipisteessä oli omat paikalliset palvelimet.

Insinöörityön alussa käytiin läpi suunnitteluvaihe, jossa kunkin toimipisteen palvelimien tilanne tutkittiin sekä tehtiin migraatiosuunnitelma. Palvelimet olivat rakenteiltaan samankaltaisia, mutta erosivat toisistaan yksityiskohtien sekä toimipisteiden käyttäjien käytäntöjen mukaisesti. Kunkin toimipisteen palvelimien alasajoa käsiteltiin omana projektina, mutta eteneminen tapahtui yhtäaikaisesti kaikissa toimipisteissä.

Projekti vaati paljon koordinoimista ja kommunikointia eri osastojen kanssa useissa eri maissa. Projektin laajuuden takia edettiin vaihe vaiheelta palvelimen siirto kerrallaan. Pyrittiin myös löytämään ratkaisuja siihen, miten uudet muutokset saataisiin automatisoitua IT:n sekä käyttäjien ajan säästämiseksi.

Lopputuloksena oli, että kaikki määrätyt paikalliset palvelimet saatiin siirrettyä pilvipohjaisiin ratkaisuihin ja paikalliset palvelimet purettua. Onnistuminen näkyi siten, että käyttäjät eivät juuri havainneet muutosta ja työnteko sujui ongelmitta.

Projektin edetessä tehtiin ohjekirjaa, johon merkattiin tehokkaampia ratkaisuja ajamaan paikalliset palvelimet alas. Tällä opastetaan maailmanlaajuisesti muita IT-osastoja toimimaan samoin heidän vastuualueeseen kuuluvien palvelimien alasajossa. Tämä raportti pyrkii myös toimimaan mahdollisena ohjeena niille, jotka ovat kiinnostuneita siirtymään modernimpaan pilvipohjaiseen infrastruktuuriin.

Avainsanat: palvelin, pilvipalvelut, konesali, palveliton

Contents

List of Abbreviations

AD: Active Directory. A directory service used to manage permissions to different network resources.

AD DS: Active Directory Domain Services. A role within Active Directory that stores and manages information about the resources in the network.

CAB: Change Advisory Board. A group of people that ensure changes to the IT environment are planned and executed the best way possible.

CMDB: Configuration Management Database. An asset management tool that stores information on the organization’s hardware and software.

CSP: Cloud Service Provider. A third-party company that offers cloud computing services.

DC: Domain Controller. A server managing users’ security authentications within the computer network.

DNS: Domain Name System. Resolves IP addresses into human-recognizable Internet addresses.

EOL: End of Life. The period a product reaches where the manufacturer no longer supports it.

GPO: Group Policy Object. A collection of policy settings set within Group Policy.

iLO: integrated Lights-Out. HP proprietary server management system embedded to HP servers.

ITSM: IT Service Management. The process of delivering and managing IT services the best way possible to meet the needs of the business.

MFT: Managed File Transfer. A technology used to securely, reliably and efficiently transfer data between systems.

NAS: Network Attached Storage. A device used to store data and share it through the network.

OU: Organizational Unit. Directories used to group objects within the Active Directory based on logical reasons.

PaaS: Platform as a Service. A cloud service model where the service provider delivers a framework for app development and management. developing and managing applications.

PXE: Preboot eXecution Environment. A process used to boot a device from the network.

RODC: Read-only Domain Controller. A server hosting a non-writable replica of the Active Directory database.

SaaS: Software as a Service. A cloud service model where the service provider delivers applications via the Internet.

SCOM: System Center Operations Manager. A system used to monitor and report devices' statuses.

UPS: Uninterrupted Power Supply. A device that protects from power issues and temporarily provides power during an outage.

VM: Virtual Machine. A software-defined computer that can run an OS without being tied to a physical machine.

VPN: Virtual Private Network. A private network with two or more clients connected through the public network.

1 Introduction

Cloud technology is by no means a new concept but to many, it still brings a level of uncertainty even to this day. The term “cloud” is used more and more frequently while its definition remains a bit dim. To put it simply, the cloud refers to a network of servers that manages the infrastructure and is used to deliver computing services, e.g., files and applications, based on the user’s current needs.

The reason for cloud solutions’ increasing popularity in organizations is the ever-growing list of advantages it brings compared to the old-fashioned on-premises approach. By letting a third party manage the underlying infrastructure, the organization can get better scalability and elasticity, often resulting in noticeable cost-savings. However, migrating the entire infrastructure (or parts of it) to the cloud can be a long and complex process.

This is what the company, Global Blue, had decided to go for. Global Blue is a tax-free company, having a global footprint across all continents and over 51 countries, delivering tax free shopping and added-value payment solutions. Many of these countries have at least one site hosting local servers, meaning a lot of servers are distributed across the world. This on-premises approach was getting ever costlier and more troublesome to maintain, especially with the remote working caused by COVID-19. A business decision was made to decommission most, if not all, of the local servers and migrate to the cloud. This is what the thesis focuses on.

There are two goals for this thesis. The main goal was to relocate all locally hosted servers within the Nordic and Baltic regions away from on-premises. Another goal of the project was to get enough knowledge of all the moving parts to guide the rest of the company to do the same around the world. This is what a runbook was written and shared for. This thesis report also aims to work as potential guidance for others hoping or curious to migrate to the cloud.

2 Theory

This thesis report will first start by going through a brief theory about the key elements of this project; servers and cloud. With some background information, the next chapter will give a better idea of what the infrastructure was beforehand and what the plans were to go serverless while glimpsing at the obstacles found along the way and how they were tackled. And at the end is the outcome of the project as well as some future insights.

2.1 Servers And Virtualization

Many of the things we use computers for, like applications, services, or files, are running or were obtained from a server. Servers are computer programs with the primary purpose of running their designated service, often used to share its resources with users. Most commonly known servers would have to be the web- and fileservers – both hosting files for different services and ensuring the proper files are given to the right users.

Servers can be created and run on pretty much any device, but servers requiring more processing power often run on server-dedicated hardware that can better handle the computing required to efficiently deliver the services needed for the users.

IT organizations have different approaches to dealing with the server-side, but for Global Blue, there was additional hardware to make it easier to monitor and manage the servers and have higher service uptime. The hardware portion will be covered better in chapter 3. This traditional way of running servers became troublesome and costly, so virtualization environments began to kick in.

Virtualization is a very popular solution in IT infrastructures and is widely used in the cloud [1]. Virtualization allowed creating multiple virtual machines (VM) behind one physical computer. To end-users, VMs were just like any other computer because the operating system and applications worked the same. For

IT, however, it offers better manageability and is a more cost-efficient solution because the same hardware could be used for hosting multiple operating systems, resulting in less hardware needed to be purchased, upgraded, or maintained [2].

VMs can also be called virtual servers. The only difference to a traditional one is that multiple server OSes are running on the same hardware. Refer to figure 1 for a comparison between the traditional architecture compared to a virtualized one.

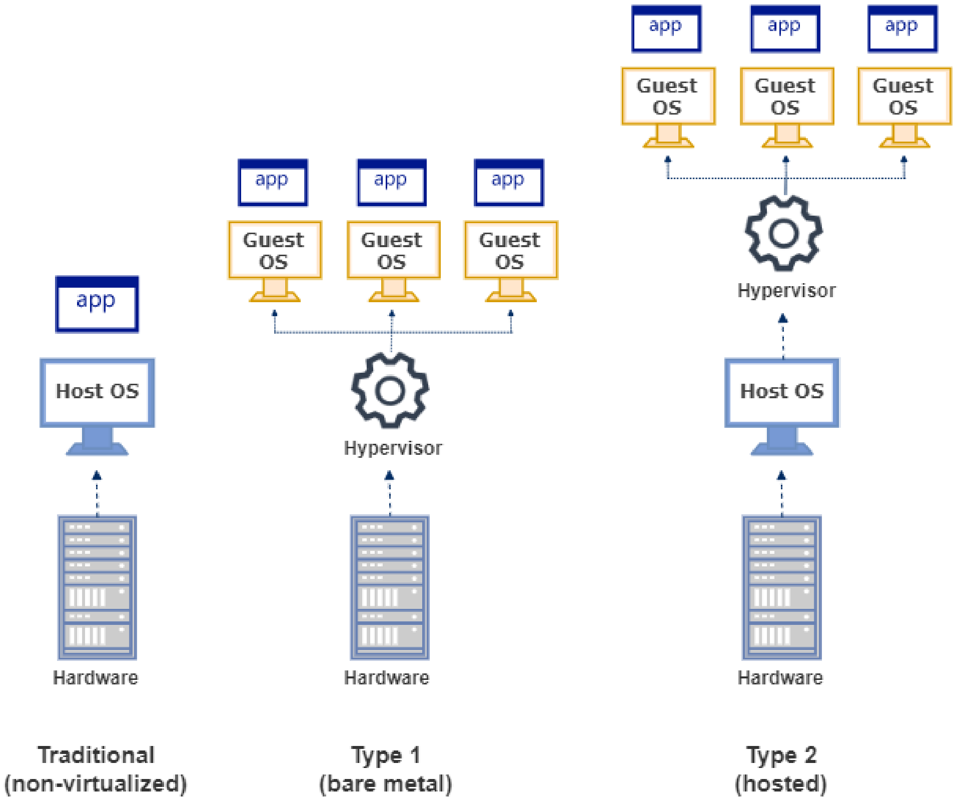

Figure 1. Traditional architecture vs. different virtualization types.

VMs are essentially virtual hard drive files running on the VM host, and due to this nature, they can be moved to another VM host. This could be for hardware upgrade purposes or if the hardware was to fail or reach its end-of-life (EOL). The VMs can be moved to another physical computer with a hypervisor.

A hypervisor is software that creates and manages the virtualized environment. There are two types of hypervisors. Refer back to figure 1 – the type 2 (also known as hosted) hypervisor has a regular OS to which the hypervisor is installed and running on top. [3.]

Type 2 hypervisors have no direct access to the physical hardware and have fewer management options compared to type 1 hypervisors. They are often used for virtual desktops rather than virtual servers as they have more security and latency concerns due to having a separate OS hosting the environment. [3.]

Type 1 (also known as “bare metal” or native) hypervisor eliminates the need of having an operating system in the middle to host the VMs. This reduces resource usage and eliminates the security risks that come with the regular OS. The hypervisor acts as a lighter version of the virtualized environment and is often used to run virtual servers at data centers. [3.]

Oracle VirtualBox and VMware Workstation are examples of type 2 hypervisors, whereas the most common type 1 hypervisors are VMware ESXi and Microsoft Hyper-V. VMWare was sometimes used for special cases, like in one Finnish server, but Global Blue uses Hyper-V for all its standard virtualized servers. Hyper-V will be covered more in this thesis.

Hyper-V can easily be mistaken for a type 2 hypervisor because it is often found as additional software on a Windows Server OS. Hyper-V hypervisor is installed with access directly to the hardware, making it a type 1 hypervisor. [4.] See figure 2 below.

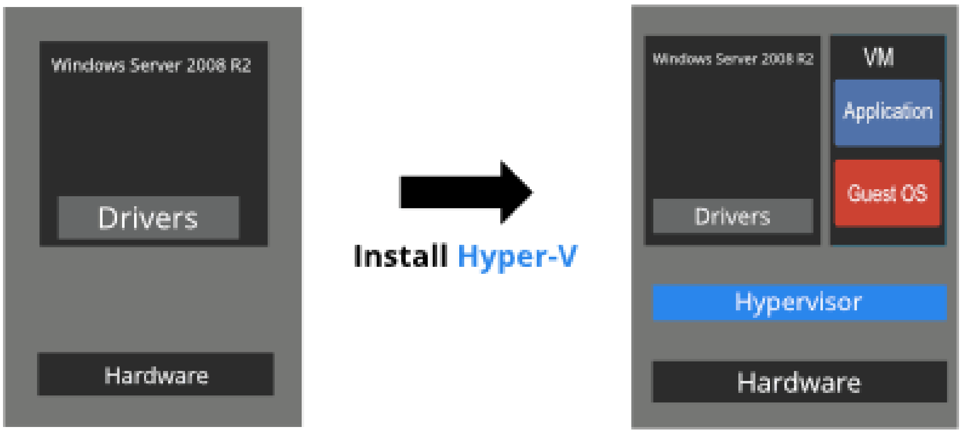

Figure 2. Installing Hyper-V creates an additional virtualization layer below the original OS [4].

While on-premises, the local IT has to worry about scaling the hardware based on the server or user needs, which might require purchasing and replacing new components, and the server cooling and the space it takes at the office.

You can start to see how this can be troublesome and expensive for the IT and the company as a whole. What if there was a way to shift the responsibility of managing the infrastructure (or parts of it) to someone else?

2.2 Cloud Services and Serverless Architecture

If servers are needed to run specific services, then what does serverless mean? Based on the name, one might assume the servers are entirely removed from the equation, but that is not quite true. The serverless architecture simply means that the organization does not have servers of their own but instead has the services running on servers owned by a third party, also known as cloud service providers (CSP). There are a lot of CSPs, each with a different set of services, features, and prices, but the biggest three are Google Cloud, Microsoft Azure, and Amazon Web Services [5].

The term “cloud” is often used very lightly, but for something to qualify as cloud, it needs to meet five criteria. The hardware resources need to be shared or pooled to multiple users to make use of the hardware available better, and they also need to be easily accessible to the users, usually via the Internet. The services or resources provided need to be available on-demand (which often shows as 99.99% uptime) and they need to scale based on the demand. Finally, all these services are charged based on how much the services or resources are used instead of the more traditional monthly or annual fees. [6.]

Instead of the organization worrying about the hardware and servers, they can move them to the cloud and only get the service they need. Typical cloud services are Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). Figure 3 illustrates the set of responsibilities each service method includes. [7.]

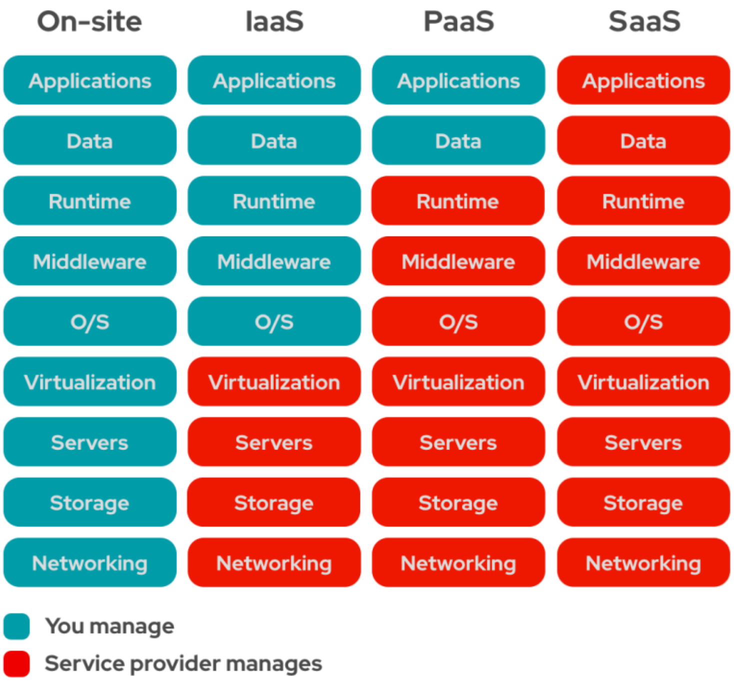

On-site	IaaS	PaaS	SaaS
🟦 Applications	🟦 Applications	🟦 Applications	🟥 Applications
🟦 Data	🟦 Data	🟦 Data	🟥 Data
🟦 Runtime	🟦 Runtime	🟥 Runtime	🟥 Runtime
🟦 Middleware	🟦 Middleware	🟥 Middleware	🟥 Middleware
🟦 O/S	🟦 O/S	🟥 O/S	🟥 O/S
🟦 Virtualization	🟥 Virtualization	🟥 Virtualization	🟥 Virtualization
🟦 Servers	🟥 Servers	🟥 Servers	🟥 Servers
🟦 Storage	🟥 Storage	🟥 Storage	🟥 Storage
🟦 Networking	🟥 Networking	🟥 Networking	🟥 Networking

🟦 You manage
🟥 Service provider manages

Figure 3. Responsibilities based on service model [7].

With IaaS, the CSP handles all the hardware- and connection-level responsibilities at the base layer while giving the organization the ability to manage the infrastructure at the OS level. Deploying VMs is a general example of an IaaS. [7.]

PaaS is used when only the platform and applications are under the organization’s management, with CSP handling everything below that, from the OS to the hardware. These are primarily for developers that have their applications or websites to manage. [7.]

If an organization simply needed access to an application, they would not need to worry about the hardware or the platform the application is hosted on. Even software changes and releases and done by the CSP. They can simply pay for the software delivered by the CSP based on its usage. This is what SaaS is. Gmail is a popular example of SaaS. [7.]

Global Blue has a lot of different applications, some of which were previously running on their virtual servers but were then migrated to either Global Blue’s data center or the public cloud. The applications are provided to the users as a SaaS. These processes were part of the thesis and will be covered more in-depth in chapters 3 and 5.

With these cloud services, organizations can benefit a lot on numerous fronts. It depends heavily on which services are used, but scalability, flexibility, and availability are the main benefits. [8.]

Resources scale easily based on the demand: when you need more storage space or computing power, the CSP will offer it and simply charge based on the usage [8]. Sometimes the usage might temporarily spike and this would not be an issue with cloud services, whereas, for on-premises, the IT would need to choose to upgrade the hardware to support the temporary spikes and have an overkill setup for the rest of the time, or save on the hardware but have the system lag or crash during temporary spiking.

Cloud services are almost always accessed via the Internet, and thus it gives better flexibility to customers and employees to access the services no matter where they are, or on what device, just as long they have Internet access. [8.]

No doubt, a big CSP is often able to provide higher availability and uptime than on-premises. The top CSPs have a lot of data centers around the world [9] and if one were to fail, the other one would be used in the meantime as a backup.

The cloud services can differ based on the cloud deployment model the organization is using. There are a few different ones, and these depend on where the infrastructure and services are hosted. The models are called public, private, hybrid, and multi-cloud, as illustrated in the figure below.

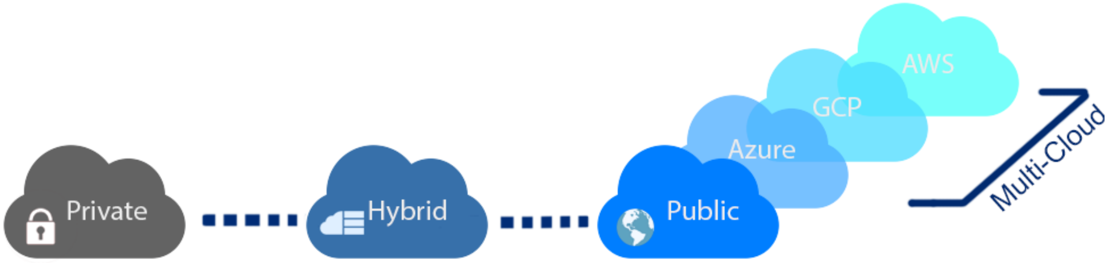

Figure 4. The four cloud deployment models: private, public, hybrid, and multi-cloud.

The previously mentioned CSPs and cloud services are known to be in the public cloud, which most people associate cloud with. As the name implies, the CSP offers the services to the public, which means while your data is separated from other organizations' data via encryptions and authentications, they are still all running on the same infrastructure. If the organization wants to only share the resources and services within itself, it would need to look into the private cloud. [8.]

Private cloud uses the same principles as the public cloud, but instead of the service being shared publicly, it is only provided for the customer alone. The private cloud could also refer to an organization using a data center of its own to deliver services in the same manner. This keeps all the data within the organization itself for higher security, but it also allows the organization to better manage the infrastructure however they please, as opposed to the public cloud, where the CSP manages the bottom layers of the infrastructure. Of course, having your hardware also means being responsible for scaling and maintaining it. [8.]

A hybrid cloud is a combination of private and public. This requires additional maintenance and management from the organizations to ensure that the private cloud infrastructure is in sync with the other part of the infrastructure hosted in the public cloud. This option is often chosen as it offers more flexibility and control over some services hosted on-premises while still having the benefits of a public model. [8.]

Multi-cloud is a bit newer deployment model where more than one CSP is used. CSPs have started introducing more services to make it easier to move or share infrastructure with other CSPs. This allows organizations to choose which CSP to use for which services depending on their needs and budget. [8.]

Global Blue uses a hybrid cloud model where they have their private data centers and host many computing services within the public cloud. The cloud deployment model plays a huge part in the serverless project as some of Global Blue’s local services are not optimal nor possible to migrate to either private or public cloud. These services will be covered next.

3 Project Overview

This chapter will go more in-depth on what the infrastructure looked like at Global Blue before the serverless solution was started, the drawbacks and bottlenecks of the traditional on-premises setup, and where each service was planned to be migrated. It should give a better idea of why Global Blue decided to decommission most of the on-premises servers and why said solutions were found most optimal.

3.1 The Starting Point

Global Blue offices are split into branch and refund offices. The branch offices mainly consisted of moving sales staff, whereas refund offices had stationary employees delivering tax-free refund services. Fifteen sites within the Nordic and Baltic regions had servers, and all planned to be decommissioned. Figure 5 below shows the regions and sites that got their servers decommissioned.

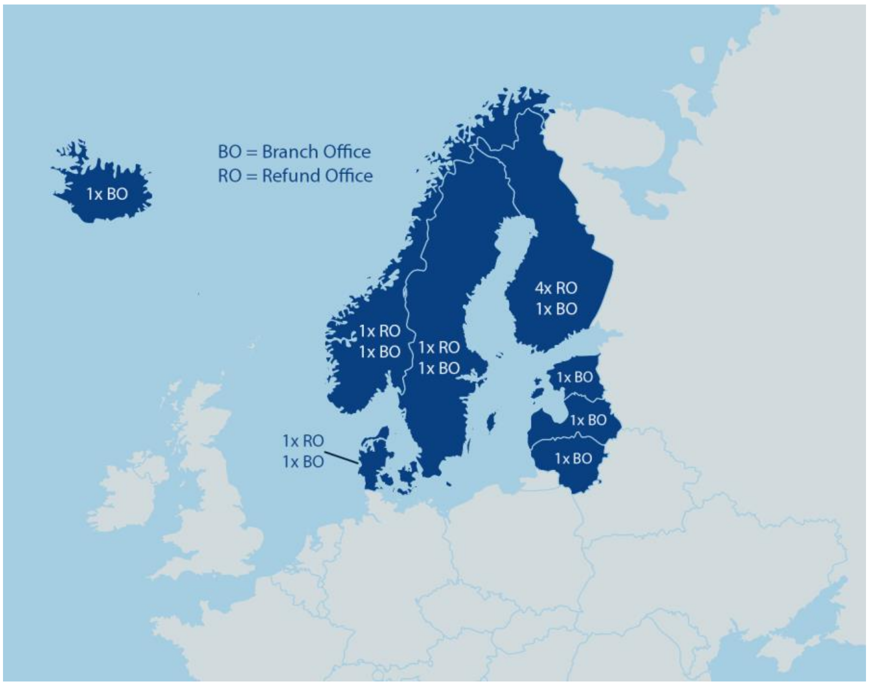

Figure 5. Affected sites with their servers.

Branch offices often held more data and had different, more complex server configs than the refund offices. While there were differences between each site, all branch and refund offices followed roughly the same baseline, with each having unique changes or additions based on business, technological or legal reasons.

The general baseline was that offices had two physical servers, both running as Hyper-V hosts, with the other one taking a primary role and the second one a backup role. There was a constant replication between the servers, so if the primary server malfunctioned and got offline, the backup server would have the identical setup and would take the primary role.

Hyper-V was used to host two main virtual servers; one for files, prints, and configuration manager (CM), the other for Domain Host Control Protocol (DHCP), Domain Name System (DNS), and Read-Only Domain Controller (RODC), with refund offices, also having a third virtual server for an application used for refund. The figure below illustrates this layout better.

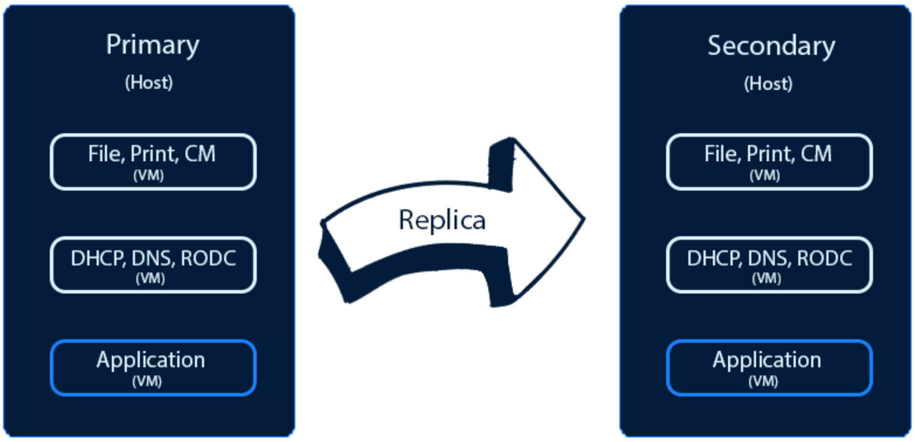

Figure 6. Generic branch and refund office server structure.

Due to COVID-19, it was better noticed that offering cloud services was much needed as employees needed to access and share data more easily regardless of their location, even outside the corporate network.

As mentioned in the previous chapter, Global Blue uses a hybrid cloud deployment model that uses the benefits of both private and public cloud models. The private portion is hosted in Global Blue's distributed data centers, whereas the public cloud portion is hosted in Microsoft Azure.

Some services were easier and better to migrate to the public cloud, like some files and applications, whereas others were chosen to keep within the organization – in the private cloud. The figure below illustrates this better.

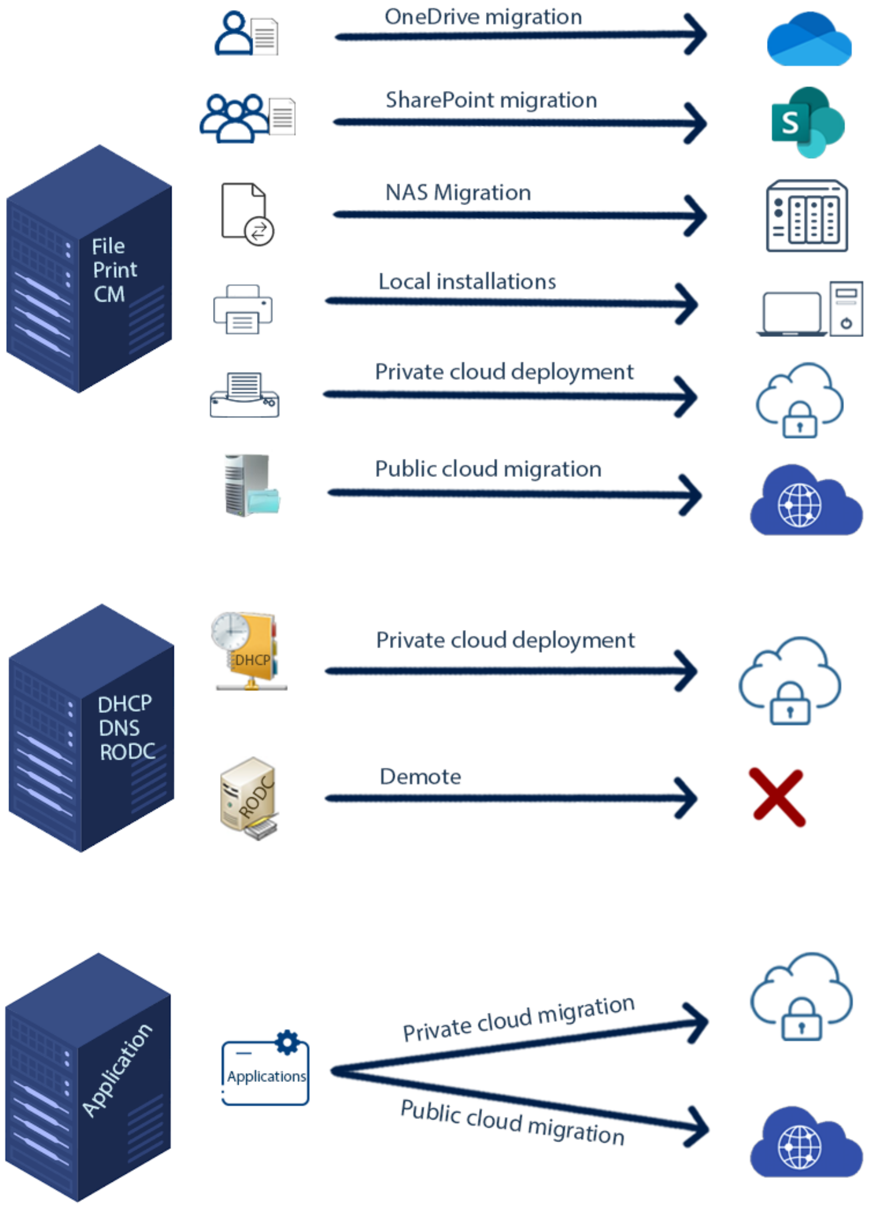

Figure 7. The serverless project plan.

Each of these services had plenty of reasons to migrate to the cloud, whether private or public. The drawbacks of the old on-premises solution, why everything was migrated as illustrated in the figure above, and some general information, will be covered next.

3.2 The Local Services

File and Print Servers

The primary tasks of file and print servers were to host all office employees' files and manage and share the local printers. These work similarly, where a share is created for the folders and the printers, which gives a network path for the users to use to access them.

The files hosted were split into team and user files. Along with the traditional home drive, the users would also either only have their Favorites or also other known Windows folders (e.g., Desktop, Documents, and Pictures) redirected to the file server via Group Policy. This was useful because the files were always accessible regardless of the computer as they were hosted in the server and not on each local device.

Refund office (RO) users shared computers within the team and often had fewer reasons to have their files in Desktop and Documents folders, so they were not hosted in the server, unlike Branch office (BO) users who were moving a lot and had their laptops, usually with more files as well. Refer to the figure below for a clearer picture of the folders and drives hosted in the server depending on the user.

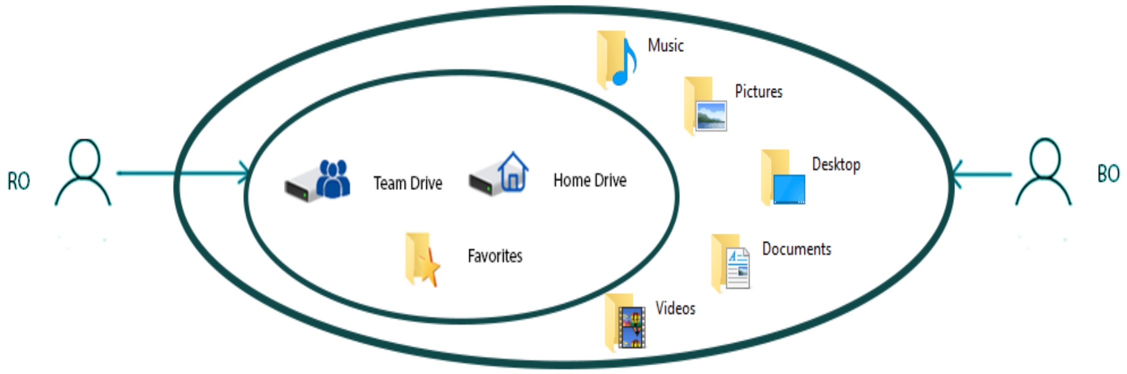

Figure 8. Drives and folders hosted in the file server based on the office type.

There were also “special” file shares like the accounting records data, periodically uploaded to the file share via a script. This accounting data needed to stay within the country for legal reasons, so it was not possible to move them to the cloud, like the rest of the data. These were planned to be moved to a NAS (Network-Attached Storage). A NAS is used to host and share files across the network.

The issue with local file servers is that the storage is very limited. Servers often triggered SCOM (System Center Operations Manager) alerts of storage capacity going below the allowed threshold. Because users need to have the ability to create, copy, or download more data, this would require periodic file clean-ups or storage upgrades to keep the storage capacity in check, but neither was always ideal. Much of the data had to be held there for archiving or legal reasons, so simply deleting everything X-years old was not a great idea either.

The second known issue was users not accessing their files from specific devices or locations even with or without a VPN (Virtual Private Network). It was noticed how much easier a public cloud approach would be to host the employees’ ever-growing data storage and access them regardless of their location or device. The team data was to be migrated to SharePoint Online, whereas user data to OneDrive.

There are several benefits to this approach. The main ones are accessibility regardless of location as well as scalability. Most users have a 1 TB size limit for their personal OneDrive and, for SharePoint, there are no size limits. There are also 15 years of data backup and versioning available, making it easy for users to go back in history to view deleted or changed data.

With SharePoint, the channel and folder permissions were more efficient to manage as the Team Owner (team managers) had the channel owner permissions. They often had the best knowledge of who should be able to access what resource.

Some services, however, were not as ideal for migrating to the public cloud. For example, Dynamic Host Configuration Protocol (DHCP) servers or service-related printers and applications were seen best to keep within the organization, i.e., in the private cloud.

A print server would host all printers and their drivers. Its main job is to share the printers to the users, point them to the drivers needed to make the printers functional, and managing them. Global Blue has two types of printing services: a regular one you'd find in most offices for general office prints and a system-specific, service-related printer for printing cheques, forms, and reports. This was mainly only found in branch offices.

Because the service-related printing was behind an application config and there were fewer (roughly 1 per branch office), it was easier to manage with a centralized print server in the data center. Regular office printer drivers and queues were constantly changing, and having a centralized server to host all company's printers would be a huge hassle to manage and update. Instead, it was found that regular printers should only be installed directly to local devices via their respective IP addresses.

There were the possibilities of using a public cloud service to manage the printers for countries that already had a specialized print server for the regular printers as well. This public cloud, however, was costlier and only for exceptional and necessary cases.

The file server also acted as a distribution point for the System Center Configuration Manager (SCCM). This role is used to host source files for client devices for software installations and updates as well as PXE boots (Preboot eXecution Environment), where the computer boots up from the network [10]. Because the SCCM repository was migrated to the cloud as part of the larger cloud migration project, this distribution point did not need migration of its own.

DHCP and RODC Servers

The second virtual server was a DHCP and an RODC server. DHCP is the protocol used to deploy IP addresses to end devices automatically. This automation process makes getting new or existing devices connected to the network easier. They are also more efficiently managed than manually setting static IP addresses, usually only given to printers, network devices, and servers.

It also hosted a local DNS server to turn IP addresses into human-readable formats, like Google.com. Global Blue already had DNS servers hosted in the data center, so the local DNS service was removed along with the RODC.

DC servers are in charge of controlling all domain users, groups, and computers. In Windows environments, these are often managed with Active Directory (AD), a network management service used to organize and manage the organization’s devices and users and permissions within organizational units (OU).

Devices that get promoted to host the Active Directory Domain Services (AD DS) are called DCs. An RODC acts as a relay of the DC server. Global Blue had deployed read-only versions of the DC to all sites. They are mainly used for security reasons as it would be unwise to set a writable DC server to each location as it would be a massive issue within the whole network if a device were to get compromised.

RODC also offers flexibility. It is sometimes deployed to a location with poorer WAN (Wide Area Network) links, so authentications could take long, especially if multiple users are at the office.

The RODC uses the DC to get replicates of the changes and for authenticating users. RODCs store cached user logins, and this means recent user logins are faster as the authentication is given by the RODC, which is hosted in the LAN, and not the DC in the data center [11]. For an illustration of how RODCs operate, refer to the figure below.

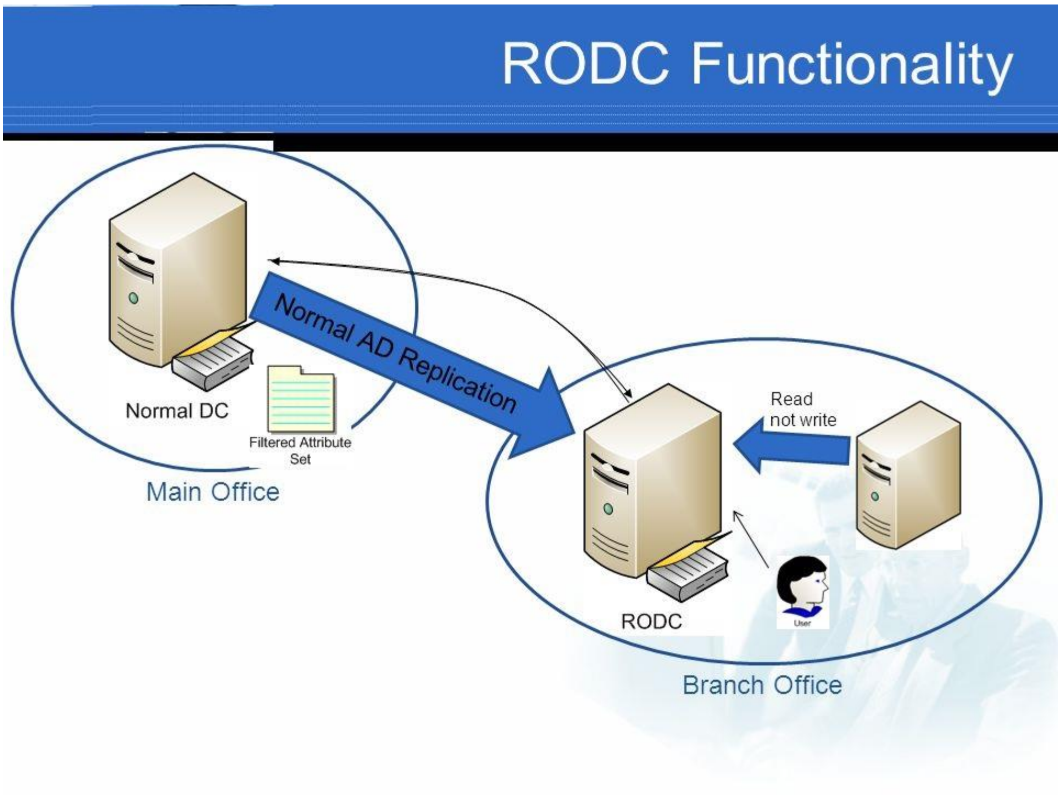

Figure 9. RODC and DC functionalities [11].

Having a DHCP server on each LAN makes sense as DHCP is a protocol generally kept within the LAN but because they can be passed through the WAN as well, having a centralized server to keep track of all locations makes it a tad more convenient and efficient for the IT team as well. Hence, the goal was to create new DHCP scopes in the DHCP servers in Global Blue's data centers and remove the existing ones hosted locally. Though the IP fetching for the devices might be slower, the users should not notice this.

The same goes for the RODC – while it is more secure and can be faster, making more frequent changes to the domain could take some time before the RODC has the newest replica, depending on the WAN links. The IT team sometimes noted that the RODCs did not always replicate quickly and reliably.

RODC was deemed unnecessary to keep after all other services were decommissioned, so all RODCs were demoted, and all clients would then always connect the primary DCs.

Application Servers

The applications were trickier as, on top of the server and application usage differed between sites, the applications were unique and needed a different approach. The most common application was used by all refund offices and agents around the world.

This used to run on a VM on each refund office, and this was a very resource-heavy option as it meant the near-identical application was running on a server of its own. They had to be manually updated every time, but they did offer the chance of offline-working, and they were faster to load and use. [12.]

A lot of applications Global Blue uses Global Blue proprietary, or otherwise difficult to migrate to a public SaaS. So, instead, they are being migrated to the private cloud and offered to employees via the Internet, like the application used for refund services. Instead of opening the app through a local server, employees now use the same app offered through Global Blue data center with unique settings and configs saved behind user accounts within the database [12].

Some countries have their unique applications running on separate servers planned to be migrated to the data center instead of a public SaaS. One application that is used in Finland is called Solotes. It is primarily used for human resource planning that includes features like staff work time scheduling, holidays, salaries, and payroll reporting.

Solotes was instead planned to be replaced with a public SaaS solution, Tamigo. It originally ran on type 2 hypervisor on a separate physical server, but as all local server infrastructure were to be decommissioned, migrating it to a public SaaS was the best approach.

Other smaller applications were run on local devices as executable files. These executable files were hosted on the local servers for employees through file shares, but after the servers were decommissioned, an alternative way for employees to access the applications was required.

3.3 The Server Hardware

When it comes to the hardware, in an ideal situation, the servers’ hardware is relatively new, their health statuses are checked periodically to make sure they are in good health, warranties are always active in case of a hardware malfunction or failure, and upgrades or downgrades are installed when needed.

After technology (hardware or software) reaches its EOL, it needs to be replaced or upgraded to keep up with the modern standards and requirements. Many Global Blue’s local servers in branch and refund offices were soon to reach their EOL, so a business decision had to be made if servers were to be kept and upgraded, or Global Blue would go for a serverless, hybrid cloud model.

The decision was made in late 2020, and the set goal was to have all required servers decommissioned by early 2023 [12]. In the meantime, the on-premises servers need to be kept operational while moving towards serverless.

Many server hardware manufacturers offer a remote server management tool to easily monitor and manage the server for checking up on the servers. Global

Blue used HP ProLiant servers, and so the server management tool used was HP Integrated Lights-Out (iLO).

Servers reaching their EOL are more likely to run into hardware failures, which were periodically checked from iLO. iLO can show which hardware parts are not functioning correctly and likely need a replacement. This can be a costly matter, and servers showing signs of retirement prioritized decommissioning.

Along with the iLO checks, Hyper-V replication status was checked weekly to lessen the chance of servers malfunctioning and losing data.

Office moves were costly and stressful from an IT point of view as the servers required additional criteria and budget increases to ensure the server hardware had an appropriate location for proper security, cooling, and power.

Essentially, the servers in on themselves required all these things, but most sites also had a KVM (Keyboard, Video, Mouse) and one or two UPS (uninterruptible power supply) for the servers alone.

The KVM is used to control and manage multiple devices, like servers or PCs. The KVMs are easy to swap between devices and are often used to check and manage the servers locally if remote means were unavailable, e.g., server start-up or connection issues.

UPS is a device used to feed the hardware power in case of temporary, sudden power outages. They are also used to ensure stable power delivery in spikes, which prolongs the hardware life expectancy. These UPS had to be kept stable and healthy as well, which takes time and money.

Depending on the UPS model, a single UPS battery alone can cost from a hundred to a thousand euros to replace [13], not to mention all the other hardware replacements required with the traditional on-premises solution, especially with older hardware. So, with servers getting decommissioned, the additional hardware was decommissioned as well. UPS and KVMs were disconnected and either reused elsewhere or disposed of.

The goal on the hardware level was that no metal was left on-premises from the infrastructure, apart from the necessary network devices or NAS.

4 Tools And Methodology

This chapter will go through the tools and resources that were available both from and outside Global Blue. It will also touch on the methods used to research solutions, gather information, and make everything progress throughout the thesis. The actual steps and the project flow will be covered in the next chapter.

The thesis was done alongside other teams within Global Blue. File transfers often required a lot of communication with the end-users as this one affected the users the most, and they had the best idea of how their file structures, especially the team files, were constructed.

While it was my job to handle the decommission steps and manage and coordinate the decommissioning processes, some key changes to specific services, like creating DHCP scopes and allowing them to pass through WAN, were other departments', like the Compute Services and Network, responsibilities. This meant that a lot of communication was required from both ends to make sure the other teams made the right changes at the right time.

ITSM Tool

Now, general office products like Microsoft Outlook and Teams were necessary for general communication, but one of the key platforms for making this collaboration between teams happen was BMC Remedy's ITSM (IT Service Management) tool. ITSM practices are used by IT organizations to better deliver more valuable IT services, which are part of the set of practices covered in the ITIL library [14].

ITSM platforms or tools like BMC Remedy are commonly associated with creating and handling incidents, service requests, or tickets in general. Though they pose many more features and tools, depending on the provider, the essential components were service requests and change management for this thesis.

Service requests are used to ask different departments to do some type of task or service. End-users might commonly use this format to request software updates or installations, but IT departments often use it to request a different IT department to do something out of their area. [14.]

Change management ensures the organization follows a set of steps for changes within the IT infrastructure to minimize issues like service downtimes [14]. The change requests were created for each site’s server decommission and for miscellaneous stuff, such as requesting additional privileges for the server decommissions and other security-related changes.

Part of the change process is the Change Advisory Board (CAB), where other IT departments review the change to see it is justified and catch any mistakes or errors [15]. The CAB played a key role in the final part of the server decommission, where the last checks and change approval are done to ensure everything has been taken care of to demote completely and shut down the servers.

Remote Tools

Another way of gathering information was via numerous remote apps. Because the servers and workstations were distributed across all Nordic and Baltic countries, getting remote access was vital as visiting the site, especially during COVID-19, was not realistic. Several remote apps were available, but only two were used for the thesis: pcvisit and Remote Desktop.

pcvisit is the remote app used to connect to users’ computers. This was essential for general IT-support tasks but the thesis as well. Testing different solutions, troubleshooting, and deploying changes related to the thesis were done on users’ workstations via pcvisit. An example of this was DHCP migration testing, where users’ computers were checked that the IP addresses were successfully obtained from the data center instead of the local server.

Remote Desktop is Microsoft proprietary remote app used to connect to Windows devices. This was massively used to connect to the servers remotely for regular server maintenances, including fixing the SCOM alerts and checking weekly Hyper-V replication status, and for the server decommission steps.

Using Remote Desktop for checking file server data structures, printer configs on the print server, and RODC replica status was often needed to get background information of the site. This worked great as a tool to allow easier server decommissions and gather information beforehand and throughout the project.

Group Policy

Another good example of using a tool for both information gathering and for applying changes was Group Policy. Group Policy is a Microsoft tool often used in domain environments to deploy a different set of computer and user settings via Group Policy Objects (GPO) without manually setting them up on each host. GPOs are a collection of different sets of policies in Group Policy. It is used in the AD environment to target different GPOs to a specific device or user groups within OUs.

Both AD and Group Policy are hosted and managed on DC servers (writable one) and changes made there are replicated to the RODCs that were locally residing in each office.

Group Policy made it possible to see the current deployment and structure with the end-users and end-devices, which network drives and types of printers were deployed to specific users, and which folders were redirected to which teams. While the GPO layouts were roughly following the same patterns in each OU

After tearing down the old GPOs, it became useful for automating new changes as well. Examples of these were enforcing OneDrive synchronizations to backup user folders as well as connecting to the respective SharePoint channels. While not as easy as the shared printer deployments, installing local printers via IP addresses were also done to users via Group Policy to some extent. More on these in the next chapter.

Documentation Sources

Windows OS (both server and workstation), AD and Group Policy, RODC, OneDrive, and SharePoint are all Microsoft products to which Microsoft provides a lot of documentation [16]. Microsoft Docs was a common place to understand the current setup more thoroughly, research new solutions, and troubleshoot obstacles.

Global Blue also uses Confluence as a knowledge base. Knowledge bases and management are within the set of practices in ITIL [17]. It provides respective employees information on different products and services owned by Global Blue, infrastructure details, third-party services, tools, and troubleshooting steps. It is maintained and updated constantly to keep the teams informed and efficient, avoiding the “reinventing the wheel” idea.

The knowledge base that is built is huge, but fortunately, it has a logical structure. It was used a lot for checking the background information on each site, such as the site infrastructure layout.

5 Going Serverless

Because the thesis covered 15 different sites across eight countries, going through all of them at once was not a realistic goal. Some sites had more complex data structures or additional setups, and some were simply more urgent than others.

The aim was to keep progressing all sites consistently, but the priority was given to sites that either had servers EOL approaching or the site had an office move ahead of them. In these cases, decommissioning the servers means users will not notice any downtime if servers start to malfunction after EOL, and it would also avoid the stress of having to worry about the additional office criteria for the servers.

Fortunately, no server reached EOL to a point where they got unstable or started to malfunction. Two Nordic sites had an office move coming, so the servers had more urgent needs to be decommissioned. Some services, however, are easier and quicker to migrate than others. That said, let’s go through each of the steps taken to decommission the servers.

5.1 File Server

The first and foremost important step for each site was to get all files migrated to the cloud (and some to a local NAS). The file storage was often the most troublesome migration step, and it was the only part that the end-users had a noticeable effect as users needed guiding to show how the data could now be accessed and shared.

One way to migrate the files straight from the file server onto Microsoft 365 was via the Migration Manager tool. It’s a Microsoft tool designed for different data migrations. It offers plenty of source and destination options, but in this thesis, the only relevant ones are from on-premises file share to SharePoint Online or

OneDrive. This requires installing an agent for the computer to access the folder share and then specifying the path where the data is to be migrated. [18.]

The OneDrive migration requires installing the migration agents on each computer, which would have required further investigating to do it efficiently. It was not as optimal nor necessary for this case, albeit a viable option. [18.]

However, this tool was not used for this thesis as it is completely managed from the SharePoint admin center, and in Global Blue, the SharePoint admin center is managed by a separate department. The Migration Manager was discovered later but is worth noting for the runbook for the consideration of other IT teams responsible for their countries' migrations.

Before migrating any data, an assessment "phase" was done where essentially the file structure was checked to see what type of files were being migrated (e.g., accounting files or applications), what the permissions look like in terms of which members of the team can access which folders, and what might not be necessary to move. Many locations, for example, had an older "IT" type of folder that was no longer relevant, so there was no point in migrating that as well.

After the assessment, it was time to start migrating the files. The flow usually started with the team data, followed by the user data, and finally, any unique setups.

Team Data

The data shared across the teams was known early on to be migrated to SharePoint Online to replace the traditional network drives. SharePoint can be used for many different team collaborations such as calendars, knowledge bases, widgets, news [19] but most teams mainly use it as a cloud file storage. The storage can be accessed via different means, such as a browser, file explorer, and Microsoft Teams. This meant any smart device could be used to access the files, so they were no longer locked behind a network drive on certain PCs.

Migrating the team data was the priority as it affected all users as a whole. The "how" the files were migrated was not necessarily important as long as the files were securely and reliably transferred. A lot of different approaches and methods were researched, tested, and used in an attempt to find the most efficient way possible. The best way to migrate files to SharePoint was through OneDrive synchronization, as not only is it quick and reliable, but it can also be done via the file explorer (as opposed to moving them through a browser) for better convenience [20].

The issue with OneDrive synchronization is that when migrating the files from the device itself, the files will reside locally in the respective SharePoint folder while being uploaded and synced to the cloud. This meant that the files would be in both the library folder and the original location. The files could be moved (instead of copied) to the library folders, but due to the complexity of some sites' permissions and file structures, it was best to keep them in the old location, not only as an additional contingency plan but also for quick access in case some settings needed adjusting.

The first approach was to archive all the team files into a zip file via 7zip, extracted them to the SharePoint library folder, and then uploaded them to the cloud via the OneDrive sync. The zip files were used to keep the file metadata, like the file creation and last-modified dates, unchanged [21].

The zipping method worked, but the servers it was tested on had smaller team data libraries size-wise. The other file servers had a lot more data in them and were low in available disk space. This meant that with zip files, the files were doubled in the SharePoint folder and original location and the zip file. The zip file would be deleted after extraction, but all three locations would still take up the space simultaneously, so a lot of extra disk space was required, and this was quickly found not to work with other servers.

One solution was to map a SharePoint library onto a network drive. This meant that the file transfer process was quicker as there were no zipped files involved, and the files were not hosted locally. But, while this was a relatively easy thing to set up in most cases, the SharePoint drive mapping was a legacy method, no longer recommended by Microsoft [20], and often required additional feature installations or configuration to support it on older servers, some of which could not support it at all.

In the end, the optimal solution was between mapping SharePoint to a network drive or migrating files through a “middle-man” computer via the OneDrive sync. These “middle-man” computers had plenty of available disk space and were within the same LAN as the servers, so the file transfer was relatively quick and seamless with proper network interface cards (NICs). But instead of copy-paste or file zipping, a Microsoft tool like Robocopy (or “Robust File Copy”) was used.

Robocopy is an enhanced command-line tool to move, purge or copy files reliably and efficiently. It poses additional features by default, such as date checking, task summarizations, but it can also be given plenty of parameters depending on the needs. One can include or exclude empty folders, permissions, or specific files/folders, encrypt files, error-checking, preserve file metadata and specify how many threads to use. [22.]

Threads are virtual CPU cores, which can be used to perform computing tasks in parallel. The more threads one specifies, the more files the computer attempts to transfer in parallel [23], making Robocopy a lot quicker to execute file transfers [22].

The date check feature was vital as some sites had complex team data structures where some files were already partially migrated to SharePoint by the respective team members. This meant some of the data resided in both old and new locations, and it was somewhat random which folders were used in the old and which in the new location. Therefore, having Robocopy automatically compare the file modification dates and overriding the older files with newer ones saved a lot of manual work.

Before migrating the files, users were notified that they could not access the team files while the migration occurred. This is to prevent anyone from using or writing data in the old location. The migration was usually done early in the morning, so few, if any, were affected.

This is where the Group Policy came in to tear down the older methods like disconnecting the respective team network drives. Sometimes, however, it was also necessary to remove the folder share permissions altogether as users working remotely might not get these Group Policy changes applied to their computers in time.

After the files were successfully migrated to SharePoint, it was just a matter of getting users to use it. As mentioned before, the files can be accessed via MS Teams and browser (which were the default ones) but also via file explorer through OneDrive sync.

The file explorer option was what users were most familiar with, and so the SharePoint libraries (that the user was part of) were synchronized to a folder view via OneDrive.

There was the possibility of mapping the SharePoint libraries to users as a network drive as well so they would feel little difference from the previous file server network drives. Though, this solution, on top of being a legacy method, prevents key cloud features like file collaborations where multiple users could see the file edits in real-time. It also prevented the files from being used offline as the network drive always required an Internet connection [19].

As mentioned earlier, the SharePoint map driving would have required additional installations or configurations, which could have been possible via Group Policy but would have required a period SharePoint website visit would have been required to keep the connection working, not to mention the removal of IE would have also called for readjustments to the GPOs. [19.]

The clear solution for this was the OneDrive sync feature. This works nicely, but due to the number of users and libraries, setting it up for each can take a toll on the IT team's resources. Fortunately, Group Policy can be used to automate this process as well. Something to keep in mind with this one was that Group Policy should not be used for libraries with more than 5000 items or applied to more than 1000 devices per library [24].

Group Policy has an Administrative Template for Microsoft 365 services, including OneDrive. These templates have a set of policies for easier management and automation of OneDrive settings. The only things required for the domain users and computers to work with Azure AD (AAD) are the Microsoft 365 tenant ID found in Azure or SharePoint and the Administrative Templates imported to the DC [24].

Each OU can have its own GPO to deploy these OneDrive settings. Many settings are available, and using them depends heavily on the IT department, such as bandwidth restrictions, automatic logins, start-ups, etc. See the figure below for a complete list of OneDrive settings available on the computer side.

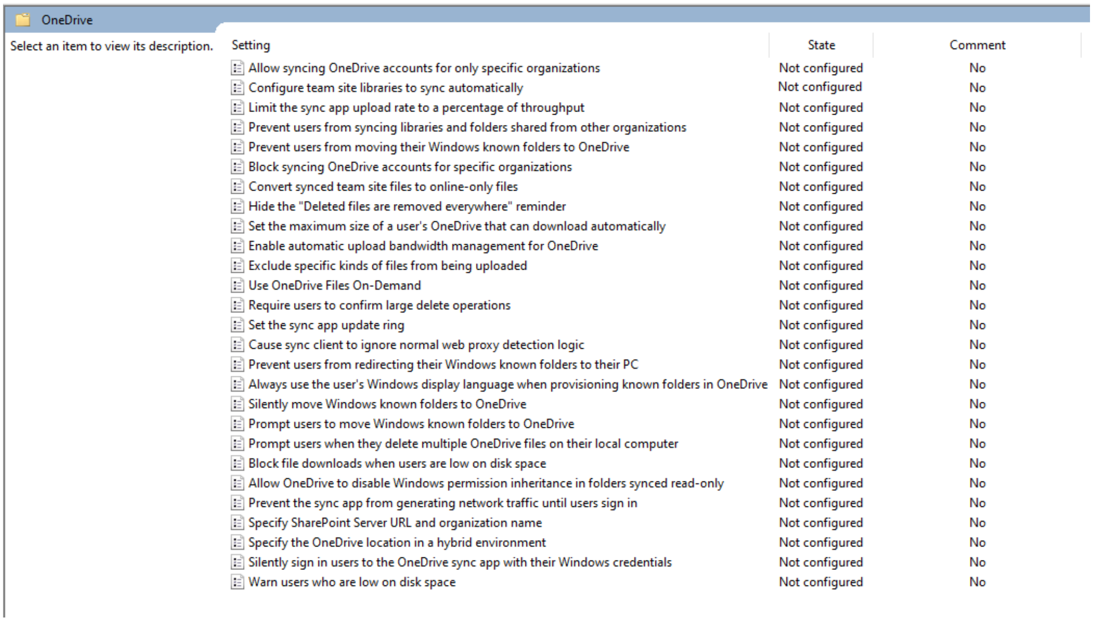

OneDrive

Select an item to view its description.

Setting	State	Comment
Allow syncing OneDrive accounts for only specific organizations	Not configured	No
Configure team site libraries to sync automatically	Not configured	No
Limit the sync app upload rate to a percentage of throughput	Not configured	No
Prevent users from syncing libraries and folders shared from other organizations	Not configured	No
Prevent users from moving their Windows known folders to OneDrive	Not configured	No
Block syncing OneDrive accounts for specific organizations	Not configured	No
Convert synced team site files to online-only files	Not configured	No
Hide the "Deleted files are removed everywhere" reminder	Not configured	No
Set the maximum size of a user's OneDrive that can download automatically	Not configured	No
Enable automatic upload bandwidth management for OneDrive	Not configured	No
Exclude specific kinds of files from being uploaded	Not configured	No
Use OneDrive Files On-Demand	Not configured	No
Require users to confirm large delete operations	Not configured	No
Set the sync app update ring	Not configured	No
Cause sync client to ignore normal web proxy detection logic	Not configured	No
Prevent users from redirecting their Windows known folders to their PC	Not configured	No
Always use the user's Windows display language when provisioning known folders in OneDrive	Not configured	No
Silently move Windows known folders to OneDrive	Not configured	No
Prompt users to move Windows known folders to OneDrive	Not configured	No
Prompt users when they delete multiple OneDrive files on their local computer	Not configured	No
Block file downloads when users are low on disk space	Not configured	No
Allow OneDrive to disable Windows permission inheritance in folders synced read-only	Not configured	No
Prevent the sync app from generating network traffic until users sign in	Not configured	No
Specify SharePoint Server URL and organization name	Not configured	No
Specify the OneDrive location in a hybrid environment	Not configured	No
Silently sign in users to the OneDrive sync app with their Windows credentials	Not configured	No
Warn users who are low on disk space	Not configured	No

Figure 10. OneDrive computer settings available in Group Policy.

The GPO can also be used to add the group’s respective SharePoint libraries to users. Some users are part of SharePoint channels that others in the same OU are not, so this would not have the desired fully automatic effect.

After researching and testing, it was noticed that one GPO could be used to control them all. All the desired SharePoint libraries can be added behind a single GPO for easier and efficient management. Only the libraries the user had access to will be added to the user. This meant that one GPO alone could manage the entire Nordic and Baltic region’s SharePoint libraries, along with the generic OneDrive setting policies.

There were also some applications being hosted in the team files, as mentioned in chapter 3. Most of these files were simply migrated to SharePoint or downloaded directly to the users’ local devices.

One group of applications used by the refund offices in Finland were challenging to migrate to the public cloud due to the nature of how the application and database were structured. These were best-kept on-premises but moved from the file server to a separate computer or NAS. Some applications were running on the server itself will be covered later in the Application Server chapter.

Other things could be easy to miss that would affect the uses of some services that the users are custom to. Apart from the smaller applications, services like scanning affected the users. Employees were used to scanning important documents back to the file server, and so all said offices needed a Scan to SharePoint -setup on their devices.

User Data

Usually, after the team data was successfully migrated to SharePoint, the user data was up next. As mentioned in chapter 3, the user data hosted in the file servers differed based on the country and type of office (branch or refund). In all cases, the first step was to download the redirected files from the file server onto the local device before migrating to OneDrive.

For refund office users, the goal was only to get Favorites-folder and everything in their Home-drive to OneDrive and guide them to start using the OneDrive folder as a place to store important files that they want to share or backup. This was due to the smaller storage limit.

For branch office users, the three known Windows folders (Documents, Desktop, and Pictures) were planned to be migrated to OneDrive along with the Favorites folder and everything in Home-drive.

Likewise, with team data, the user data was hosted in local file servers in each location, and the files were accessed via paths determined in Group Policy. The data would be migrated to Microsoft 365 environment, and the users can therefore easily access their files regardless of their device or location. User data migration was trickier as some bigger obstacles needed tackling, all relating to users’ way of working due to COVID-19.

Branch office users had practically all their data redirected to the file servers, and while it was easy to revert the file redirection change in Group Policy, the change would not take effect when users were not in the corporate network. This is because of how Windows caches information, how (RO)DCs and Global Blue VPN works [25].

As mentioned in the theory chapter, all local user authentications go through the RODC. When users log on, they normally authenticate to the nearest DC, and along with the authentication approval, the DC sends the updated version of the Group Policy. But, because branch office users work remotely, they need to use a VPN to connect and access company resources, including the RODC. [25.]

The VPN, however, has a 2-step authentication, and users can log in to it only after they have logged in to Windows. This means that no VPN connection is available when users log on to the computer, resulting in no connection to the RODC. In this case, Windows can only use the old, locally cached information [25].

This is why Group Policy changes do not update for branch office users working remotely, and inherently why disabling folder redirection does not work until user login within the company network. This was discovered much later in the thesis, and a workaround was needed for users who could not go to the office before server decommission.

The known Windows (or user shell) folder paths are determined in Registry. These folders were redirected to the file server behind the user profile folder. These folder paths needed to be manually changed back to the local device. See the figure below to see what the tool looks like after the folders have been moved back to the local device.

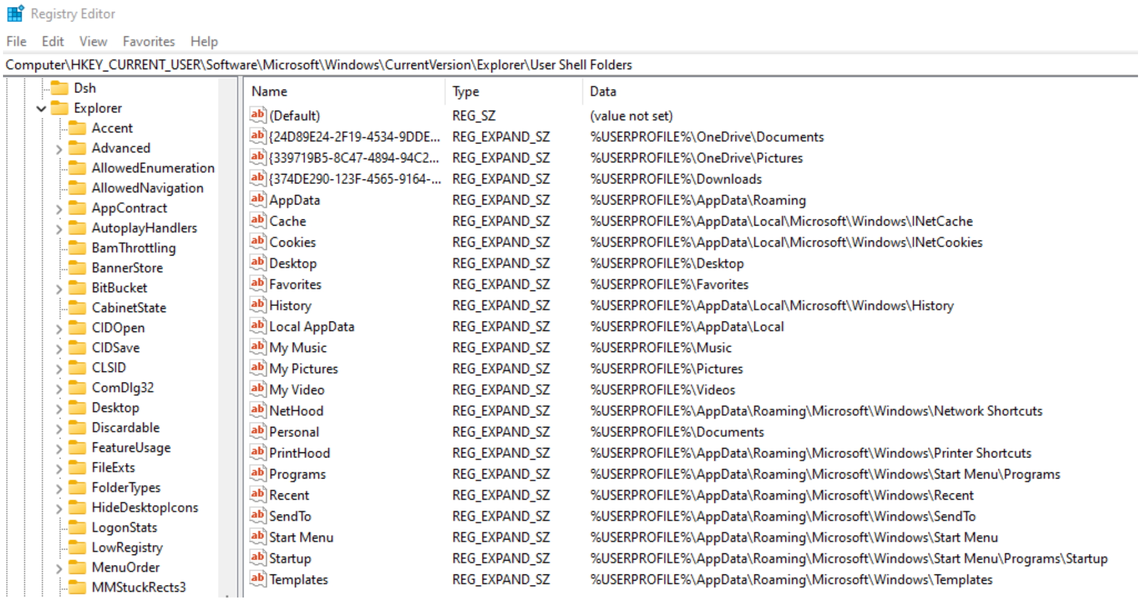

Computer\HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders

Name	Type	Data
(Default)	REG_SZ	(value not set)
{24D89E24-2F19-4534-9DDE...	REG_EXPAND_SZ	%USERPROFILE%\OneDrive\Documents
{339719B5-8C47-4894-94C2...	REG_EXPAND_SZ	%USERPROFILE%\OneDrive\Pictures
{374DE290-123F-4565-9164-...	REG_EXPAND_SZ	%USERPROFILE%\Downloads
AppData	REG_EXPAND_SZ	%USERPROFILE%\AppData\Roaming
Cache	REG_EXPAND_SZ	%USERPROFILE%\AppData\Local\Microsoft\Windows\INetCache
Cookies	REG_EXPAND_SZ	%USERPROFILE%\AppData\Local\Microsoft\Windows\INetCookies
Desktop	REG_EXPAND_SZ	%USERPROFILE%\Desktop
Favorites	REG_EXPAND_SZ	%USERPROFILE%\Favorites
History	REG_EXPAND_SZ	%USERPROFILE%\AppData\Local\Microsoft\Windows\History
Local AppData	REG_EXPAND_SZ	%USERPROFILE%\AppData\Local
My Music	REG_EXPAND_SZ	%USERPROFILE%\Music
My Pictures	REG_EXPAND_SZ	%USERPROFILE%\Pictures
My Video	REG_EXPAND_SZ	%USERPROFILE%\Videos
NetHood	REG_EXPAND_SZ	%USERPROFILE%\AppData\Roaming\Microsoft\Windows\Network Shortcuts
Personal	REG_EXPAND_SZ	%USERPROFILE%\Documents
PrintHood	REG_EXPAND_SZ	%USERPROFILE%\AppData\Roaming\Microsoft\Windows\Printer Shortcuts
Programs	REG_EXPAND_SZ	%USERPROFILE%\AppData\Roaming\Microsoft\Windows\Start Menu\Programs
Recent	REG_EXPAND_SZ	%USERPROFILE%\AppData\Roaming\Microsoft\Windows\Recent
SendTo	REG_EXPAND_SZ	%USERPROFILE%\AppData\Roaming\Microsoft\Windows\SendTo
Start Menu	REG_EXPAND_SZ	%USERPROFILE%\AppData\Roaming\Microsoft\Windows\Start Menu
Startup	REG_EXPAND_SZ	%USERPROFILE%\AppData\Roaming\Microsoft\Windows\Start Menu\Programs\Startup
Templates	REG_EXPAND_SZ	%USERPROFILE%\AppData\Roaming\Microsoft\Windows\Templates

Figure 11. Paths for known Windows folders in Registry Editor.

This approach, however, does not automatically download the data as the Group Policy change would. Instead, the users' folders will all be empty because there is nothing on the local device as all the files are still in the file server. Fortunately, the users' folders in the file server can be accessed separately and mapped to a network drive like any shared folder. Then a Robocopy script can be run to copy the files from the file server onto the folders found in the local device.

The second obstacle had to do with some users who were on vacation or leave and thus were not actively logging in to their computers. The Group Policy changes were only possible to do when the user was logged in to the computer, and so backups of the users' files were made to other locations, such as SharePoint or a server that was not being decommissioned yet.

The first step to migrate the user files was to adjust the existing folder redirection GPO, found in the User Configuration -> Policies -> Windows Settings -> Folder Redirection, and ensuring that on policy removal, the folders will be redirected back (see figure below). By default, it is set to leave the contents.

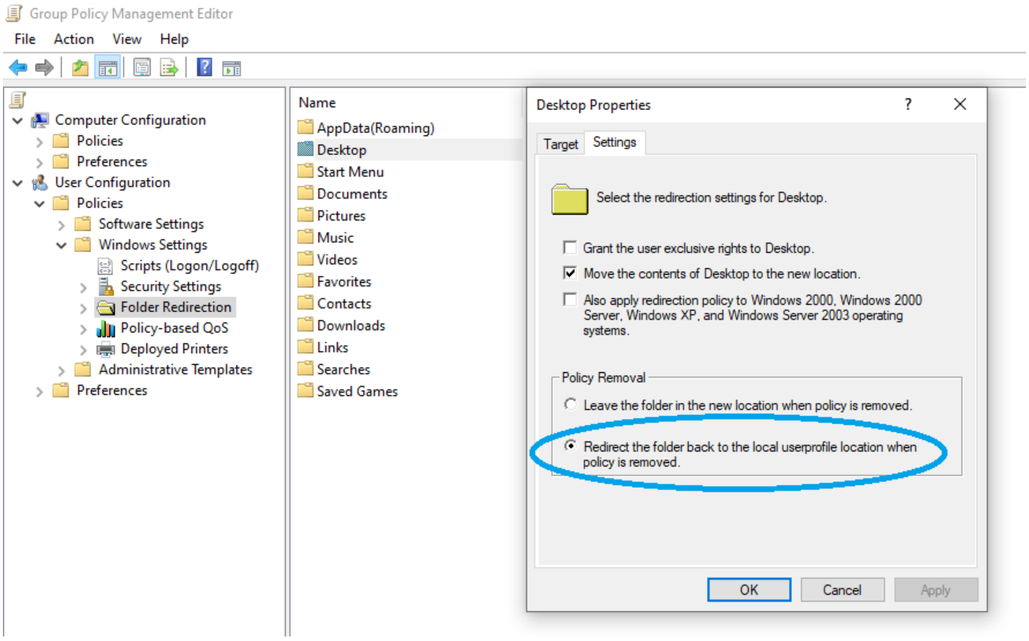

Desktop Properties

Target | Settings

Select the redirection settings for Desktop.

☐ Grant the user exclusive rights to Desktop.

☑ Move the contents of Desktop to the new location.

☐ Also apply redirection policy to Windows 2000, Windows 2000 Server, Windows XP, and Windows Server 2003 operating systems.

Policy Removal

○ Leave the folder in the new location when policy is removed.

● Redirect the folder back to the local userprofile location when policy is removed.

Figure 12. Policy Removal setting in Folder Redirection GPO.

This was crucial as the folder paths would be changed to the local folders without this step, but no files would be downloaded from the file server [26]. After that, the GPO can be disabled from users, and the next time users log onto the computer with the updated GPO, their files will be downloaded locally during logon.

Something to note here was if users would have had more data in the file server than the local device had available on the disk space, this method would not have been possible. In these cases, the Registry method would have sufficed with only parts of the data being migrated to OneDrive at once. Fortunately, this was never something to worry about, but it’s good to keep in mind if attempting this method.

After the user files were downloaded from the file server to the local device, the next step was to backup it all to OneDrive. This essentially meant that these files were hosted locally as well, making them also accessible offline.

Favorites and files in the Home drive were copied to OneDrive using Robocopy to keep the file metadata intact [22], similar to the team data migration to SharePoint.

This also meant using OneDrive’s backup feature for branch office users as it was both quick and simple to set up. See an example picture below.

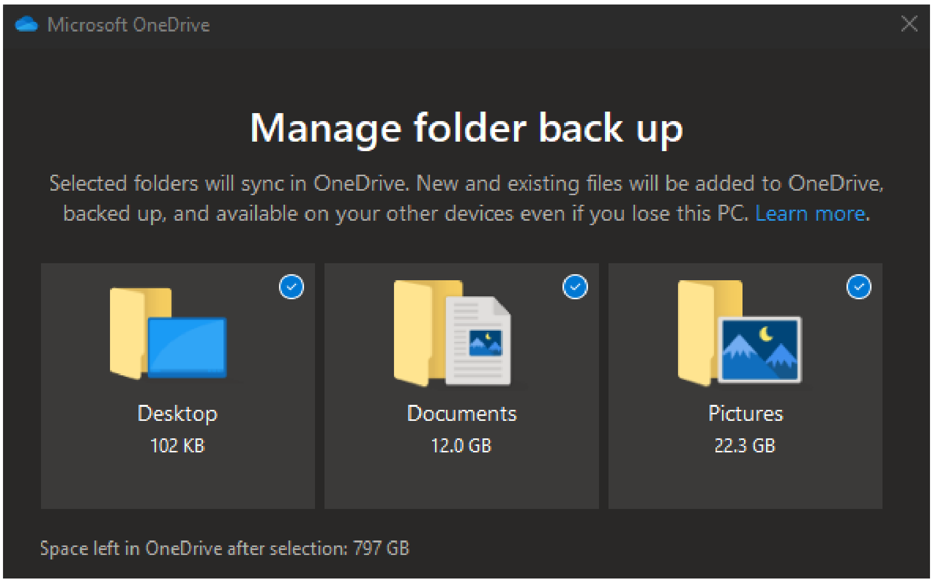

Microsoft OneDrive

Manage folder back up

Selected folders will sync in OneDrive. New and existing files will be added to OneDrive, backed up, and available on your other devices even if you lose this PC. Learn more.

Desktop 102 KB

Documents 12.0 GB

Pictures 22.3 GB

Space left in OneDrive after selection: 797 GB

Figure 13. Example of OneDrive folder backup window.

The status for each user was always checked to ensure the files were successfully migrated to their OneDrive but using the folder redirection and OneDrive GPOs, as mentioned earlier, fully automates the entire process. The known Windows folder backup is behind the policy “Silently move Windows known folders to OneDrive” [24].

User files were automatically downloaded to the local device during login, and after login was done, the OneDrive GPO ensures OneDrive is started and logged on automatically, after which it will start backing up the folders to OneDrive.

Accounting Records

Not all data were possible to migrate the same way as user and team data. The accounting records in some file servers had unique use-cases and restrictions that forced a different approach. The users did not use this data but were simply kept on local servers as a clone due to the country’s legal reasons.

Norway, Iceland, and Latvia were the only countries out of the eight that had a legal obligation to keep the accounting records within the country [27, 28, 29].

Only Norway could host the data at another Nordic country [29], but it was chosen to be kept in Norway anyway, for simplicity’s sake.

This posed an issue because all other files were migrated to the public cloud and the region for Global Blue’s SharePoint is in the Central Europe region. The Microsoft 365 services, like SharePoint and OneDrive, run in Azure infrastructure [30], and these data centers are neither found in all three countries [9] nor would it be ideal for each country to have its own region storage database. The remaining ideas revolved around physical computers hosting the data on-site.

A NAS or a general spare computer capable of hosting the data and supporting remote connections were ideal for not only a budget’s sake but also for convenience. A NAS is generally used for such cases, but because the accounting data is relatively small in size and very infrequently accessed, purchasing a NAS was unnecessary as there were capable, spare computers around that were already in the domain and easily accessible via remote access tools.

To prepare the computers, they were first wiped clean via Windows reinstallation or PXE boot, and the latest updates were installed. A folder share was created to host the upcoming accounting data. This share needed read/write permissions to the user accounts that were running the file transfer scripts. Without the permissions, the script would not access the share or add any files there.

Next to the What, Where, and Why, we have How. The log files were no simple text files, nor did they use a basic script to simply copy the files from server A to B. Global Blue uses a technology called Managed File Transfer (MFT) to copy and move accounting data to other servers (like the local servers) via scheduled tasks.

MFT is a more reliable, efficient, and secure way to transfer data as opposed to common protocols like HTTP or FTP (File Transfer Protocol). It has security functionality like SFTP (Secure File Transfer Protocol), but MFT can protect the data even outside transit (i.e., at rest) and offers a broad selection of tools crucial for transferring this type of data. This includes (but is not limited to) alerts and reports of the status of transferred files as well as easier task automation and scheduling. [31.] The list of benefits is big, but the key differences can be summarized as shown below.

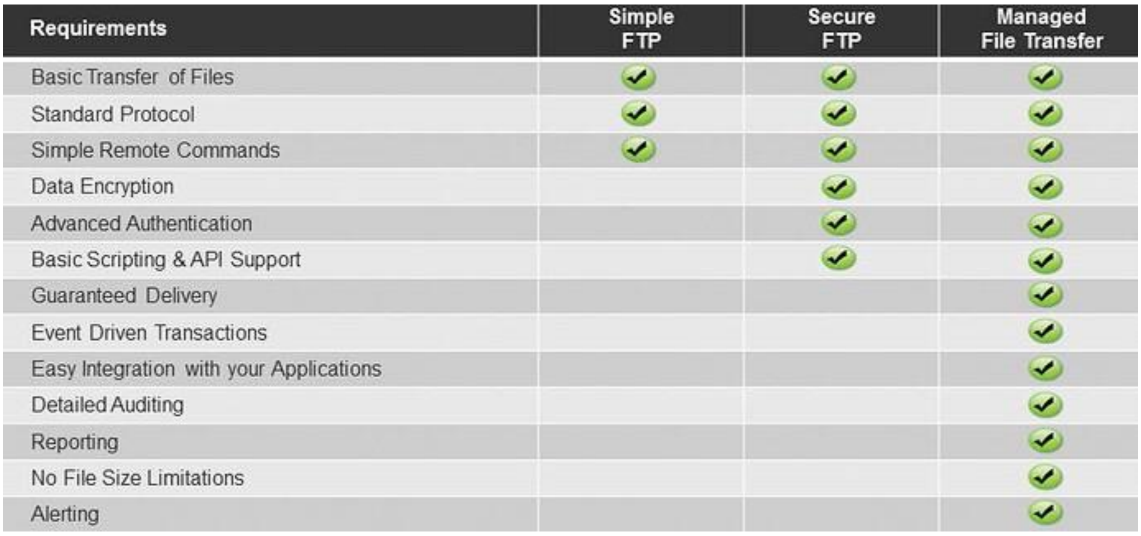

Requirements	Simple FTP	Secure FTP	Managed File Transfer
Basic Transfer of Files	✓	✓	✓
Standard Protocol	✓	✓	✓
Simple Remote Commands	✓	✓	✓
Data Encryption		✓	✓
Advanced Authentication		✓	✓
Basic Scripting & API Support		✓	✓
Guaranteed Delivery			✓
Event Driven Transactions			✓
Easy Integration with your Applications			✓
Detailed Auditing			✓
Reporting			✓
No File Size Limitations			✓
Alerting			✓

Figure 14. Differences between MFT, FTP and SFTP [32].

After the computers were set and file share created with proper MFT read & write permissions, the Service Center, responsible for the MFT tasks, were requested to make a clone of the old task via the ITSM tool to clone the same task (with the same schedule and properties) to the new folder share location.

After confirming the new accounting data was successfully being copied to the new share, the original MFT task could be disabled, and at this point, the “special data” and the whole file server portion were completed.

5.2 Print Server

The office printers were split into service-related and normal office printers. Both regular and service-related printers were hosted and shared from the same local print servers except for most refund offices that only had the regular printers.

The regular printers were agreed not to migrate to the cloud, as opposed to the service-related printers that were migrated to the Global Blue data center. The goal was to install the printer drivers/queues on each device via TCP/IP, i.e., using the printer's IP address as opposed to the original share path.

For some sites, however, there were options to migrate the printing to a cloud service. There are plenty of SaaS options to choose from, such as PrinterLogic, designed for serverless-approach and would allow printers to be managed, deployed, and used via the platform [33]. See the PrinterLogic migration flow below.

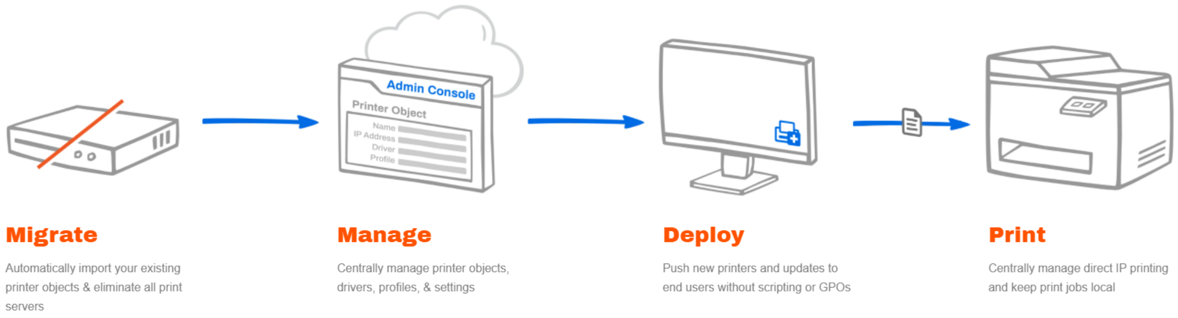
Admin Console
Printer Object
Name
IP Address
Driver
Profile
Migrate
Automatically import your existing printer objects & eliminate all print servers
Manage
Centrally manage printer objects, drivers, profiles, & settings
Deploy
Push new printers and updates to end users without scripting or GPOs
Print
Centrally manage direct IP printing and keep print jobs local

Figure 15. PrinterLogic SaaS solution for cloud-based printing [33].

The SaaS option was for countries with a special need for a centralized, managed printing system, but for this thesis, the plan was to only install regular printers locally to each device.

The printers would be installed from Control Panel and selecting Add Printer -> TCP/IP. So, while the process of installing printers to devices is quick and straightforward, this was a time-consuming process to do for all clients and printers. A more efficient solution was needed.

As mentioned before, Group Policy was used to deploy the shared printers. After the servers were being decommissioned, Group Policy was planned to be used to install the printers via TCP/IP as well. Though this solution did not work as intended because it requires a UNC (Universal Naming Convention) path to the driver distribution point from where the drivers could be installed [34] and without a server (or separate computer to host them) this was not possible.

Fortunately, there are workarounds to this. An easy solution would be to have a stationary computer within the office to share the printers as the server would have, but a basic computer might not be able to reliably deliver the print services that bigger offices might require. This is an option for the future, however.

Group Policy could also be used to install the printers, not via the general printer deployment policy but via a PowerShell script, ran on logon. PowerShell has plenty of modules that can be used to do various tasks straight from the command line – one of which was the PrintManagement that can be used to install, manage, view, or delete printers installed on local or remote computers [35] and these commands were used to create a script file to install office printers. An example script is shown below.

```
Add-PrinterPort -Name “192.168.100.100” -PrinterHostAddress “192.168.100.100”
Add-PrinterDriver -Name “HP Universal Print Driver”
Add-Printer -Name “HP M404n - Front Office” -DriverName “HP Universal Print
Driver” -PortName “192.168.100.100”
Invoke-CimMethod -InputObject $(Get-CimInstance -Class Win32_Printer -Filter
"Name='HP M404n - Front Office'") -MethodName SetDefaultPrinter
```

Listing 1. A PowerShell script structure to install a network printer.

As seen on the listing, the TCP/IP port is created along with the driver. If these already exist, then they will simply fail, and the script will continue. After that, the printer itself gets installed with the given name, port, and driver, and the last line makes the printer default – this is simply to avoid potential IT-support cases.

The script was deployed to the selected users via GPO targeting as a scheduled task after logon and was set to run only once. This solution does not work for every situation but is intended to lighten the task of installing common printers to existing computers in bigger offices.

There were cases where the drivers in the computers’ driver store did not work with the local printer, or there were printers that only some users were allowed to use. Both cases were rare, and in these cases, the printers were simply installed manually.

The service-related printing in itself is not anything too out of the ordinary but rather the backend side and purpose of it is. All branch offices had an additional printer for printing cheques and form documents. These types of prints were done via Global Blue’s system.

Unlike the normal office prints, this printing was more standardized and important, so the decision was made to have a centralized printing server in the private data center to manage all of them. This way, they are a lot easier to manage and troubleshoot should there be any issues with the service, either on the application or the printing side.

These printer queues were re-created in the data center’s print server using the printers’ IP addresses. Some adjustments to the printer queue properties were needed as, even though the system used to print was standardized, the printers themselves slightly differed based on the country.

The printer queues, once installed, were all given a logical name based on the printer's physical location, and a test print was conducted. If the test print (through printer properties window) worked, this meant the printer queue was operating properly and the only things left were making the change to the system and asking respective users to make a proper test print.

A different department is responsible for maintaining the company’s in-house system. Collaboration was needed via the ITSM tool, where the department was given the required information regarding the new printer queue, such as the location, old queue path, and the new path, and they would apply the changes.

After this, users were asked to test if the documents were still printed and had the proper, necessary pre-set data. After successfully printing the desired forms or cheques, the printer server portion was completed.

5.3 DHCP Server

Much like the service-related print services, the DHCP servers were migrated to the private cloud, to the Global Blue data center. There was the possibility of letting the local router act as the DHCP server if there were any restrictions for having the DHCP at the data center, but routers, in general, are not meant for this as it requires additional processing, which may affect the network. For this thesis, all local DHCP servers were migrated to the Global Blue data center.

The DHCP portion required a lot of collaboration with other departments. The DHCP servers, while accessible via local IT, were managed by Compute Service team.

The first step was to obtain the current DHCP setup by accessing the local DHCP server and noting down the network and gateway addresses as well as the IP range, including any excluded addresses from the DHCP service tool and the DNS addresses from the command prompt. These were given to the team by creating a service request using the ITSM tool.

The DHCP scopes were then done, but no DHCP packets were going there because previously there was no need to allow DHCP packets from the router to the WAN. When the client sends a DHCP request to ask for the DHCP server address, the request sent is a broadcast, which means they are distributed to everyone in the same subnet. This is why by default the routers (or level 3 switches) will not send the packets outwards.

This change required some firewall and network configurations from the Network team. The firewall needed to whitelist the DHCP traffic from a certain network to the DHCP server. After that, the local router needed a DHCP relay agent role (done via *IP helper-address* command), which means it will forward the requests to the destination, i.e., the DHCP servers [36]. These changes were also requested via the ITSM tool. See the figure below for an illustration.

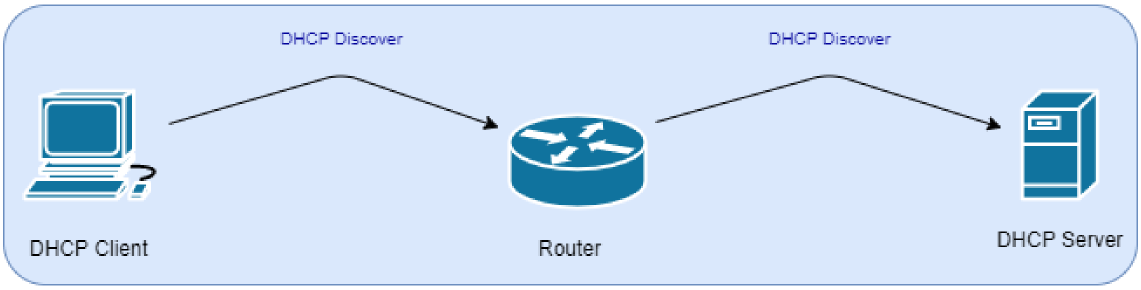

Figure 16. Local router as the DHCP relay agent.

After the Network team reported back that the configuration was done, the DHCP service could be shut down from the local DHCP server, and any client or device within the network could be used to test the new network change. This could be done via either a device restart or via command prompt using *ipconfig /release* and *ipconfig /renew* to give up the DHCP lease and request a new one, respectively [37]. This is where the pcvisit remote tool, for example, became handy.

It is promising to see if the device can connect to the network, but to ensure there are no static IP addresses or any old connections, it is best to run the *ipconfig* command again and see the DHCP and DNS addresses point to the servers in the data center. In this case, the DHCP portion could be checked from the list.

5.4 Application Server

Global Blue uses a lot of different applications for different business services. Several applications were found in the local servers, but the most common one that all refund office servers had was a Global Blue proprietary application used for refund services. This application ran on its own virtual server in each refund office and was migrated to a web-based service hosted in the Global Blue data center [12].

Previously the refund office computers would run this as a desktop application straight from the local servers. The server was migrated to a web service – using browsers to access the application from the data center, which also hosts the configurations and databases. [12.] This migration was done before this thesis but was essential to know as the old application servers still existed in Hyper-V.

What was included in this thesis was the staff scheduling system, Solotes, hosted in the Finnish branch office server. It’s running on a VM and on separate hardware from the rest of the VMs. Instead of the regular Windows Server and Hyper-V setup, Solotes runs on Windows 7, hosted on VMWare hypervisor.

Because the application and its database run fine on the workstation, the VM could move to a physical computer with upgraded OS and hardware. VMWare offers tools to convert the VM image into a physical, normal OS image [38]. Which would mean the application server would simply change on the infrastructure level.

There were pros to this, such as being able to more easily migrate it from the server, allowing the decommission of all physical servers, and end-users would not be affected by this, nor would they need to learn to use a new application.

However, it would have required worrying about the computer hardware to make sure it could manage the storage, bandwidth, and processing power needed to host the application and database within the WAN. Contingency plans would have also been needed in case of hardware or power failure.

So, the plan was to decommission the server and migrate the staff scheduling service to a SaaS from Tamigo. This later turned into a separate project from this thesis. The SaaS approach would mean Global Blue does not need to worry about the hardware or platform where the application and the database are hosted.

Tamigo is a more modern application, used as a SaaS, that offers more features [39] and is already used by the Scandinavian countries, so having this centralized application for Finland as well was seen as the better option.

This migration required a lot of planning to clone the entire employee base and configuration as well as the concern regarding the large and old database of all the previous records. The database cannot be imported from Solotes to Tamigo, requiring additional planning for how long and where the old data from Solotes needs to be kept.

This is a large project in itself due to planning with various Global Blue and third-party departments, budgeting, and actual implementation with proper configurations, so, unfortunately, it could not be completed in this thesis.

Every application is different, and when considering the application migration from on-premises, many variables need to be considered. The refunding application, for example, was better kept in the Global Blue data centers as it was already more integrated into the rest of the system, and there were no SaaS alternatives for this purpose. Unlike Solotes, for example, there had more modern alternatives offered by the public CSPs, like Tamigo.

There are also alternatives to host and develop the company’s applications within the public cloud as a PaaS, but due to the sensitivity of some applications they were best kept within the company data centers.

5.5 RODC Server

The RODC server, unlike any of the other servers, was not migrated to the cloud but was demoted in its entirety instead. As mentioned in chapter 3, the RODC is a read-only version of the primary DC. They are distributed to remote locations for security and flexibility reasons, so there is no need to host RODC at the data center where the DCs are.

After demoting the RODC (and DHCP), users would both authenticate and obtain their IP addresses from the virtual servers within the private cloud.

The RODC was set up and configured in each location using the same logic, so the demote process was the same for each site. The end goal is to have a server with AD DS, DHCP, and DNS server roles removed, as shown below.

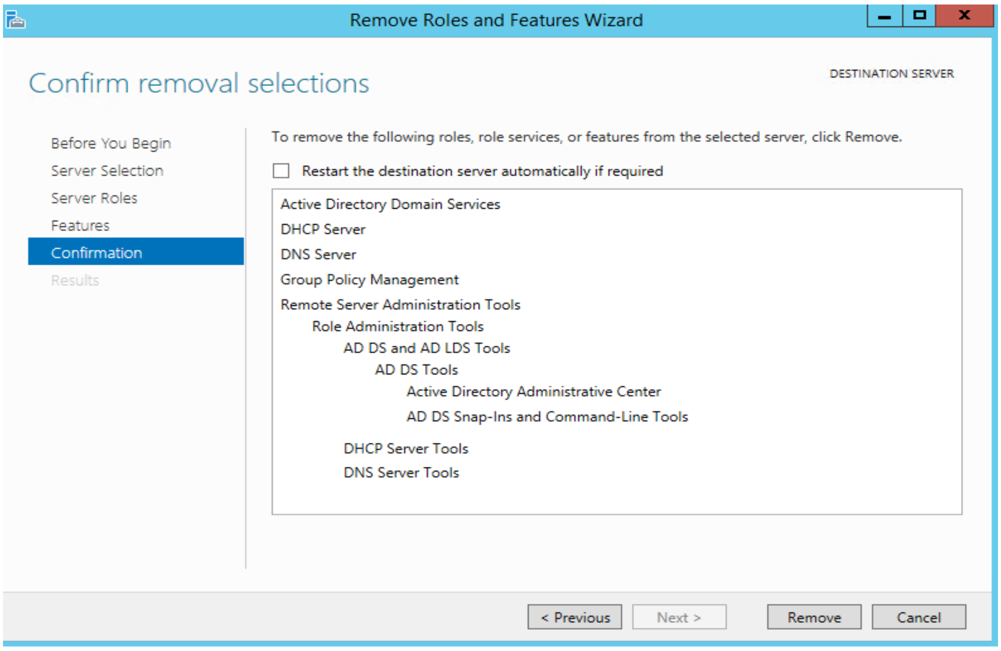

Confirm removal selections

To remove the following roles, role services, or features from the selected server, click Remove.

□ Restart the destination server automatically if required

Active Directory Domain Services
DHCP Server
DNS Server
Group Policy Management
Remote Server Administration Tools
    Role Administration Tools
        AD DS and AD LDS Tools
            AD DS Tools
                Active Directory Administrative Center
                AD DS Snap-Ins and Command-Line Tools
        DHCP Server Tools
        DNS Server Tools

Figure 17. AD DS, DHCP, and DNS server roles are being removed.

Because the DHCP service has already been disabled and migrated to the data center, removing the DHCP role is as simple as unchecking the checkboxes in Remove Roles, and Features Wizard found from the Server Manager [40].

The AD DS role, however, cannot be removed while the RODC is still in active mode. By demoting the RODC and not simply shutting the server down, the replicated changes, like user password change requests, are sent to the DC, and no metadata is left [41].

Much like with the other migration paths, there are also several ways to demote the RODC. A common way to previously install and remove AD DS roles was to use a tool called Active Directory Domain Services Installation Wizard (or dcpromo), but this has been deprecated with the release of Windows Server 2012 [42].

With Windows Server 2012, the AD DS became more integrated into the Server Manager tool. The Server Manager tool (or PowerShell) could be used to both install and remove the AD DS role with fewer steps. These are the official recommended methods from Microsoft. [41.]

While trying to find the most effective and efficient methods wherever possible, such as running a CLI (command-line interface) script via PowerShell for a quick automation process, having an easy-to-follow GUI (graphical user interface) via Server Manager was more ideal for this purpose. The whole AD DS demote process is illustrated below.

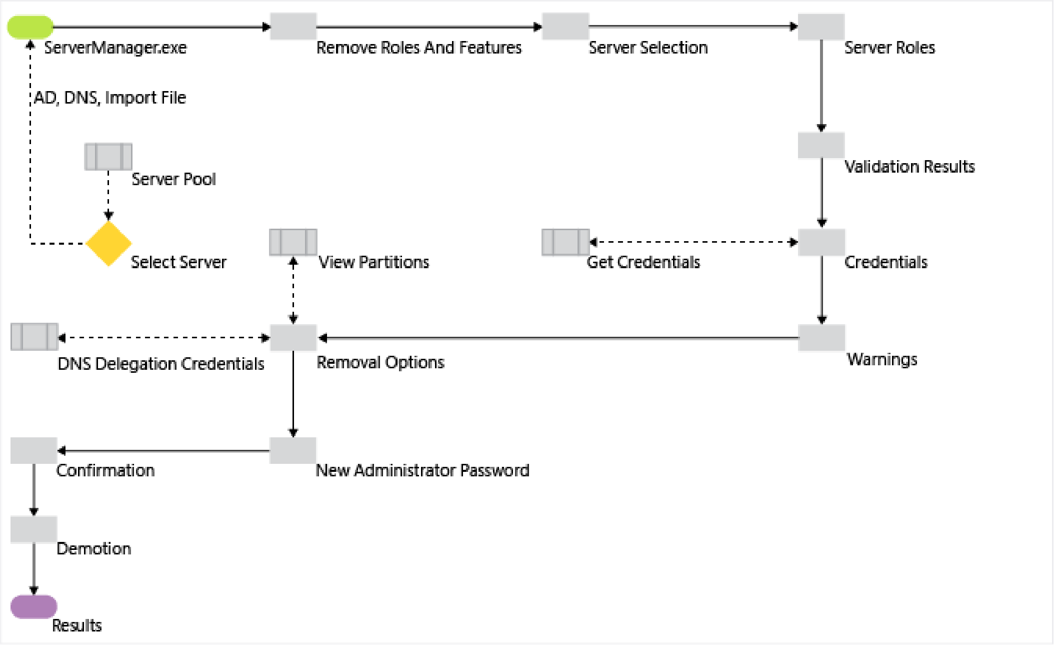

Figure 18. RODC demote process flow [41].

Starting the RODC demote process is similar to how the DHCP and DNS roles were removed: via Remove Roles and Features from the Server Manager tool and selecting the AD DS role from the desired server.

This will prompt a validation result window that says the AD DS role cannot be removed until the RODC has been demoted. As shown in Figure 18 above, there will be several confirmations and then a local Administrator password creation page for the demotion. This local Administrator account can be used to log in and remove the server roles after the demotion when the server is no longer an RODC.

The demote process has an option to retain the metadata, and it is important not to check the “retain domain controller metadata”. If there were some metadata stored, the demote would not be completed fully as it would leave remains of the RODC to the AD database. This tells the other DCs in AD DS that the RODC is still there, which causes replication errors, for example. [41.]

The metadata will also be kept if the “force the removal of this domain controller” is checked, but this option is only to be used if there are any connection issues to the DC [41]. In one of the cases, while demoting one of the RODCs, this checkbox was accidentally checked, and the metadata was kept in the AD, resulting in the RODC showing up as “unoccupied Domain Controller” in the AD.

To solve this, the metadata needs to be manually cleaned when deleting the RODC from AD Users and Computers. When deleting the RODC from the AD, it asks about resetting user passwords and exporting the cached users to a file, but because the RODC was not compromised, there is no need to reset the cached passwords. [43.]

After the RODC has been successfully demoted, the computer will be restarted, and the AD DS can then be removed from the role list along with DHCP and DNS (see Figure 17). The AD object should no longer have an RODC role and should simply exist there as a normal computer, in which case the RODC decommission has been completed.

Because this demote process requires lengthier server restarts, to avoid any DHCP service downtime or SCOM alerts, it was seen best to do the RODC decommission last – right before the servers were shut down for good. This brings us to the last part of the decommissioning.

5.6 Shutting It Down

After all virtual servers were either migrated to the cloud or decommissioned, it was time to shut everything down. The virtual servers were shut down first and then the physical Hyper-V host servers. This could be done by simply pulling the power cord but, to avoid any data losses, they were all shut down gracefully, just in case.

Before this, change requests were made via the ITSM tool for approval and to ensure everything had been taken care of. After the change request was approved in CAB, the last remaining prerequisites had to be done. The figure below shows the entire process cycle.

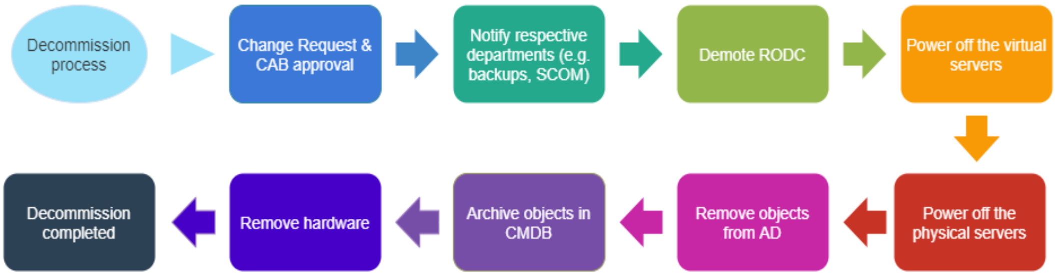
Decommission process
Change Request & CAB approval
Notify respective departments (e.g. backups, SCOM)
Demote RODC
Power off the virtual servers
Power off the physical servers
Remove objects from AD
Archive objects in CMDB
Remove hardware
Decommission completed

Figure 19. The decommission process flow.

This included informing all the necessary departments about the servers going offline, and some changes had to be made to avoid unnecessary alerts.

For example, the server monitors in SCOM had to be disabled. You can imagine how many alerts this would keep throwing when four servers (and all the services and monitors in them) suddenly became unreachable.

The second thing was the Networker backup system that Global Blue uses. It is regularly doing a file backup on the file server and needed to be disabled for the respective file server prior to going offline.

So, it is crucial to make sure all parties are well informed of the upcoming change. Some locations had additional setups (like the computers hosting the accounting records) that needed informing, as with the servers going down, the computers were required to stay up and running for legal reasons.

After the prerequisites, the steps that were remaining were pretty straightforward. As with most interactions with the servers, the shutting down process was done via Remote Desktop. The virtual servers were shut down first and then the secondary physical server.

By shutting down the secondary server before the primary one, the unnecessary replication can be prevented as, otherwise, shutting down the primary server first would initiate the secondary server taking up the primary role.

After secondary and primary servers were powered off, the devices were deleted or archived from AD and asset management. To further follow the ITIL’s practices, Global Blue uses CMDB (Configuration Management Database), a library for all the company's hardware and software [44]. Any changes made to the hardware or software needs to be updated to the CMDB. So, when servers get decommissioned, they have to be archived in the CMDB.

That was it – the only thing remaining was the hardware side. While the hardware setup was different in each site, the principal was the same: unmount the hardware from their server racks and prepare for disposal, with hard drives being kept for security and backup reasons.

Because of COVID-19, traveling to other Nordic and Baltic countries was not easy. Therefore, the end-users were guided to do these steps for sites that had an office moving coming. For other offices, the hardware will be picked up after COVID-19 allows easier travel. In any case, they were as good as decommissioned successfully.

6 End Results and the Future

There were many obstacles and setbacks that required reassessing the approach and researching better alternatives, such as with the file and printer migrations. As mentioned in the previous chapter, there were inefficient or inoperative approaches taken into consideration and even tried, but each of them was a learning step for a more successful end.

For every obstacle, the runbook was updated with an improved solution. It was, however, never possible nor the goal to achieve a runbook that would completely cover all cases because each country and office has different setups due to business, technological, or legal reasons.

Users' data and user experience were important factors to consider, and users were guided before and after the migration on how they could use the newer ways. This became even more fluent after using Group Policies so the users or local IT did not have to make the changes manually. A lot of them were already ready to use the newer approach, and many found it to be relieving.

The file services were the only "front-end" services that the user could notice, whereas the "back-end" services like printers, DHCP and DC often went unnoticeable by the users, unless there was a fault, of course.

The most complex location was the Finnish branch office that acted as the "admin center" for the Nordic and Baltic countries. A lot of different applications, databases, files, and backups were hosted on the Finland branch office, and it also had by far the most users. The servers in this branch office were saved till last as due to its "admin center" nature, it was debated a lot if the server would stay and be upgraded instead of being decommissioned like the rest were.

After thoroughly going through everything that was running and hosted on the server and documenting it, it was deemed a better, albeit slower, option to migrate to the cloud and decommission the server. Because of this not all services within these servers could be fully decommissioned within this thesis.

So, what comes to the future – in the short-term, the last remaining server that could not be finished within this thesis, like the Solotes server, will be decommissioned and migrated to the cloud later as a separate project.

The runbook written through the thesis will likely assist other IT departments worldwide in decommissioning their local servers. As more specialized cases and efficient alternatives are found, it will likely be further improved.

These local server decommissions are only a part of the bigger picture of the cloud migration process Global Blue has decided to undertake. A lot of servers and services are planned to be migrated to the hybrid model.

7 Conclusion

The thesis set out to decommission the Nordic and Baltic servers in each of the 15 offices, which consisted of various common servers found in all sizes of infrastructure, and from the thesis point of view, this has been done successfully with all said offices residing in the hybrid cloud model.

Another goal of the thesis was to get enough material and knowledge off of the Nordic and Baltic server decommissions to aid better other countries' IT departments to do the same with the help of a runbook. This has already been published and is used by many regions.

Because Global Blue uses a hybrid model, they will also have servers to host private cloud services. So, not all server hardware was removed from the company but considering there were up to 50 countries with most having local servers; it's safe to say the total investment will be worth it for Global Blue.

This might not be the case for every organization's IT infrastructure. It is also clear that most organizations have different infrastructure and different services that likely need a different set of approaches to migrate, or they might have different restrictions to not being optimal to migrate in the first place. Cloud services are becoming more popular and richer in features, but they have their cons, so having a full assessment of the pros and cons of cloud migration is essential.

Hopefully, this thesis can be of use to obtain a better understanding of what needs to be taken into consideration with the serverless approach, how some services could be migrated to the cloud and what can be the benefits.

References

1 Compute. Online. Microsoft. <https://azure.microsoft.com/en-us/product-categories/compute>. Accessed 6.10.2021.

2 Brief History of Virtualization. Online. Oracle. <https://docs.oracle.com/cd/E26996_01/E18549/html/VMUSG1010.html>. Accessed 5.10.2021.

3 Hypervisor. Online. VMWare. <https://www.vmware.com/topics/glossary/content/hypervisor>. Accessed 2.10.2021.

4 Posey, Brien. Top 11 Microsoft Hyper-V Terminologies you need to know. 2018. Online. Vembu. <https://www.vembu.com/blog/top-11-microsoft-hyper-v-terminologies-you-need-to-know>. Accessed 27.9.2021.

5 Richter, Felix. Amazon Leads $150-Billion Cloud Market. 2021. Online. Statista. <https://www.statista.com/chart/18819/worldwide-market-share-of-leading-cloud-infrastructure-service-providers>. Accessed 8.10.2021.

6 Final Version of NIST Cloud Computing Definition Published. 2018. Online. National Institute of Standards and Technology. <https://www.nist.gov/news-events/news/2011/10/final-version-nist-cloud-computing-definition-published>. Accessed 2.10.2021.

7 IaaS vs PaaS vs SaaS. 2020. Online. Red Hat. <https://www.redhat.com/en/topics/cloud-computing/iaas-vs-paas-vs-saas>. Accessed 2.10.2021.

8 What are public, private, and hybrid clouds? Online. Microsoft. <https://azure.microsoft.com/en-us/overview/what-are-private-public-hybrid-clouds>. Accessed 26.9.2021.

9 Mehta, Kiran & Singh, Shatakshi. Cloud Data Center Locations for Top 3 CSPs (AWS, Azure & Google). 2019. Online. Techyaz. <https://techyaz.com/cloud/cloud-data-center-locations-for-top-3-csps-aws-azure-google>. Accessed 1.10.2021.

10 Plan for site system servers and site system roles in Configuration Manager 2021. Online. Microsoft. <https://docs.microsoft.com/en-us/mem/configmgr/core/plan-design/hierarchy/plan-for-site-system-servers-and-site-system-roles>. Accessed 28.9.2021.

11 Kumar, Vipan. Why we need Read-only domain controllers (RODC). 2019. Online. Windows Techno.

<https://www.windowstechno.com/why-we-need-read-only-domain-controllers-rodc>. Accessed 29.10.2021.

12 Makkonen, Teemu. 2021. Area Manager Field Service CENTRAL & NORDICS. Global Blue - Administration Center North Oy. Interview. 4.10.2021.

13 UPS Replacement Battery Selector. Online. APC. <https://www.apc.com/shop/fi/en/tools/replacement-battery-selector>. Accessed 20.9.2021.

14 IT service management (ITSM). 2020. Online. IBM. <https://www.ibm.com/cloud/learn/it-service-management>. Accessed 2.10.2021.

15 Watts, Stephen. ITIL Change Advisory Board (CAB) Explained. 2017. BMC. Online. <https://www.bmc.com/blogs/itil-change-advisory-board-cab/>. Accessed 24.9.2021.

16 Technical documentation. Online. Microsoft. <https://docs.microsoft.com/en-us/documentation/>. Accessed 20.9.2021.

17 ITIL® SKMS & Knowledge Management. 2016. Online. BMC. <https://www.bmc.com/blogs/itil-knowledge-management>. Accessed 22.9.2021.

18 Overview: Migrate your file shares to Microsoft 365 with Migration Manager. 2021. Online. Microsoft. <https://docs.microsoft.com/en-us/sharepointmigration/mm-get-started>. Accessed 8.10.2021.

19 SharePoint. Online. Microsoft. <https://www.microsoft.com/en-us/microsoft-365/sharepoint/collaboration>. Accessed 20.9.2021.

20 Map a network drive to a SharePoint library. Online. Microsoft. <https://support.microsoft.com/en-us/office/map-a-network-drive-to-a-sharepoint-library-751148de-f579-42f9-bc8c-fcd80ccf0f53>. Accessed 20.9.2021.

21 Ball, Craig. Preserving MAC Times Collecting Files in E-Discovery. 2018. Online. CraigBall. <https://craigball.net/2018/06/20/preserving-mac-times-collecting-files-in-e-discovery>. Accessed 3.10.2021.

22 Robocopy. 2021. Online. Microsoft. <https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/robocopy>. Accessed 28.9.2021.

23 What is Virtual core and how does it different from Physical core? 2018. Online. Geekboots. <https://www.geekboots.com/story/what-is-virtual-core-and-how-does-it-different-from-physical-core>. Accessed 28.9.2021.

24 Use OneDrive policies to control sync settings. 2021. Online. Microsoft. <https://docs.microsoft.com/en-us/onedrive/use-group-policy>. Accessed 26.9.2021.

25 Group membership changes don't update over some VPN connections. 2021. Online. Microsoft. <https://docs.microsoft.com/en-us/troubleshoot/windows-client/group-policy/group-membership-changes-not-updating-over-some-vpn-connections>. Accessed 2.10.2021.

26 Using Folder Redirection in Group Policy. 2016. Online. Microsoft. <https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-r2-and-2012/dn789199(v=ws.11)>. Accessed 28.9.2021.

27 Accounting policy . 2020. Online. Ministry of Finance Republic of Latvia. <https://www.fm.gov.lv/en/accounting-policy>. Accessed 1.10.2021.

28 The Accounting Act No. 145/1994. 1994. Online. Stjórnarrádid. <https://www.stjornarradid.is/media/atvinnuvegaraduneyti-media/media/Acrobat/The-Accounting-Act-No.-145-1994.pdf>. Accessed 1.10.2021.

29 Kveine, Gunhild. Article: Digital storage of accounting records. 2018. Online. BDO. <https://www.bdo.no/en-gb/insights/digital-storage-of-accounting-records>. Accessed 1.10.2021.

30 Foley, Marry Jo. Microsoft moves closer to running all of its own services on Azure. 2021. Online. ZDNet. <https://www.zdnet.com/article/microsoft-moves-closer-to-running-all-of-its-own-services-on-azure/>. Accessed 30.9.2021.

31 What is managed file transfer? Online. IBM. <https://www.ibm.com/topics/managed-file-transfer>. Accessed 29.9.2021.

32 Managed File Transfer. 2018. Online. Software AG. <https://tech.forums.softwareag.com/t/managed-file-transfer/237787> Accessed 29.9.2021.

33 Deliver a Serverless Printing Infrastructure. Online. PrinterLogic. <https://www.printerlogic.com>. Accessed 4.10.2021.

34 <https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-r2-and-2012/dn581925(v=ws.11)>. Accessed 4.10.2021.

35 PrintManagement. Online. Microsoft. <https://docs.microsoft.com/en-us/powershell/module/printmanagement>. Accessed 4.10.2021.

36 IP Addressing: DHCP Configuration Guide, Cisco IOS Release 15SY. Online. Cisco. <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_dhcp/configuration/15-sy/dhcp-15-sy-book/dhcp-relay-agent.html>. Accessed 25.10.2021.

37 Ipconfig. 2017. Online. Microsoft. <https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/ipconfig>. Accessed 27.9.2021.

38 Virtual Machine to Physical Machine Migration. Online. VMWare. <https://www.vmware.com/support/v2p/doc/V2P_TechNote.pdf>. Accessed 4.10.2021.

39 Workforce management in ONE solution. Online. Tamigo. <https://www.tamigo.co.uk/solution>. Accessed 7.10.2021.

40 Disable or remove the DHCP Server service installed on any domain controllers. 2020. Online. Microsoft. <https://docs.microsoft.com/en-us/services-hub/health/remediation-steps-ad/disable-or-remove-the-dhcp-server-service-installed-on-any-domain-controllers>. Accessed 4.10.2021.

41 Demoting Domain Controllers and Domains. 2018. Online. Microsoft. <https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/deploy/demoting-domain-controllers-and-domains--level-200->. Accessed 7.10.2021.

42 What's New in Active Directory Domain Services Installation and Removal. 2018. Online. Microsoft. <https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/deploy/what-s-new-in-active-directory-domain-services-installation-and-removal>. Accessed 7.10.2021.

43 Clean up Active Directory Domain Controller server metadata. 2018. Online. Microsoft. <https://docs.microsoft.com/en-US/windows-server/identity/ad-ds/deploy/ad-ds-metadata-cleanup>. Accessed 7.10.2021.

44 Tambralli, Kishan. Configuration Management Database (CMDB). 2021. Online. ITIL Docs. <https://www.itil-docs.com/blogs/configuration-management/configuration-management-database-cmdb>. Accessed 6.10.2021.